

Abstract

Title of dissertation: Exploring Equilibrium Systems with Nonequilibrium Simulations

Andrew J. Ballard, Doctor of Philosophy, 2012

Dissertation directed by: Professor Christopher Jarzynski
Department of Chemistry and Biochemistry
and Institute for Physical Science and Technology

Equilibrium sampling is at the core of computational thermodynamics, aiding our understanding of various phenomena in the natural sciences including phase coexistence, molecular solvation, and protein folding. Despite the widespread development of novel sampling strategies over the years, efficient simulation of large complex systems remains a challenge. While the majority of current methods such as simulated tempering, replica exchange, and Monte Carlo methods rely solely on the use of equilibrium techniques, recent results in statistical physics have uncovered the possibility to sample equilibrium states through nonequilibrium simulations.

In our first study we present a new replica exchange sampling strategy, “Replica Exchange with Nonequilibrium Switches,” which uses nonequilibrium simulations to enhance equilibrium sampling. In our method, trial swap configurations between replicas are generated through nonequilibrium switching simulations which act to drive the replicas towards each other in phase space. By means of these switching simulations we can increase an effective overlap between replicas, enhancing the probability that these moves are accepted and ultimately leading to more effective sampling of the underlying energy

landscape. Simulations on model systems reveal that our method can be beneficial in the case of low replica overlap, able to match the efficiency of traditional replica exchange while using fewer processors. We also demonstrate how our method can be applied for the calculation of solvation free energies.

In a second, separate study, we investigate the dynamics leading to the dissociation of Na^+Cl^- in water. Here we employ tools of rare event sampling to deduce the role of the surrounding water molecules in promoting the dissociation of the ion pair. We first study the thermodynamic forces leading to dissociation, finding it to be driven energetically and opposed entropically. In further analysis of the system dynamics, we deduce a) the spatial extent over which solvent fluctuations influence dissociation, b) the role of sterics and electrostatics, and c) the importance of inertia in enhancing the reaction probability.

Exploring Equilibrium Systems with Nonequilibrium Simulations

by

Andrew J. Ballard

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2012

Advisory Committee:

Professor Christopher Jarzynski, Chair/Advisor
Professor Silvina Matysiak
Professor Garegin Papoian
Professor Sergei Sukharev
Professor John D. Weeks

© Copyright by
Andrew J Ballard
2012

To my Family and Friends

Acknowledgments

I begin by thanking my advisor, Professor Christopher Jarzynski, for his guidance during my graduate studies. Chris has taught me the importance of critical thinking and objectivity in scientific research, and I am truly grateful to him for his dedication and ongoing support over the years. I would also like to thank members of my research group, including Suri, Jordan, Dibyendu, Sebastian, and Zhiyue for their stimulating scientific conversations and encouragement.

Secondly, I would like to acknowledge Professor Christoph Dellago and computational physics group at the University of Vienna for hosting me during my Fulbright fellowship. In addition to gaining a deeper understanding and appreciation for the physics of rare events, with the Dellago group I also learned many practical computational skills which will continue to help me in my scientific career. My year in Austria was enriched by the warm welcome given by Christoph and his group members Georg, Philip, Clarion, Jaffar, and Ernesto, who made me feel at home overseas and for whom I am truly grateful.

I would also like to thank all of my friends who have supported me in so many ways over these past years. Thank you to Maschenka and Megan, Justine, Michelle, and Katie for your encouragement. Thank you Ina for the inspiration you gave, and Shell and the Columbia Heights Sangha for your compassion.

Finally, I thank my parents for inspiring me to ask questions and pursue my dreams. Without you I would not be here today.

Funding is gratefully acknowledged from the National Science Foundation grant CHE-0841557, the US Student Fulbright Program, and the University of Maryland.

Table of Contents

List of Tables	vi
List of Figures	vii
List of Abbreviations	x
1 Introduction	1
1.1 Background	1
1.2 Fluctuation theorems and nonequilibrium sampling strategies	11
1.3 Outline of Thesis	13
2 Replica Exchange with Nonequilibrium Switches: theoretical development	16
2.1 Description of method	17
2.2 Derivation: Deterministic dynamics	20
2.3 Derivation: Stochastic dynamics	26
2.4 Velocity rescaling and other escorted switching dynamics	34
2.5 A consistency check for equilibrium sampling	37
2.6 Summary	40
3 Replica Exchange with Nonequilibrium Switches: numerical studies	41
3.1 Using RENS to enhance simulation efficiency	42
3.2 Example 1: 1D system	44
3.3 Example 2: Alanine Dipeptide	52
3.3.1 Metastable states and correlation times	52
3.3.2 Simulation efficiency	57
3.3.3 Switching protocols and CPU costs	60
3.4 Discussion	62
4 Small molecule solvation	66
4.1 Introduction	66
4.2 RENS implementation	72
4.3 Solute of study: Monoethanolamine	73
4.4 Cavity formation	74
4.5 RENS simulation setup	79
4.6 Solvation free energy calculation	80
4.7 Metastable states	82
4.8 Summary	84

5	Dynamics of complex systems: finding the mechanism of ionic dissociation in water	85
5.1	Introduction	85
5.2	Model	88
5.3	Thermodynamics of ionic dissociation	89
5.4	Transition path analysis	93
	5.4.1 Constrained interionic distance r_{ion}	96
	5.4.2 Constrained solvent	97
5.5	Timescales of p_{B} fluctuations	101
5.6	Inertial effects	103
5.7	Discussion and conclusions	105
6	Summary and outlook	108
A	Derivation of RENS for deterministic dynamics with Andersen thermostat	112
B	RENS implementation with the Nosé-Hoover thermostat	117
C	Code development	120
	Bibliography	123

List of Tables

3.1	Empirical CPU times needed for switching trajectories.	59
-----	--	----

List of Figures

1.1	Schematic picture of a rugged (free) energy landscape, with dashed lines representing typical system energies for various temperatures. While large barriers separate metastable states at low temperatures, at higher temperatures the system easily surmounts these barriers, readily exploring the landscape with little difficulty.	4
1.2	Tackling the overlap problem: The shaded circles represent the equilibrium distributions of replicas A and B , which share little phase space overlap. In REM (left panel), sufficient overlap is ensured by adding replicas between A and B , depicted by unfilled circles. With our method (right panel), increased overlap is facilitated by driving A and B out of equilibrium and towards the equilibrium ensembles of replicas B and A , respectively.	9
2.1	Replica Exchange with Nonequilibrium Switches: Between segments of equilibrium sampling (solid red intervals), replicas A and B undergo switching simulations (dashed blue intervals). In this example the trial move $(x, y) \rightarrow (y', x')$ at time t_2 is accepted.	19
2.2	Choosing a sufficiently small timestep: $\Delta L(w)$ is plotted for REM simulations of alanine dipeptide with various timesteps. A flat line indicates Crooks's fluctuation theorem is satisfied (see Eqs. 2.50 and 2.51), suggesting data is sampled from equilibrium.	39
3.1	Mock-up of a rough potential energy landscape, adapted from Ref. [41], Chapter 14. An asterisk marks the fourth well ($x \geq 1.25$). Ordinary replica exchange works well for one particle, but encounters difficulties when $n_p = 10$	45
3.2	Observed occupations of wells 1 - 4, obtained by following a single tagged particle in the output trajectory, for each of the 25 test runs. The solid horizontal lines are exact values determined by integration of the single-particle Boltzmann distribution.	46
3.3	Average reduced work (filled circles) and observed acceptance frequency (open circles), as functions of the fraction of simulation time devoted to switching intervals. $\langle P_{acc} \rangle$ is defined as the fraction of attempted replica exchanges that were accepted, in a given test run. At $f_{sw} = 0$, corresponding to ordinary (instantaneous) replica exchange, $\langle P_{acc} \rangle \approx 0.003$	47
3.4	$n_4(t)$ is plotted over an interval of time, for output trajectories obtained with REM (upper trace, shifted for clarity) and RENS (lower trace). The two traces represent roughly the same number of attempted replica exchanges (≈ 1700), but reveal substantially different acceptance rates.	48

3.5	Sample cost t^* plotted against f_{sw} , both with and without adjustment for increased relative cost of work simulations (see text). The circles identify the run at $\tau = 2.0$. The dashed line is the sample cost of REM with $M = 4$ replicas.	51
3.6	Ramachandran plot of dihedral angles ϕ and ψ for alanine dipeptide at $T_A = 300$ K. The plots were calculated from REM (left panel) and RENS (right panel) simulations, and are shown on a log scale. The boxed region in the left panel indicates conformer C7 _{ax}	53
3.7	Calculated values of t_c , in units of swap attempts, for various values of T_A . The red, green and blue curves correspond, respectively, to switching protocols with $\tau = 0, 10,$ and 100 fs.	55
3.8	Average replica exchange acceptance probability as a function of T_A for various values of the switching time: $\tau = 0$ (filled red circles), 10 fs (open green circles) and 100 fs (filled blue squares). Increasing τ systematically increases $\langle P_{\text{acc}} \rangle$	56
3.9	The CPU cost, t^* , for the simulations of Fig. 3.7.	60
3.10	Discretized switching protocol: The evolution of the system out of equilibrium proceeds via instantaneous changes $\Delta\lambda$ in the control parameter, followed by relaxation steps of duration Δt . Here Δt is much larger than the integration timestep δt	62
4.1	RENS solvation free energy calculations assisted by cavity growth: Replicas progress from the fully noninteracting to fully interacting system through a series of solvent cavities of growing size. Switching on solute-solvent interactions (rightmost pair of replicas) becomes easier in the presence of a large cavity.	71
4.2	Cavity growth simulations: During the switching simulations between replicas i and $i + 1$ the cavity size R is grown / shrunk at a finite rate.	73
4.3	The solvent density profile. <i>Lower Panels:</i> Solvent density profile relative to bulk, as a function of the distance r from the center of the cavity. The noninteracting (lower left panel) and interacting (lower right panel) solute-solvent systems behave differently for small R . For $R > 6\text{\AA}$, however, they show similar behavior. <i>Upper Panels:</i> The external potential $V_{\text{cav}}(r; R)$ (Eq. 4.10) for various cavity sizes R	76
4.4	Replica overlap associated with instantaneous solute insertion and deletion for various cavity sizes. <i>Upper Panel:</i> The average total work and overlap measure C as a function of cavity size R . <i>Lower Panel:</i> Work distributions of conjugate insertion (solid lines) and deletion (dashed lines) switching simulations. Replica overlap increases with R	77
4.5	Calculation of ΔF_{solv} from a thermodynamic cycle with cavity growth.	81
4.6	<i>Left:</i> Free energy cost of cavity growth as a function of cavity size R for replicas in absence of solute-solvent interactions (red) and in presence of solute-solvent interactions (green). <i>Right:</i> Free energy of cavity growth in presence of solute, relative to that in absence of solute	82
4.7	Potential of mean force of MEA along dihedral angle ψ	83

5.1	Thermodynamics of ionic dissociation. The free energy (red) as a function of r_{ion} displays a stable associated state at $r_{\text{ion}} = 2.7 \text{ \AA}$, separated from the dissociated state by a free energy barrier of $5 k_{\text{B}}T$. Also plotted are the average energy (green) and negative entropy (blue) as a function of r_{ion} . Dissociation is driven energetically and opposed by entropy. The inset shows the free energy, the energy and the entropic contribution for an implicit solvent model, in which the electrostatic interaction between the two ions is reduced by a factor $\varepsilon = 80$	90
5.2	<i>Top:</i> Average number $\langle n_{\text{Na}^+} \rangle$ and $\langle n_{\text{Cl}^-} \rangle$ of water molecules in the first solvation shell of Na^+ and Cl^- . Shown in blue is the number $\langle n_s \rangle$ of waters common to the solvation shells of both ions. <i>Bottom:</i> Sum $\langle n \rangle = \langle n_{\text{Na}^+} \rangle + \langle n_{\text{Cl}^-} \rangle$ of the number of water molecules in the first solvation shells of Na^+ and Cl^- and total number of water molecules $\langle m \rangle = \langle n_{\text{Na}^+} \rangle + \langle n_{\text{Cl}^-} \rangle - \langle n_s \rangle$ in the combined first solvation shell.	92
5.3	Entropy S as a function the number $\langle n \rangle = \langle n_{\text{Na}^+} \rangle + \langle n_{\text{Cl}^-} \rangle$ of close contacts between the ions and water molecules in the first solvation shell. The dotted line denotes a linear fit to the data, displaced vertically for better visibility.	94
5.4	Distribution of p_{B} values for equilibrium configurations x restricted to $r_{\text{ion}}(x) = r_{\text{ion}}^*$, corresponding to the top of the free energy barrier shown in Fig. 5.1. The bimodal behavior indicates that r_{ion} alone is not a sufficient reaction coordinate.	97
5.5	Depiction of the first three solvation shells of the ion pair, which were selectively constrained to investigate solvent influence on ionic dissociation.	99
5.6	Committer analysis applied to configurations containing “frozen” (identical) solvent coordinates in (a) the first, (b) first two, and (c) first three solvation shells. In each figure, the colors distinguish between different sets of frozen solvent, chosen near the associated state (red), transition region (green), and dissociated state (blue). Note that the variance of the committer distributions shown in panel (c) is consistent with sharp p_{B} values in the respective sets of configurations and should just be due to the statistical uncertainty in the p_{B} calculations [92].	101
5.7	The effect of water orientations on p_{B} . All configurations within each color contain identical oxygen coordinates but have varying orientations of the water molecules in the first three solvation shells. We analyzed three sets of oxygen positions, chosen near the associated (red), transition (green) and dissociated (blue) states.	102
5.8	Time-dependence of $p_{\text{B}}[x(t)]$, for trajectories generated from dynamics with (a) constrained r_{ion} and (b) constrained r_{ion} and first solvation shell.	103
5.9	Inertial effects near the transition state. The calculated transition path probability p_{TP} is plotted against the committer probability p_{B} for configurations constrained to $r_{\text{ion}} = r_{\text{ion}}^*$. Our results under Hamiltonian dynamics, shown in red, show a deviation of the observed p_{TP} from the analytic result under diffusive dynamics [52]. These inertial effects enhance p_{TP} near the transition state.	105

List of Abbreviations

M	number of replicas / thermodynamic states
x, y	points in configuration or phase space
\bar{x}	conjugate point to x with momentum inversion
H	Hamiltonian
h	reduced Hamiltonian
h_A	reduced Hamiltonian of replica A
\mathcal{R}_A	replica A
p_i^{eq}	equilibrium distribution of replica i
N	system size
MC	Monte Carlo
MD	Molecular Dynamics
REM	replica exchange method
RENS	replica exchange with nonequilibrium switches
λ	control parameter for switching simulations
$\lambda(t)$	switching protocol
P_{acc}	acceptance probability for replica exchange move
τ	switching time
γ	trajectory in configuration or phase space
w_A	reduced work of replica A
q_A	reduced heat of replica A
$\gamma, \tilde{\gamma}$	a pair of conjugate trajectories for switching simulations
v	system dynamics
u	flow field for escorted dynamics
J_A	Jacobian associated with replica A
d	phase / configuration space integration measure
\mathcal{D}	path space integration measure
$P(y', x' x, y)$	transition probability associated with move $(x, y) \rightarrow (y', x')$
M_A	mapping under deterministic equations of motion in replica A
$\pi_A(\gamma)$	probability of generating trajectory γ in replica A conditioned on initial point
$K_{\text{acc}}(y', x' x, y)$	acceptance kernel associated with move $(x, y) \rightarrow (y', x')$
$K_{\text{rej}}(y', x' x, y)$	rejection kernel associated with move $(x, y) \rightarrow (y', x')$
δt	discretization (or simulation) timestep
U	potential energy
$g(\gamma_A, \gamma_B)$	transition probability in path space
\mathbf{p}_i	momentum vector of particle i
\mathbf{q}_i	position vector for of particle i
N_d	number of system degrees of freedom
$\rho_A(w)$	distribution of reduced work values for replica A
Δf	reduced free energy difference
n_p	number of particles
$\bar{\tau}_{\text{eq}}$	average equilibrium sampling duration between switching intervals
f_{sw}	fraction of simulation time devoted to switching simulations
t_c	correlation time of output trajectory
$C(t)$	time correlation function
t^*	simulation sample cost
AD	alanine dipeptide
ϕ, ψ	dihedral angles

$C7_{ax}$	metastable state of alanine dipeptide
$Y(x)$	indicator function for state $C7_{ax}$
T_{sw}^{CPU}	CPU time associated with switching segment
T_{eq}^{CPU}	CPU time associated with equilibrium segments
Δt	relaxation timestep for switching protocol
$\Delta \lambda$	increment interval of λ
U_u	solute-solute interaction energy
U_v	solvent-solvent interaction energy
V_{int}	solute-solvent interaction energy
ΔF_{solv}	solvation free energy
$\Delta F_{solv}(x)$	solvation free energy in frozen solvent configuration x
QCT	quasichemical theory
V_{cav}	solvent cavity potential
R	size of solvent cavity
$w_{i \rightarrow i+1}$	reduced work performed in replica i to $i + 1$
$\lambda_{i \rightarrow i+1}$	switching protocol between replicas i and $i + 1$
MEA	monoethanolamine
V	simulation volume
$\rho(r)$	solvent density distance r from solute center
C	Bennett overlap measure
ΔF_{cav}	free energy associated with cavity growth in presence of solute
ΔF_{cav}^0	free energy associated with cavity growth in absence of solute
ΔF_{int}	free energy associated with solute insertion
r_{ion}	ion-ion separation
$F(r_{ion})$	free energy as a function of interionic distance
$U(r_{ion})$	average potential energy as a function of interionic distance
$S(r_{ion})$	entropy as a function of interionic distance
n_{Na}	hydration number in the first solvation shell of Na^+ ion
n_{Cl}	hydration number in the first solvation shell of Cl^+ ion
$p_B(x)$	committor probability of configuration x
$p_{TP}(x)$	transition path probability of configuration x
$r(x)$	reaction coordinate
$P(p_B)$	distribution of committor values
m_i	mass of particle i
ζ	Nosè-Hoover bath variable
K	system kinetic energy

Introduction

1.1 Background

Over the past half century or so, computer simulations have become an increasingly indispensable tool for basic scientific research. Our knowledge of the natural world has been clarified by the ability to solve equations of motion numerically, providing information about the behavior of a system in atomistic detail. Since the introduction of the Metropolis Monte Carlo method [80] in 1953, molecular simulations have advanced our understanding of various systems in physics, chemistry, and biology for which theories are only approximate, including liquids and liquid structure [125], the nature of protein folding [31, 110], glassy systems [118], critical phenomena [16], and hydrophobicity [22] to name a few.

The systems that are amenable to study through simulation are often large and complex in nature, evolving in a high-dimensional phase space in which rugged energy landscapes give rise to multiple, distinct conformationally-stable states separated by high (free) energetic barriers. An ion channel, for instance, can exist in “open” and “closed” states, with transitions between the stable states involving a complicated collective reori-

entation of entire protein substructures. The process of protein folding from an unfolded state to a native structure can proceed via Markov-like transitions between metastable intermediaries, and through multiple pathways [86]. Even an ion pair in solution contains stable associated and dissociated states, with transitions promoted by nontrivial reorganization of the solvent molecules in the vicinity of the ions [43]. Simulations have helped to understand structural properties of such systems, for example the classification of metastable states, as well as dynamical properties like reaction rate constants and reaction coordinate determination. The inherent complexity of such systems is often what makes them so interesting to study.

Despite the success of simulations in describing various physical phenomena, the effective simulation of complex systems proves to be challenging. Due to their rugged energy landscapes, complex systems tend to remain trapped in locals well for long times. In many biomolecules, for instance, transitions between metastable states can occur in times up to 100 ns or more, orders of magnitude larger than timescales which are accessible in standard Molecular Dynamics simulations. As a result, estimation of equilibrium properties can be inaccurate at best. Complications also arise for large systems, as integration times typically grow faster than the system size.

Advances in computational power and resources have partially been able to address this issue. Perhaps most impressive example is Anton [109], a supercomputer hard-wired for Molecular Dynamics simulations produced by D.E. Research, which is capable of orders of magnitude speedup with respect to standard machines. The Folding@home project¹, spearheaded by the Pande lab at Stanford, has also pushed the limit of biomolecular simulation, and is able to accomplish computations of similar scale in a parallelized fashion.

¹folding.stanford.edu

Here users throughout the globe donate a fraction of their computational resources to a small bit of a large calculation. Despite these pioneering efforts, the vast majority of scientific computing resources are unable to handle such large and lengthy computations.

To address this problem from a theoretical point of view, much research has been devoted to the development of novel simulation methods that are effective at sampling complex energy landscapes. While the complete understanding of a physical system requires knowledge of static as well as dynamical properties, most simulation methods are only able to retrieve one at the expense of the other. It is then useful to distinguish between methods that are designed to obtain static thermodynamic information, i.e. quantities that can be written as ensemble averages, such as metastable state probabilities, specific heats, and free energy differences, and methods for calculating dynamic quantities, such as time correlation functions, reaction rates and reaction coordinates. Trajectory-sampling methods such as transition path sampling [17] can be very efficient at sampling rare events, and preserve dynamical information needed for the determination of reaction coordinates and transition pathways. However, they only sample a subset of the canonical distribution pertaining to transition path segments. Most thermodynamic sampling methods, on the other hand, are able to achieve enhanced sampling of distributions by introducing collective moves in Monte Carlo simulations, or applying configurational swap moves between trajectories running in parallel. Although these methods efficiently sample the desired distribution, because the employed moves alter the system dynamics, they do not preserve the physical dynamics pertaining to the underlying system. In this thesis we will investigate both thermodynamic and dynamical aspects of systems in equilibrium. The first and largest part will be concerned with the development of a sampling method for accurate determination of static equilibrium quantities. In the second part, where we

investigate the dissociation of $\text{Na}^+\text{-Cl}^-$ in water, the dynamics are of central interest; this project is self contained within its own chapter. In preparation for the introduction of our sampling method, we now focus on thermodynamic sampling methods.

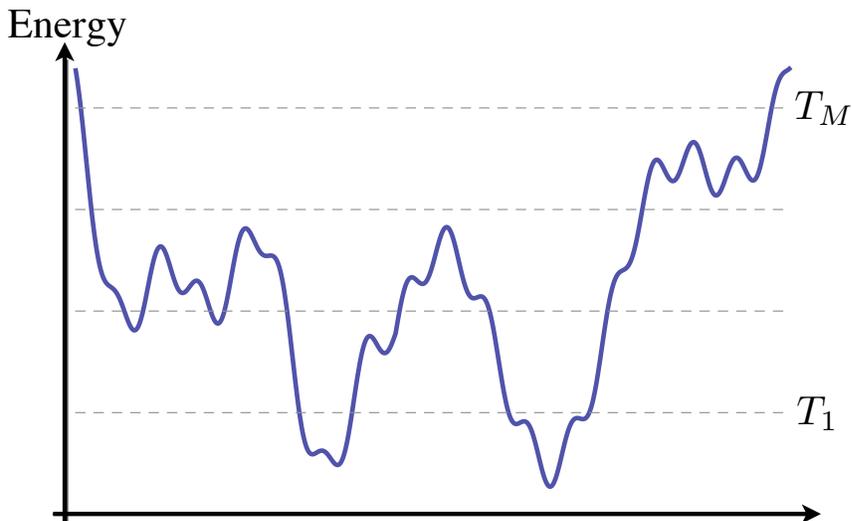


Figure 1.1: Schematic picture of a rugged (free) energy landscape, with dashed lines representing typical system energies for various temperatures. While large barriers separate metastable states at low temperatures, at higher temperatures the system easily surmounts these barriers, readily exploring the landscape with little difficulty.

Many novel methods have been developed for the efficient sampling of equilibrium distributions of complex systems. Out of this has emerged a general class of extended state methods in which *multiple* thermodynamic states of a given system are simulated in tandem, in order to enhance the sampling of one (or more) of them. In general one considers M thermodynamic states, typically of the same system but at different temperatures $T_1 < T_2 < \dots < T_M$. These methods prescribe Monte Carlo-like moves that are attempted between these thermodynamic states. The key to these methods is the realization that, while at a low temperature T_1 the system may be stuck in a given configuration, at higher

temperatures the enhanced thermal fluctuations allow for easier exploration of its configuration space (see Fig. 1.1). By means of these Monte Carlo moves, lower temperature states are exposed to the broader distributions sampled by the higher temperature states.

These multistate methods have proven to be powerful and flexible tools for sampling complex systems, and can be implemented in a variety of serial and parallel implementations. Before proceeding to a discussion of a parallel version, which will be the main focus of this thesis, it is useful to briefly introduce a serial implementation, Simulated Tempering (ST). In the Simulated Tempering [73,75] method, a single copy of a given system samples M distinct thermodynamic ensembles at temperatures T_1, \dots, T_M . Here the system evolves under Molecular Dynamics (or other dynamics) at a given temperature and for a specified amount of time. Periodically this evolution is interrupted by moves that attempt to alter the temperature $T_i \rightarrow T_j$. If the move is accepted, the system continues MD sampling but now at the new temperature T_j ; if the move is rejected, MD sampling resumes at T_i . In this way, a single trajectory performs a random walk along a ladder of temperatures in addition to sampling equilibrium at each fixed temperature. Importantly, these Monte Carlo-like moves between temperatures are accepted with a detailed balance-preserving condition, guaranteeing that sampling within a given temperature preserves the respective equilibrium distribution. The gain in the method comes from the fact that, while at lower temperatures the system may be trapped in a local well, at higher temperatures it can readily cross barriers and explore other regions of phase space. Other variants allow for switching of parameters other than temperature, such as pressure [91] or terms in the Hamiltonian [70].

While ST has been successful at sampling complex systems, its parallel equivalent, the Replica Exchange Method [33,41] (REM), is arguably the more commonly used of the

thermodynamic parameters. In its original development due to Swendsen and Wang [116], REM was used in a temperature implementation with MC dynamics for efficient simulation of spin glass systems. Since its development with Molecular Dynamics [82], it has been a common and powerful tool for simulation of complex systems in chemistry and physics, especially biomolecular systems [44].

There is an advantage to using REM over serial methods like ST: For an efficient ST implementation, weights must be assigned to the individual ensembles to ensure roughly equal sampling of each thermodynamic state. Because the optimal weight corresponds to the free energy difference between two neighboring ensembles, estimation of effective weights can be quite expensive computationally. With replica exchange, however, these weights are unnecessary since all ensembles are sampled equally by construction. However, provided one has the optimal weights, the serial algorithm always accepts more moves on average [91].

While REM is in principle valid for any sequence of replicas (T_1, \dots, T_M) , an effective implementation demands that a decent fraction of exchange attempts are accepted, which can only be achieved when neighboring replicas share a significant overlap in phase space [64]. If two replicas share little overlap, most swap attempts will be rejected simply because the swap configurations are not representative of the equilibrium ensembles they are swapped into. Consider, for example, an implementation with replicas A and B , where the equilibrium distribution of replica A at temperature T_A is largely confined to locally-metastable regions. If $T_B \gg T_A$, replica B will diffuse through an effectively uniform energy landscape, sampling a very broad distribution and spending very little time in local wells which A is confined to. In this case of low overlap, most swaps will be rejected, simply because most configurations sampled in replica B lie outside the local

wells which contribute most to the equilibrium of A . To prevent such a situation, the spacings $\Delta T = T_B - T_A$ between neighboring replicas must be chosen small enough so that decent overlap is achieved.

This overlap requirement ultimately leads to a replica scaling relation: the number of replicas needed grows with the system size N as \sqrt{N} [51]. This is attributed to the fact that the average system energy scales like N , yet the width of the distribution scales like \sqrt{N} . In order to maintain a certain amount of overlap between energy distributions (which determines the acceptance rate), $\mathcal{O}(\sqrt{N})$ replicas must be added to fill this gap. This scaling relation can be a particularly difficult reality for e.g. simulations of large biomolecules in explicit solvent, which usually involves thousands of degrees of freedom.

This overlap problem has been addressed both through semi-analytic [64, 103, 104] analysis of REM as well as novel REM methods [21, 34, 36, 50, 71, 82, 101, 124]. For example, research has been devoted to finding optimal temperature spacings of parallel tempering simulations which guarantee uniform acceptance rates across all replica pairs. Kofke has found that for systems with roughly-constant heat capacity, replica temperatures should follow a geometric progression [64]. Adaptive schemes for determination of optimal replica spacings have also been developed [99, 106]. Many REM methods have been proposed where overlap is enhanced by sampling from Tsallis-like distributions [47, 54, 126], or in combination with multicanonical [21, 34, 82] and expanded ensembles [36, 101]. Berne and coworkers [50, 71, 124] have also developed Replica Exchange with Solute Tempering, a version of parallel tempering in which only a subset of the system is tempered. This strategy, particularly useful for solutes in explicit solvent, involves a deformation of the system Hamiltonian such that the effective temperature of the solvent molecules is constant across all replicas. By tempering only the solute degrees of freedom, the acceptance

criteria becomes independent of the solvent-solvent energy contributions and hence also independent of the number of solvent molecules.

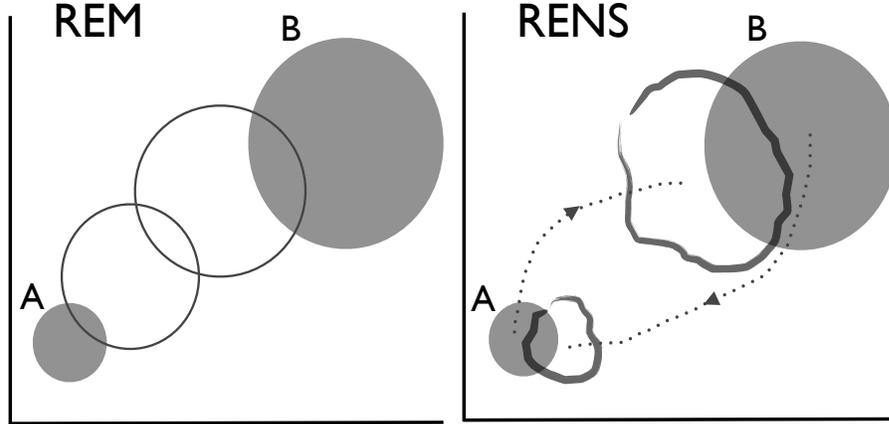
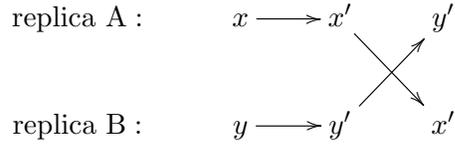


Figure 1.2: Tackling the overlap problem: The shaded circles represent the equilibrium distributions of replicas A and B , which share little phase space overlap. In REM (left panel), sufficient overlap is ensured by adding replicas between A and B , depicted by unfilled circles. With our method (right panel), increased overlap is facilitated by driving A and B out of equilibrium and towards the equilibrium ensembles of replicas B and A , respectively.

In this thesis we develop an alternative strategy, Replica Exchange with Nonequilibrium Switches (RENS), which uses nonequilibrium simulations to increase replica overlap [7, 8]. Instead of performing instantaneous exchanges between replicas A and B , as is the case in REM, in our method the configurations to be swapped are first generated from finite-time “switching simulations,” during which the system evolves as the temperature of the thermostat (and/or the system Hamiltonian) is parametrically switched between its values corresponding to the equilibrium distributions of A and B . At the culmination of a switching simulation, an attempt is made to exchange the final points generated in A

and B .



Although the replicas are driven out of equilibrium during the switching simulations, by accepting the move with the probability

$$P_{\text{acc}} = \min\{1, e^{-w_A - w_B}\} \quad (1.2)$$

we guarantee that equilibrium in each replica remains undisturbed. The quantities w_A and w_B represent generalized reduced work values in replicas A and B , whose particular form depends upon the dynamics used to model the system (see e.g. Eqs. 2.5 and 2.28 of Chapter 2). By means of the switching simulations, the system in replica A is brought nearer to the region of phase space associated with replica B , and vice versa (see Fig. 1.2); at the same time, the duration of the switching simulation is a parameter set by the user which gives direct control over this overlap. While each replica incurs an additional CPU cost due to the switching simulations, this comes in exchange for an increased exchange acceptance rate. Related replica exchange methods have also incorporated nonequilibrium moves: The annealed swapping [89] and C-Walking [18] methods, for instance, apply heating and cooling protocols to generate trial swap configurations. Other methodologies [6, 30, 83–85, 113] have also incorporated nonequilibrium simulations to enhance equilibrium sampling, including serial multistate methods [30, 85].

Our method, which we will describe and derive in detail in the coming chapters, relies upon recent theoretical results in nonequilibrium statistical physics that have made deep connections between systems driven out of equilibrium and their corresponding equilibrium states [59]. In order to understand our method in detail, it is important to first introduce

two key results, Crooks’s fluctuation theorem (Eq. 1.3 below) and the nonequilibrium work relation (Eq. 1.5 below), which will be exploited in later chapters.

1.2 Fluctuation theorems and nonequilibrium sampling strategies

Over the past 15 years or so, it has become clear that equilibrium information is encoded in systems driven out of equilibrium. Various theoretical results have shown that it is possible to extract equilibrium quantities such as free energy [55, 56] and entropy differences [57] and equilibrium statistics [53, 55] from systems driven out of equilibrium. Jarzynski [55], and Hummer and Szabo [53], for instance, showed that equilibrium distributions can be recovered from systems driven out of equilibrium via proper reweighting of nonequilibrium trajectories based upon the work performed. This can be utilized: In experiments, for example, one can calculate the equilibrium potential of mean force of a DNA hairpin as a function of its end-to-end distance from a series of force spectroscopy experiments [46] in which the hairpin is repeatedly pulled apart at a finite rate. In computations, these ideas have also gained traction, and recently methods have incorporated nonequilibrium simulations to enhance sampling of equilibrium states. Our RENS method is one such example, where replica exchange trial moves are generated by driving the replicas out of equilibrium.

The ability to extract equilibrium information from systems out of equilibrium can be understood in terms of a deep symmetry between conjugate processes driven out of equilibrium. This symmetry relation, first shown by Crooks [26–28], provides an important tool for development of equilibrium sampling methods which incorporate nonequilibrium trajectories. To explain this, we consider a system initially in equilibrium at a temperature T , which is driven out of equilibrium by variation of an external control parameter λ . We

denote the *forward process* to be one in which the system, initiated in equilibrium at $\lambda = A$, is driven out of equilibrium by variation of λ from $\lambda = A$ to $\lambda = B$ in a time τ . In the *reverse process* let the system be prepared in equilibrium state B , and driven out of equilibrium through the time-reversed protocol $\tilde{\lambda}(t) = \lambda(\tau - t)$ from $\lambda = B$ to $\lambda = A$. (Within the context of our RENS simulations, the forward process could denote the heating of the system in replica A from T_A to T_B , and the reverse process the cooling in replica B from T_B to T_A .) Crooks's fluctuation theorem relates the probability of observing a trajectory γ in the forward process to the probability of observing its time-reversed trajectory $\tilde{\gamma}$ in the reverse process:

$$\frac{\mathcal{P}_F[\gamma]}{\mathcal{P}_R[\tilde{\gamma}]} = e^{(W[\gamma] - \Delta F)/k_B T} \quad (1.3)$$

Eq. 1.3 quantifies the time asymmetry arising from such irreversible processes in terms of the work dissipated into the environment, $W - \Delta F$, and can be viewed as an extension of detailed balance to systems driven out of equilibrium. As such, this fluctuation theorem is ultimately at the heart of the validity of nonequilibrium sampling methods such as RENS and will be exploited in the coming chapters.

As Eq. 1.3 suggests, it is also possible to recover free energy differences ΔF from systems driven out of equilibrium. It has long been known that the average work W performed during an irreversible process is bounded from below by the free energy difference $\Delta F = F_B - F_A$ between the two equilibrium states A and B :

$$\langle W \rangle \geq \Delta F, \quad (1.4)$$

This is a statement of the second law of thermodynamics [37], and the equality is reached when the process is performed reversibly and isothermally. In fact, this inequality follows

from a stronger equality, first shown by Jarzynski in 1997 [56]:

$$\langle e^{-W/k_B T} \rangle = e^{-\Delta F/k_B T}. \quad (1.5)$$

This nonequilibrium work relation, which can be seen as a consequence of Eq. 1.3, states that the free energy difference between two equilibrium states can be obtained from an ensemble of work values performed by driving that system out of equilibrium. The average in Eq. 1.5 (and Eq. 1.4) is taken over infinitely many repetitions of the forward process described above in which the system is initially prepared in equilibrium at $\lambda = A$ and driven out of equilibrium through varying λ from A to B . In addition to the physical insight it brings, the nonequilibrium work relation is useful in a practical sense, as it provides a free energy estimator: For a simulation with n such trajectories, the free energy can be estimated as

$$e^{-\Delta F/k_B T} \approx \frac{1}{n} \sum_i e^{-W_i/k_B T}, \quad (1.6)$$

where W_i is the work value obtained during trajectory i and the approximation becomes an equality in the limit $n \rightarrow \infty$. Since the introduction of the nonequilibrium work relation, Eq. 1.5 (and variants thereof) has been used to estimate free energy differences from computer simulations. This will be used in Chapter 4 for calculation of solvation free energies.

1.3 Outline of Thesis

This thesis describes two distinct research projects. The first project, presented in Chapters 2-4, is the development of RENS, our replica exchange method which implements nonequilibrium switching simulations. In Chapter 2 we focus on the theoretical development of RENS, describing and deriving the method in detail. Because the particular implementation of RENS depends on the dynamics used to model the system, we

present separate derivations for deterministic and stochastic equations of motion. We then describe novel dynamical schemes which can be useful for enhancing simulation efficiency, and present a consistency check for equilibrium sampling. In Chapter 3 we test RENS on two model systems: a simple one-dimensional system of particles, and alanine dipeptide, with the goal of finding the regimes in which it is useful to use RENS over REM. To this end we quantify the gains in sampling quality and efficiency for RENS as a function of the simulation time devoted to the nonequilibrium segments, a parameter set by the user. Our findings confirm that our method can be beneficial when the replicas share little overlap. We proceed in Chapter 4 to apply RENS to a different problem, the calculation of the solvation free energy of a small solute. Whereas in the previous examples our replicas were defined by different temperatures, in this study the free energy of solvation is calculated from a sequence of replicas with unique Hamiltonians. Here the two end replicas sample the fully interacting and noninteracting solute-solvent system, and are connected through a sequence of intermediately-solvated replicas which contain a solvent cavity surrounding the solute. We use nonequilibrium-based free energy estimators (Eq. 1.5) to calculate the solvation free energy of the small organic molecule monoethanolamine.

In the second project, discussed in Chapter 5, we shift gears and focus on rare event sampling and dynamics leading to transitions between two metastable states. In particular, we investigate the dynamics which promote the dissociation of Na^+ and Cl^- in water. It is known that the water molecules surrounding the ion pair play a non-trivial role during the dissociation process [43]. To further elucidate this mechanism, we first show through thermodynamic analysis that dissociation is driven energetically, and opposed entropically, where the loss of entropy is explained by an increasing number of solvent molecules entering the highly-coordinated solvation shell. By using techniques developed within rare

event sampling fields, we then investigate the dynamical aspects leading to dissociation, including characterizing the spatial extent to which water influences this reaction, estimating timescales associated with solvent rearrangements leading to dissociation, as well as the effect of inertia near the transition state.

We note that Chapters 2, 3, and 5 of this thesis are based upon the following papers:

1. A.J. Ballard and C. Jarzynski. “Replica exchange with nonequilibrium switches,” *Proc. Natl. Acad. Sci. U.S.A.*, 106:1222412229, 2009.
2. A.J. Ballard and C. Jarzynski. “Replica exchange with nonequilibrium switches: Enhancing equilibrium sampling by increasing replica overlap”, *J. Chem. Phys.*, 136:194101, 2012.
3. A.J. Ballard and C. Dellago, “Towards the mechanism of ionic dissociation in water”, *in revision*.

Replica Exchange with Nonequilibrium

Switches: theoretical development

¹In the previous chapter we introduced the replica exchange method (REM), a strategy to enhance sampling of complex systems. An effective replica exchange implementation requires a phase space overlap in order to sample effectively. Within REM this overlap is satisfied by adjusting the number of replicas in a given simulation. Our method of Replica Exchange with Nonequilibrium Switches (RENS), which we briefly introduced last chapter and describe in detail below, achieves an increased overlap through nonequilibrium simulations which drive the replicas closer to one another in phase space. In doing so, we forgo the need for many replicas at the expense of additional simulation time devoted to the nonequilibrium simulations in each replica. In this chapter we present the theoretical framework for RENS, describing its implementation and deriving its validity

¹This chapter is based on the papers “Replica exchange with nonequilibrium switches“, A. J. Ballard and C. Jarzynski, *Proc. Natl. Acad. Sci. U.S.A.*, 106 (30):12224 (2009) and “Replica Exchange with Nonequilibrium Switches: Enhancing Equilibrium Sampling by Increasing Replica Overlap“, A. J. Ballard and C. Jarzynski, *J. Chem. Phys.*, 136 194101 (2012).

for various dynamical schemes.

2.1 Description of method

To begin, consider a collection of M replicas of a physical system, labeled $\mathcal{R}_1, \dots, \mathcal{R}_M$. The thermodynamic state of \mathcal{R}_i is described by Hamiltonian H_i and temperature T_i , allowing for Hamiltonian and/or temperature replica exchange. As in Ref. [91], we define the reduced Hamiltonian, $h_i(x) \equiv H_i(x)/k_B T_i$, where x (and later y) denotes a point in the phase space of the system. It is convenient to view the set of M replicas as a single, extended system, characterized by a reduced Hamiltonian $\mathcal{H}(x_1, \dots, x_M) = \sum_{i=1}^M h_i(x_i)$. When each \mathcal{R}_i is sampled according to its equilibrium distribution $p_i^{\text{eq}}(x) \propto \exp[-h_i(x)]$, then the extended system samples the distribution $\mathcal{P}^{\text{eq}} \propto \exp[-\mathcal{H}]$. Both the traditional replica exchange method (REM) and our approach (RENS) represent strategies for sampling the extended distribution \mathcal{P}^{eq} .

Imagine first that each \mathcal{R}_i evolves under the equilibrium dynamics corresponding to h_i . Replica exchange supplements this independent evolution via swap attempts between neighboring replicas: If replicas A and B exist in x and y , respectively, at the time of an attempted swap, under REM the move $(x, y) \rightarrow (y, x)$ represents a trial replica exchange. To ensure detailed balance, this move is accepted with a probability $P_{\text{acc}} = \min\{1, \exp[-\Delta h]\}$, where $\Delta h = h_A(y) - h_A(x) + h_B(x) - h_B(y)$ is the total change in reduced energy for the replica pair associated with the trial move under consideration.

Our method, RENS, is illustrated in Fig. 2.1. From time t_0 to t_1 , replicas A and B undergo independent equilibrium sampling at h_A and h_B , respectively (solid red lines). At t_1 , a decision is made to attempt an exchange. Whereas in REM one would immediately attempt to swap the points x and y , our method prescribes the generation of new swap

points by first performing *switching simulations* (dashed blue lines): From t_1 to t_2 replica A evolves as the Hamiltonian is parametrically switched between h_A and h_B in time τ , generating a path γ_A and arriving at final point x' . Similarly, the switching simulation in replica B generates trajectory γ_B by switching from h_B to h_A , reaching a final point y' . Upon completion of the switching simulations a swap is attempted between these final points, represented by the trial move $(x, y) \rightarrow (y', x')$. The move is accepted with probability

$$P_{\text{acc}} = \min\{1, e^{-w_A - w_B}\}. \quad (2.1)$$

with w_A and w_B the reduced work performed during the switching simulations in A and B . If accepted, x' is copied into B and y' into A ; if the move is rejected, x and y are copied back into A and B , respectively, and the momenta are inverted². Whether the move is accepted or rejected, at t_2 both replicas resume their equilibrium dynamics (at fixed h) and continue sampling until the next exchange attempt. Note that while switching simulations are performed in \mathcal{R}_A and \mathcal{R}_B , the remaining replicas continue to sample at fixed h_i .

Thus, each replica alternates between *sampling intervals* at fixed h_i (Fig. 2.1, solid red lines), and *switching intervals* (dashed blue lines). In the next section will show that RENS satisfies detailed balance, in the following sense: in each replica \mathcal{R}_i , if we discard the data generated during the switching intervals, and stitch together the remaining sampling intervals, we obtain a long trajectory that samples the distribution p_i^{eq} . In effect, the acceptance criterion compensates for the fact that the system is driven out of equilibrium during the switching simulations.

²The momentum flip upon rejection also appears in e.g. the generalized hybrid Monte Carlo [69] and NCMC [85] methods, and ensures equilibrium sampling.

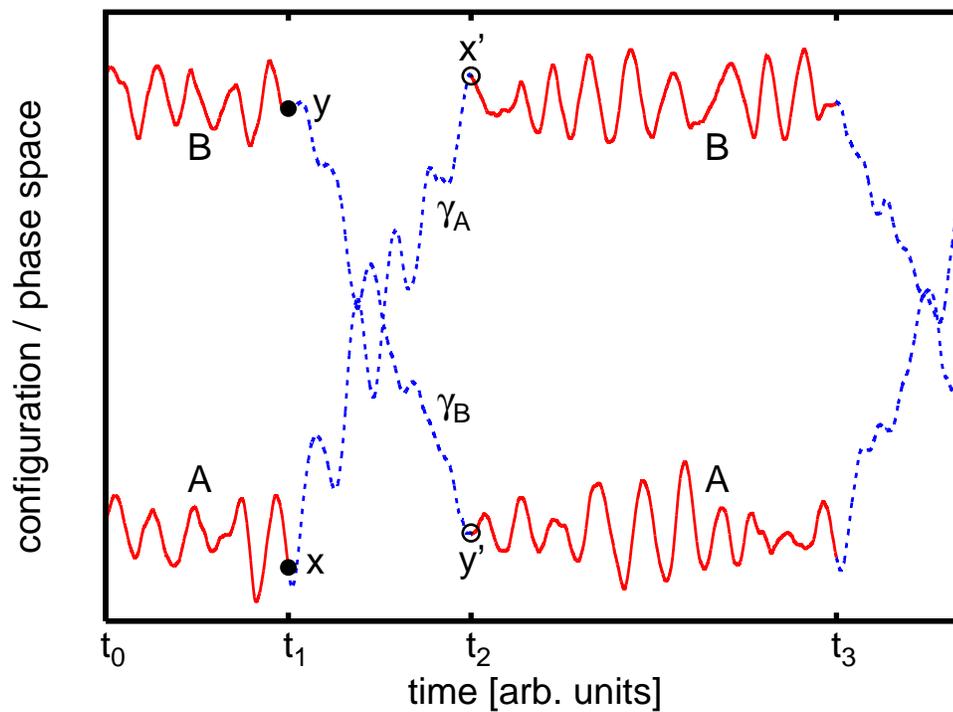


Figure 2.1: Replica Exchange with Nonequilibrium Switches: Between segments of equilibrium sampling (solid red intervals), replicas A and B undergo switching simulations (dashed blue intervals). In this example the trial move $(x, y) \rightarrow (y', x')$ at time t_2 is accepted.

The acceptance probability of a given trial move (Eq. 2.1 above) is determined by the reduced work performed during the switching simulations. While the particular form of w_A and w_B will in general depend upon the dynamics used to model the system (compare Eqs. 2.5 and 2.28 below), the reduced work can generically be thought of as the change in the system’s reduced energy due to parametrically switching its Hamiltonian between h_A and h_B through control parameter λ :

$$w = \int_0^\tau dt \dot{\lambda} \frac{\partial h}{\partial \lambda} \quad (2.2)$$

The specific form the work takes is described in more detail in the derivations below.

Our method is quite general, and ultimately traces its validity to Crooks’s extension of detailed balance to nonequilibrium trajectories [27]. In a given implementation, however, the definition of reduced work depends on the dynamics chosen to model the evolution of the system. For discrete-time Monte Carlo dynamics (MC), RENS is mathematically equivalent to the annealed swapping method of Opps and Schofield [89] and related to the C-walking algorithm of Brown and Head-Gordon [18], although the physical interpretation of RENS differs from that of Refs. [18, 89]. In the following sections we derive our method for a) deterministic, reversible molecular dynamics, and b) stochastic dynamics. In Appendix A we present a third derivation for dynamics which combine deterministic and stochastic evolution (e.g. Molecular Dynamics in conjunction with an Andersen thermostat).

2.2 Derivation: Deterministic dynamics

In this section we derive RENS for continuous-time, deterministic dynamics. Specifically, we introduce a parameter-dependent reduced Hamiltonian $h(x; \lambda)$ (see below), and will be considering system evolution under deterministic equations of motion which are

typically of the form

$$\frac{dx}{dt} = v(x; \lambda). \quad (2.3)$$

The dynamics represented by Eq. 2.3 are valid both when the parameter λ is fixed and when it changes with time, as represented respectively by the solid red and dashed blue lines in Fig. 2.1. As discussed in detail below, we will further assume that these dynamics are symmetric under time-reversal³.

To begin our derivation for deterministic systems we consider a pair of replicas \mathcal{R}_A and \mathcal{R}_B , and introduce a parametrized Hamiltonian $h(x; \lambda)$ that interpolates from $h(x; 0) = h_A(x)$ to $h(x; 1) = h_B(x)$. To implement an attempted swap between these replicas, we first specify a *switching protocol* $\lambda_A(t)$, with $\lambda_A(0) = 0$ and $\lambda_A(\tau) = 1$. In \mathcal{R}_A , starting from state $x_0 = x$, we generate a trajectory γ_A during which the system evolves under the specified dynamics as the parameter λ is varied from 0 to 1 according to the switching protocol:

$$\gamma_A : x = x_0 \xrightarrow{\lambda \rightarrow 1} x_\tau = x'. \quad (2.4a)$$

Simultaneously, in \mathcal{R}_B we generate a trajectory γ_B by varying λ from 1 to 0 under the time-reversed protocol, $\lambda_B(t) = \lambda_A(\tau - t)$,

$$\gamma_B : y' = y_\tau \xleftarrow{0 \leftarrow \lambda} y_0 = y. \quad (2.4b)$$

Here the arrows denote evolution under the deterministic equations of motion, and the arrowheads denote the direction of time. The two trajectories γ_A and γ_B are illustrated by the dashed blue segments between times t_1 and t_2 in Fig. 2.1. We assume these

³Eq. 2.3 represents deterministic equations of motion for physically-relevant evolution such as Hamilton's equations and Nosè-Hoover dynamics, but the particular form represented by Eq. 2.3 is not assumed in this derivation. The derivation to follow is also valid for generalizations of these dynamics (such as Eq. 2.44 in Sec. 2.4 below), as long as the evolution is deterministic and symmetric under time reversal.

trajectories are generated by deterministic equations of motion that are symmetric under time-reversal; for any trajectory $\gamma_A = (x_0 \rightarrow x_\tau)$ that is a solution of Eq. 2.3 under the protocol $\lambda_A(t)$, the time-reversed trajectory $\tilde{\gamma}_B = (\bar{x}_0 \leftarrow \bar{x}_\tau)$ is a solution under $\lambda_B(t)$, where \bar{x} denotes inversion of momenta, $\mathbf{p} \rightarrow -\mathbf{p}$. This assumption is satisfied by Hamiltonian, Nosé-Hoover [48, 87], and other dynamics, provided the Hamiltonian itself is time-reversal symmetric, i.e. $h(x; \lambda) = h(\bar{x}; \lambda)$.

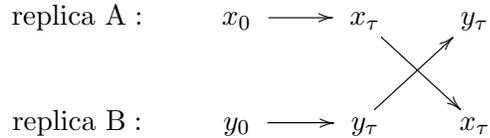
For the trajectories γ_A and γ_B generated under these deterministic dynamics, we define the reduced work as follows:

$$w_A(x_0 \rightarrow x_\tau) = h_B(x_\tau) - h_A(x_0) - \ln J_A(x_0) \quad (2.5a)$$

$$w_B(y_0 \rightarrow y_\tau) = h_A(y_\tau) - h_B(y_0) - \ln J_B(y_0). \quad (2.5b)$$

Here $J_A = |\partial x_\tau / \partial x_0|$ and $J_B = |\partial y_\tau / \partial y_0|$ are the Jacobians associated with propagating the system from the initial to the final point. Eq. 2.5 is analogous to the first law of thermodynamics, with $\ln J$ representing a heat term associated with increase of system entropy.

When both switching simulations have been completed, we attempt to swap the *final* configurations, assigning $x' = x_\tau$ to replica *B* and $y' = y_\tau$ to replica *A*. Schematically,



where the parallel arrows indicate the switching simulations, and the crossed arrows the attempted swap. The acceptance probability for the swap is given by Eq. 2.1. If it is rejected, the replicas are reset to the time-reverse of their initial states, $\bar{x} = \bar{x}_0$ and $\bar{y} = \bar{y}_0$.

To analyze our method, it is useful to think of an expanded phase space containing

two copies of the system, corresponding to \mathcal{R}_A and \mathcal{R}_B . In this space, we wish to sample the distribution $p_{AB}^{\text{eq}}(x, y) \propto e^{-h_A(x)-h_B(y)}$. The switching simulations, followed by the attempted swap, represent an elaborate trial Monte Carlo move $(x, y) \rightarrow (y', x')$ [113]. In what follows we will show that the method satisfies balance: The probability flux into state (y', x') , from all possible states (x, y) , is such that the equilibrium distribution is preserved. We write this balance condition as follows:

$$\int dx \int dy P(y', x'|x, y) p_{AB}^{\text{eq}}(x, y) = p_{AB}^{\text{eq}}(y', x'). \quad (2.6)$$

Here $P(y', x'|x, y)$ is the net transition probability associated with the move $(x, y) \rightarrow (y', x')$, including contributions from both accepted and rejected moves. To establish balance, we first show that our scheme satisfies a *detailed balance* condition for accepted moves (Eq. 2.15), which ultimately implies Eq. 2.6 and the validity of our method.

For the switching simulations in \mathcal{R}_A and \mathcal{R}_B , because our dynamics are deterministic we can treat the final microstate as a function of the initial microstate [88],

$$x_\tau = M_A(x_0) \quad , \quad y_\tau = M_B(y_0), \quad (2.7)$$

obtained by integrating the equations of motion. We will use the notation

$$\pi_A(x'|x) = \delta(x' - M_A(x)) \quad (2.8a)$$

to denote the probability to arrive at x' during a switching simulation in \mathcal{R}_A , starting from x ; and similarly

$$\pi_B(y'|y) = \delta(y' - M_B(y)). \quad (2.8b)$$

Moreover, let

$$\pi(x', y'|x, y) = \pi_A(x'|x) \pi_B(y'|y) \quad (2.9)$$

denote the joint probability for both events; and let

$$\alpha(x', y'|x, y) = \min\{1, e^{-w_A(x \rightarrow x') - w_B(y \rightarrow y')}\} \quad (2.10)$$

be the probability to accept the corresponding swap (Eq. 2.1). Finally, define

$$K_{acc}(y', x'|x, y) = \pi(y', x'|x, y) \alpha(y', x'|x, y) \quad (2.11)$$

K_{acc} is the joint probability of generating configurations (y', x') and accepting the swap, given initial points (x, y) .

The functions M_A and M_B introduced in Eq. 2.7 are related by our assumption of time-reversal symmetry. Namely, if $x' = M_A(x)$ then $\bar{x} = M_B(\bar{x}')$. This in turn implies

$$\pi_A(x'|x) = \pi_B(\bar{x}|\bar{x}')/J_A(x) \quad (2.12a)$$

$$\pi_B(y'|y) = \pi_A(\bar{y}|\bar{y}')/J_B(y). \quad (2.12b)$$

The Jacobians are necessary to ensure normalization of π_A and π_B . In Eq. 2.12a, for instance,

$$\begin{aligned} \int dx' \pi_A(x'|x) &= \int dx' \left| \frac{\partial x}{\partial x'} \right| \pi_B(\bar{x}|\bar{x}') \\ &= \int dx \pi_B(\bar{x}|\bar{x}') = 1 \end{aligned} \quad (2.13)$$

(Identifying $q = \ln J$ as reduced heat, this result is equivalent to Eq. 9 of Ref. [27] and is really a statement of a fluctuation theorem.) Finally, the reduced work (Eq. 2.5) is odd under time-reversal,

$$w_A(x \rightarrow x') = -w_B(\bar{x}' \rightarrow \bar{x}). \quad (2.14)$$

Now, combining Eqs. 2.5 and 2.9-2.14 we get

$$\begin{aligned}
K_{acc}(y', x'|x, y) &= \pi(x', y'|x, y) \alpha(x', y'|x, y) \\
&= \frac{\pi(\bar{y}, \bar{x}|\bar{y}', \bar{x}')}{J_A(x) J_B(y)} \alpha(\bar{y}, \bar{x}|\bar{y}', \bar{x}') e^{-w_A(x \rightarrow x') - w_B(y \rightarrow y')} \\
&= K_{acc}(\bar{x}, \bar{y}|\bar{y}', \bar{x}') e^{-h_A(y') - h_B(x') + h_A(x) + h_B(y)} \\
&= K_{acc}(\bar{x}, \bar{y}|\bar{y}', \bar{x}') \frac{p_{AB}^{eq}(\bar{y}', \bar{x}')}{p_{AB}^{eq}(x, y)}, \tag{2.15}
\end{aligned}$$

a detailed balance-like condition for accepted moves.

To show that our scheme satisfies balance, we first note that the transition probability has contributions from accepted and rejected moves:

$$P(y', x'|x, y) = K_{acc}(y', x'|x, y) + K_{rej}(y', x'|x, y) \tag{2.16}$$

The quantity $K_{acc}(y', x'|x, y)$, defined in Eq. 2.11 above, is the joint probability that a swap move is accepted and ends up in y', x' , having begun at (x, y) . Similarly, $K_{rej}(y', x'|x, y)$ is the joint probability that a swap move is rejected and ends up in (y', x') , having started at (x, y) . This rejection contribution is

$$K_{rej}(y', x'|x, y) = [1 - p_{acc}(x, y)] \cdot \delta(y' - \bar{x}) \delta(x' - \bar{y}), \tag{2.17}$$

where the delta functions revert the system back to the time-reverse of their original points upon rejection, as specified in the method. The factor $1 - p_{acc}(x, y)$ is the net probability to reject the swap, conditioned on initial points x and y , and

$$p_{acc}(x, y) = \int dx' \int dy' K_{acc}(y', x'|x, y). \tag{2.18}$$

Finally, to show that the balance condition is satisfied (Eq. 2.6) we integrate P over a canonical distribution of initial points. The resulting contribution from the accepted

moves is

$$\begin{aligned} \int dx \int dy K_{acc}(y', x'|x, y) p_{AB}^{\text{eq}}(x, y) &= \int dx \int dy K_{acc}(\bar{x}, \bar{y}|\bar{y}', \bar{x}') p_{AB}^{\text{eq}}(y', x') \\ &= p_{acc}(\bar{y}', \bar{x}') p_{AB}^{\text{eq}}(\bar{y}', \bar{x}'), \end{aligned} \quad (2.19a)$$

where we used Eqs. 2.15 and 2.18. The contribution from the rejected moves is

$$\begin{aligned} \int dx \int dy K_{rej}(y', x'|x, y) p_{AB}^{\text{eq}}(x, y) &= \int dx \int dy [1 - p_{acc}(x, y)] \cdot \delta(y' - \bar{x}) \delta(x' - \bar{y}) p_{AB}^{\text{eq}}(x, y) \\ &= [1 - p_{acc}(\bar{y}', \bar{x}')] p_{AB}^{\text{eq}}(\bar{y}', \bar{x}'). \end{aligned} \quad (2.19b)$$

By adding Eqs. 2.19a and 2.19b together, we see that our condition of balance, Eq. 2.6 above, is satisfied. Hence our method preserves equilibrium sampling within each replica. In Appendix B we provide an example implementation of a deterministic RENS scheme, where the system evolves under Nosé-Hoover dynamics.

2.3 Derivation: Stochastic dynamics

In this section we derive the validity of RENS for evolution under stochastic equations of motion. This derivation follows logic similar to the deterministic case, where we introduce a parameter-dependent reduced hamiltonian $h(x; \lambda)$, and show that a balance condition is satisfied for our RENS swap moves (Eq. 2.29 below), guaranteeing equilibrium sampling in each replica. In contrast to the deterministic case, however, here we will analyze the evolution in a discretized manner, where trajectories will formally be composed of a finite number of points (see Eq. 2.26), and the switching parameter λ is updated in discrete steps (see below). We will assume that the dynamics are Markovian and satisfy detailed balance when h is fixed. Our analysis is applicable to such dynamics that are *explicitly discrete* in time, such as Monte Carlo evolution, as well as *continuous-time* stochastic dynamics in which trajectories are generated through discretized equations of

motion. Before proceeding to the derivation it is useful to illustrate dynamical schemes which satisfy these assumptions.

We assume evolution under stochastic equations of motion which in general are described by the transition probability $p_\lambda(y|x)$: given a value of λ and current state x , the next configuration y a time δt later is a random sample from the distribution $p_\lambda(y|x)$. We assume that these dynamics satisfy the condition of detailed balance:

$$\frac{p_\lambda(y|x)}{p_\lambda(\bar{x}|\bar{y})} = e^{-[h_\lambda(y)-h_\lambda(x)]} \quad (2.20)$$

This condition implies that the canonical distribution $p_\lambda^{\text{eq}}(x) \propto e^{-h_\lambda}$ is stationary under these dynamics. Under discrete-time Monte Carlo dynamics, detailed balance is enforced explicitly: given current point x , a randomly-generated trial configuration y is accepted according to the Metropolis algorithm [41], which guarantees Eq. 2.20 is satisfied.

If it is desired to implement continuous-time evolution, Langevin dynamics [122] is a valid option. To illustrate this, we consider a single brownian particle of mass m moving in a potential U . In the overdamped limit, the evolution of its position q can be described by the following Langevin equation:

$$\dot{q} = -\frac{1}{m\gamma_p} \frac{dU(q)}{dq} + \eta(t), \quad (2.21)$$

While the Newtonian forces acting on the particle are described by the potential $U(q)$, the random forces are manifested through gaussian white noise term η , obeying statistics $\langle \eta \rangle = 0$ and $\langle \eta(t)\eta(t+t') \rangle = 2D\delta(t-t')$. The diffusion constant $D = k_B T/m\gamma_p$ is given by the fluctuation-dissipation theorem with friction coefficient γ_p . Langevin dynamics also satisfy detailed balance. To demonstrate this on our simple overdamped case, we construct the transition probability by discretizing the dynamics in time, showing Eq. 2.20 is satisfied for δt sufficiently small⁴. Assuming at time $t = 0$ the system exists in configuration q , its

⁴For Langevin evolution, detailed balance only holds in the strict limit $\delta t \rightarrow 0$, and so it is important

position q' a short time δt later is described by

$$q' = q - \frac{\delta t}{m\gamma} \frac{dU(q)}{dq} + \delta q_R. \quad (2.22)$$

Here the random component δq_R is gaussian distributed with mean 0 and variance $\sigma^2 = 2\frac{k_B T}{m\gamma}\delta t$. Given q then, q' is also a gaussian random variable with a shifted mean and variance σ^2 :

$$p_\lambda(q'|q) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[- \left(q' - q + \frac{\delta t}{m\gamma} \frac{dU(q)}{dq} \right)^2 / 2\sigma^2 \right]. \quad (2.23)$$

A similar expression holds for $p_\lambda(q|q')$, which describes the reverse transition. By taking the ratio of $p_\lambda(q'|q)$ and $p_\lambda(q|q')$, substituting $\sigma^2 = 2\frac{k_B T}{m\gamma}\delta t$ and keeping only the leading-order terms in δt we get

$$\frac{p_\lambda(q'|q)}{p_\lambda(q|q')} = \exp \left[-(q' - q) \left(\frac{dU(q)}{dq} + \frac{dU(q')}{dq} \right) / 2k_B T \right]. \quad (2.24)$$

If we further assume $U(q)$ is approximately linear between q and q' (which is justified for sufficiently small δt), then $(q' - q)dU(q)/dq = U(q') - U(q)$ and from Eq. 2.24 we get

$$\frac{p_\lambda(q'|q)}{p_\lambda(q|q')} = \exp \left[- (U(q') - U(q)) / k_B T \right]. \quad (2.25)$$

Such an analysis can be generalized to larger-dimensional systems, such that for $\delta t \rightarrow 0$ detailed balance is satisfied. In addition to MC and Langevin dynamics, other stochastic schemes such as the Andersen thermostat [3] satisfy the above assumptions.

We now proceed to our derivation. To perform the switching dynamics in replicas A and B , we begin by introducing a parameter-dependent Hamiltonian $h(x; \lambda)$ that interpolates between $h(x; 0) = h_A(x)$ and $h(x; 1) = h_B(x)$. Let us specify *switching protocol* $\lambda_A(t)$, for replica A , which switches between $\lambda_A(0) = 0$ and $\lambda_A(\tau) = 1$ through series of discrete steps of duration δt : $\lambda_A(t) = \{\lambda_0, \lambda_{\delta t}, \dots, \lambda_\tau\}$. As before, the switching time τ to choose a sufficiently small timestep for such simulations (see Sec. 2.5 for a discussion).

corresponds to the duration of the switching trajectory (measured in MC steps or simulation time, depending on the dynamics). When a decision is made to perform an exchange, the (reduced) Hamiltonian in \mathcal{R}_A is switched from h_A to h_B according to schedule $\lambda_A(t)$. Under these time-varying dynamics the system in A evolves from point $x_0 \equiv x$, through trajectory γ_A , reaching final point $x_\tau \equiv x'$ in a series of $N = (\tau/\delta t) + 1$ points:

$$\gamma_A : x_0 \Longrightarrow x_0 \xrightarrow{\lambda_{\delta t}} x_{\delta t} \cdots x_\tau \Longrightarrow x_\tau. \quad (2.26a)$$

The notation indicates that between updates of the parameter (\Rightarrow) the system evolves (under detailed balanced-dynamics) at fixed λ (\rightarrow). Simultaneously, as the Hamiltonian in \mathcal{R}_B is switched from h_B to h_A through protocol $\lambda_B(t) = \lambda_A(\tau - t)$, the system evolves from $y_0 \equiv y$ to $y_\tau \equiv y'$ through trajectory γ_B :

$$\gamma_B : y_\tau \Leftarrow y_\tau \cdots y_{\delta t} \xleftarrow{\lambda_{\tau-\delta t}} y_0 \Leftarrow y_0. \quad (2.26b)$$

If γ_A and γ_B are generated from MC dynamics, the evolution $x_0 \xrightarrow{\lambda} x_{\delta t}$ represents a single Monte Carlo move at fixed λ . If, on the other hand, the system evolves under continuous-time stochastic dynamics, this evolution represents propagating the system a time δt under the equations of motion (e.g. Eq. 2.22) with fixed λ , and δt is the simulation timestep.

Over the course of the switching simulations an amount of reduced work is performed in each replica. For the case of MC dynamics, we define the work in \mathcal{R}_A as

$$w_A = \sum_{i=0}^{N-1} h(x_{i\delta t}, \lambda_{(i+1)\delta t}) - h(x_{i\delta t}, \lambda_{i\delta t}), \quad (2.27a)$$

quantifying the change in “reduced energy” along γ_A due to variation of the control parameter $\lambda_{i\delta t} \rightarrow \lambda_{(i+1)\delta t}$ at fixed $x_{i\delta t}$. Similarly, in \mathcal{R}_B we have

$$w_B = \sum_{i=0}^{N-1} h(y_{i\delta t}, \lambda_{\tau-i\delta t}) - h(y_{i\delta t}, \lambda_{\tau-(i-1)\delta t}), \quad (2.27b)$$

For continuous-time dynamics, assuming we are in the small- δt limit, we can write work (and heat below) as integrals rather than sums: In \mathcal{R}_A , we define the work as

$$w_A = \int_0^\tau dt \dot{\lambda}_A \frac{\partial h(x_t; \lambda_A)}{\partial \lambda_A}, \quad (2.28a)$$

quantifying the change in “reduced energy” along γ_A due to variation of the control parameter λ_A . Similarly, in \mathcal{R}_B we have

$$w_B = \int_0^\tau dt \dot{\lambda}_B \frac{\partial h(y_t; \lambda_B)}{\partial \lambda_B}. \quad (2.28b)$$

[Of course, when calculating w_A and w_B from simulations, Eq. 2.28 is evaluated as a sum.]

We note that under certain conditions⁵ our definition of work for deterministic dynamics, Eq. 2.5 above, can be shown to be equal to Eq. 2.28 (or Eq. 2.46 under more general dynamics discussed below).

Upon completion of the switching simulations, the swap attempt is made between the final points as described above and illustrated in Fig. 2.1. As above, to verify the validity of RENS we have to show that the swap move satisfies balance,

$$\int dx \int dy P(y', x' | x, y) p_{AB}^{\text{eq}}(x, y) = p_{AB}^{\text{eq}}(y', x'), \quad (2.29)$$

such that the probability flux into (y', x') , from all possible states (x, y) preserves the distribution p_{AB}^{eq} . As before, $P(y', x' | x, y)$ is the transition probability to state (y', x') from (x, y) , and $p_{AB}^{\text{eq}}(x, y) = p_A^{\text{eq}}(x) p_B^{\text{eq}}(y)$ is the joint equilibrium distribution of replicas A and B .

To establish balance for this stochastic evolution, we follow the logic of the previous derivation. In contrast to the deterministic case, however, we must now explicitly take into account entire path probabilities when describing the evolution. We first analyze

⁵The deterministic definition for work, Eq. 2.5, can be written as Eq. 2.28 if the canonical distribution is stationary under the deterministic dynamics when λ is held fixed.

our system at the level of path probabilities, showing that our scheme satisfies a *detailed balance* condition in path space (Eq. 2.39), which, in turn, via path integration (Eq. A-13), leads to a detailed balance condition for accepted moves (Eq. 2.41). By using this detailed balance condition in combination with the contributions to the transition probability P arising from rejected moves, we arrive at our balance condition, Eq. 2.29.

To begin our path analysis, we first define

$$\pi_A(\gamma_A) = P_A(x_{\delta t}, \dots, x_\tau | x_0) \quad (2.30)$$

as the path probability of γ_A (Eq. 2.26a), conditioned on its initial point. Because of our Markov assumption, P_A reduces to a product of pairwise transition probabilities,

$$P_A(x_{\delta t}, \dots, x_\tau | x_0) = \prod_{i=0}^{N-2} p_i(x_{(i+1)\delta t} | x_{i\delta t}). \quad (2.31)$$

Expressions analogous to Eq. 2.30 and 2.31 hold for $\pi_B(\gamma_B)$. We note that in contrast to deterministic dynamics, where π_A depended solely on its initial point, here the path probability depends on the entire sequence of points specified by γ_A . Now let

$$\pi(\gamma_A, \gamma_B) = \pi_A(\gamma_A)\pi_B(\gamma_B) \quad (2.32)$$

be the probability of generating γ_A and γ_B , conditioned on initial points x and y . Also, let

$$\alpha(\gamma_A, \gamma_B) = \min\{1, e^{-w_A(\gamma_A) - w_B(\gamma_B)}\} \quad (2.33)$$

be the probability of accepting the trial replica exchange. Finally, let $g = \pi\alpha$. We can view g as the path-analogue of the transition probability P , defined above: $g(\gamma_A, \gamma_B)$ is the probability of generating the trajectories γ_A and γ_B , and then accepting the swap, given initial points x and y .

Now, for a specified trajectory $\gamma_A = (x_0 \rightarrow x_\tau)$ – which evolves as λ is varied from 0 to 1 in replica A – let us define its *conjugate twin*, $\tilde{\gamma}_B = (\bar{x}_0 \leftarrow \bar{x}_\tau)$, to be the time-reversed

trajectory, passing through the same configurations only in time-reversed order and with reversed momenta, as λ is varied from 1 to 0 in replica B . Similarly, for γ_B (Eq. 2.26b) we define a twin $\tilde{\gamma}_A$. As before we assume the Hamiltonian is time-reversal invariant, $h(\bar{x}; \lambda) = h(x; \lambda)$, i.e. the system does not contain any magnetic fields.

From Eq. 2.28 (or Eq. 2.27) we have $w_A(\gamma_A) = -w_B(\tilde{\gamma}_B)$ and $w_A(\tilde{\gamma}_A) = -w_B(\gamma_B)$.

Combining this with Eqs. 2.32 and 2.33 we get

$$\frac{g(\gamma_A, \gamma_B)}{g(\tilde{\gamma}_A, \tilde{\gamma}_B)} = \frac{\pi_A(\gamma_A)}{\pi_B(\tilde{\gamma}_B)} e^{-w_A(\gamma_A)} \cdot \frac{\pi_B(\gamma_B)}{\pi_A(\tilde{\gamma}_A)} e^{-w_B(\gamma_B)}. \quad (2.34)$$

Now, using our assumption of detailed balanced-dynamics (Eq. 2.20), along with Eq. 2.31, a given trajectory and its twin satisfy [27]

$$\begin{aligned} \frac{\pi_A(\gamma_A)}{\pi_B(\tilde{\gamma}_B)} &= \prod_{i=0}^{N-2} \frac{p_i(x_{(i+1)\delta t} | x_{i\delta t})}{p_i(x_{i\delta t} | x_{(i+1)\delta t})} \\ &= \prod_{i=1}^N \exp[h(x_{i\delta t}; \lambda_{i\delta t}) - h(x_{(i-1)\delta t}; \lambda_{i\delta t})] \\ &= e^{-q_A(\gamma_A)}, \end{aligned} \quad (2.35)$$

where

$$q_A(\gamma_A) = \sum_{i=1}^N h(x_{i\delta t}; \lambda_{i\delta t}) - h(x_{(i-1)\delta t}; \lambda_{i\delta t}). \quad (2.36)$$

is the reduced heat. In its continuous-time form, we can write the heat as

$$q_A(\gamma_A) = \int_0^\tau dt \dot{x} \cdot \nabla h(x_t; \lambda_A(t)). \quad (2.37)$$

Eq. 2.35 is the analogue of Eq. 2.12a we found for deterministic dynamics, a fluctuation theorem for conditional path probabilities originally shown by Crooks [27]. From Eqs. 2.28 and 2.37 (or Eqs. 2.27 and 2.36), we have

$$w_A(\gamma_A) + q_A(\gamma_A) = h_B(x_\tau) - h_A(x_0). \quad (2.38)$$

Using Eqs. 2.35 and 2.38 and their analogues for the conjugate pair $(\gamma_B, \tilde{\gamma}_A)$, Eq. 2.34 gives us

$$\frac{g(\gamma_A, \gamma_B)}{g(\tilde{\gamma}_A, \tilde{\gamma}_B)} = \frac{p_{AB}^{\text{eq}}(y_\tau, x_\tau)}{p_{AB}^{\text{eq}}(x_0, y_0)}, \quad (2.39)$$

a detailed balance condition in path space.

As before, to show that our method satisfies balance, we consider the contributions to the acceptance probability $P(y', x'|x, y) = K_{\text{acc}}(y', x'|x, y) + K_{\text{rej}}(y', x'|x, y)$ arising from accepted and rejected moves. The quantity K_{acc} is the joint probability of accepting the move and ending up in (y', x') , conditioned on initial points (x, y) . Similarly, K_{rej} is the joint probability of rejecting the move and ending up in (y', x') , conditioned on initial points (x, y) . Here K_{acc} is related to the g via path integration:

$$K_{\text{acc}}(y', x'|x, y) = \int \mathcal{D}'\gamma_A \int \mathcal{D}'\gamma_B g(\gamma_A, \gamma_B), \quad (2.40)$$

where the notation $\int \mathcal{D}'\gamma_A = \int dx_{\delta t} \cdots \int dx_{\tau-\delta t}$ represents an integral over all intermediate points for trajectories in \mathcal{R}_A that begin and end in the fixed configurations $x_0 \equiv x$ and $x_\tau \equiv x'$. (Analogous comments apply to $\int \mathcal{D}'\gamma_B$.) Writing the corresponding expression for the conjugate trial move, $(\bar{x}, \bar{y}) \leftarrow (\bar{y}', \bar{x}')$, we get

$$\begin{aligned} K_{\text{acc}}(\bar{x}, \bar{y}|\bar{y}', \bar{x}') &= \int \mathcal{D}'\tilde{\gamma}_A \int \mathcal{D}'\tilde{\gamma}_B g(\tilde{\gamma}_A, \tilde{\gamma}_B) \\ &= \frac{p_{AB}^{\text{eq}}(x, y)}{p_{AB}^{\text{eq}}(y', x')} \int \mathcal{D}'\gamma_A \int \mathcal{D}'\gamma_B g(\gamma_A, \gamma_B) \\ &= \frac{p_{AB}^{\text{eq}}(x, y)}{p_{AB}^{\text{eq}}(y', x')} K_{\text{acc}}(y', x'|x, y). \end{aligned} \quad (2.41)$$

Here we have used Eq. 2.39, along with the correspondence between a trajectory and its conjugate twin ($\mathcal{D}'\tilde{\gamma}_A = \mathcal{D}'\gamma_B$, etc.). Eq. 2.41 is a statement of detailed balance for accepted moves, which we also arrived at in our previous derivation, Eqs. 2.15.

To complete the proof we simply follow the last two paragraphs of our previous derivation (Eqs. 2.16-2.19), where we use the detailed balance expression above as well as

the expression for the rejected moves. As before, the rejection kernel is

$$K_{rej}(y', x'|x, y) = [1 - p_{acc}(x, y)] \cdot \delta(y' - \bar{x})\delta(x' - \bar{y}), \quad (2.42)$$

where the factor $1 - p_{acc}(x, y)$ is the net probability to reject the swap, conditioned on initial points x and y , and

$$p_{acc}(x, y) = \int dx' \int dy' K_{acc}(y', x'|x, y). \quad (2.43)$$

To show balance is satisfied, one need integrate $P = K_{acc} + K_{rej}$ over a canonical distribution of initial points. This can be shown by following the steps, Eqs. 2.19a and 2.19b in our previous derivation, together with Eqs. 2.41 and 2.42 derived for stochastic dynamics. As before we find that the equilibrium distribution is preserved, hence our scheme with stochastic dynamics samples equilibrium in each replica.

2.4 Velocity rescaling and other escorted switching dynamics

In the stochastic derivation above we have implicitly assumed that the system evolves under equations of motion that model the system's physical dynamics (e.g. Langevin dynamics). It is, however, often useful to augment the equations of motion with terms that couple the system evolution to changes in λ . Such dynamical schemes, although unphysical, can also be incorporated into RENS, and, if chosen wisely, can enhance replica overlap.

As an example of such a scheme, it is standard practice in temperature REM to rescale the system velocities upon a successful swap [115]. Here the system momenta in replica A undergo the transformation $\mathbf{p} \rightarrow \sqrt{T_B/T_A} \mathbf{p}$ upon being copied into replica B , and the inverse scaling occuring in replica B . This rescaling enhances replica overlap because immediately after the swap, the distribution of momenta is given by Maxwell-

Boltzmann statistics in the corresponding replica. This rescaling practice can readily be generalized to RENS. Consider, for example, a discretized switching protocol where updates of λ occur at discrete intervals in time (see e.g. Fig. 3.10), and let each update $\lambda_i \rightarrow \lambda_{i+1}$ be accompanied by mapping $\mathbf{p} \rightarrow \alpha_i \mathbf{p}$, with $\alpha_i = \sqrt{T_{\lambda_{i+1}}/T_{\lambda_i}}$. The scaling factor α_i artificially heats up / cools down the replica in response to a changing environment, providing a means to push the system closer to its equilibrium state by artificially changing the kinetic energy commensurate with changes in the heat bath temperature.

This velocity rescaling is one example of a general scheme in nonequilibrium simulations in which the system evolution is coupled to changes in λ . Consider equations of motion of the form

$$\frac{dx}{dt} = v(x; \lambda) + \dot{\lambda}u(x; \lambda), \quad (2.44)$$

where $v_\lambda(x) = v(x; \lambda)$ corresponds to ordinary equilibrium-preserving sampling (e.g. Langevin dynamics), and the flow field $u_\lambda = u(x; \lambda)$ acts on the system upon changes in λ . The velocity rescaling above corresponds to a particular choice of u_λ (see Eq. 2.47). The use of switching simulations that evolve according to Eq. 2.44 has been developed previously within the context of free energy simulations [58, 81, 119]. It has been shown that these dynamics satisfy a fluctuation theorem [121], with a modified definition of work.

We note that the evolution given by Eq. 2.44 can be considered a continuous-time version of the switching dynamics used in the related NCMC method [85], which consist of alternating sets of propagation steps at fixed λ and perturbation steps during discrete λ jumps.

RENS remains valid under these more general dynamics. If the dynamics described by Eq. 2.44 are deterministic and time-reversible, they satisfy the assumptions required for the derivation in Sec. 2.2, in which case the reduced work is given by (Eq. 2.5). In

replica A , for instance,

$$w^{(u)}(x_0 \rightarrow x_\tau) = h_B(x_\tau) - h_A(x_0) - \ln J_A(x_0) \quad (2.45)$$

If the dynamics are stochastic, RENS still remains valid⁶, with a modified definition of work which accounts for the presence of u_λ :

$$w^{(u)} = \int_0^\tau dt \dot{\lambda} \left[\frac{\partial h_\lambda}{\partial \lambda} + u_\lambda \cdot \nabla h_\lambda - \nabla \cdot u_\lambda \right]. \quad (2.46)$$

The first term on the right hand side of Eq. 2.46 corresponds to the normal definition of work (Eq. 2.28); the two additional terms roughly account for the change in energy and entropy-like phase space compression due to the influence of u_λ .

Consider again the velocity rescaling above with a generic switching protocol. For particle i at point $\{\mathbf{q}, \mathbf{p}\}_i$, the flow field that achieves this rescaling acts independently on each i as

$$\{u_\lambda\}_i = \{0, s_\lambda \mathbf{p}\}_i, \quad s_\lambda = (1/2T_\lambda)(dT_\lambda/d\lambda). \quad (2.47)$$

For a Hamiltonian with kinetic and potential parts, $h_\lambda = k_\lambda + v_\lambda$, the scaling of velocities achieved by the flow field is such that $\partial k_\lambda / \partial \lambda$ is exactly canceled by $u_\lambda \cdot \nabla h_\lambda$, simplifying the expression for the work in Eq. 2.46 above. The last term of Eq. 2.46 represents phase space contraction of the system, and its integral can be evaluated analytically. In replica A , for instance,

$$\int_0^\tau dt \dot{\lambda}_A \nabla \cdot u_\lambda = \frac{N_d}{2} \ln(T_B/T_A), \quad (2.48)$$

where N_d is the number of degrees of freedom. Similarly, the last term of Eq. 2.46 evaluated

⁶As in the derivations above, the validity of RENS for these new dynamics rests upon a fluctuation theorem derived by Vaikuntanathan and Jarzynski (Eq. 33 of Ref. [121])

in replica B gives us the inverse of Eq. 2.48. The work in each replica is then given by

$$w_A^{(u)} = \int_0^\tau dt \dot{\lambda}_A \frac{\partial v_{\lambda_A}}{\partial \lambda_A} + \frac{N_d}{2} \ln(T_B/T_A), \quad (2.49a)$$

$$w_B^{(u)} = \int_0^\tau dt \dot{\lambda}_B \frac{\partial v_{\lambda_B}}{\partial \lambda_B} - \frac{N_d}{2} \ln(T_B/T_A). \quad (2.49b)$$

The second term in Eq. 2.49 can be ignored for the purposes of the acceptance/rejection decision, since the two terms cancel when adding $w_A^{(u)}$ and $w_B^{(u)}$. Hence the work is determined solely by the variation of the reduced potential v with time. While this simple flow field only couples to momentum degrees of freedom, we have found it can substantially enhance the replica exchange acceptance rate. More elaborate choices of u_λ which couple to configurational degrees of freedom can enhance the efficiency of free energy calculations by orders of magnitude [58, 119, 121], and we expect such couplings to also be beneficial in RENS.

2.5 A consistency check for equilibrium sampling

In RENS, trial swap configurations are generated by driving the replicas out of equilibrium, but by accepting or rejecting the swap with a work-based criterion, the entire move satisfies the balance condition. Although in principle this guarantees equilibrium sampling, any real simulation is bound to contain errors due approximations made, such as a finite timestep, insufficient simulation length etc., not to mention possible coding mistakes. In this section we discuss a procedure⁷ to verify equilibrium sampling from finite-time RENS simulations that combines Crooks's fluctuation theorem for work distributions [26], Eq. 2.50 below, with Bennett's overlapping distribution method [13]. We then test this procedure on our simulations.

⁷See Ref. [96] for a discussion of this procedure within the context of free energy calculations.

A symmetry exists between trajectories generated during switching intervals in \mathcal{R}_A and \mathcal{R}_B : Crooks's fluctuation theorem [27], integral in the derivation above (Eq. 2.35), relates the statistical weight of a trajectory observed in A and its conjugate twin (time-reversed counterpart) in B . This symmetry leads to a fluctuation theorem [26] (FT) that relates the distribution of work values generated in \mathcal{R}_A to that of \mathcal{R}_B :

$$\frac{\rho_A(w)}{\rho_B(-w)} = e^{w-\Delta f}, \quad (2.50)$$

Here $\Delta f = -\ln [\int dx e^{-h_B} / \int dx e^{-h_A}]$ is the reduced free energy difference between thermodynamic states of replicas A and B . [Eq. 2.50 can be derived by performing an appropriate path integral over Eq. 2.35 with initial conditions sampled canonically.] An underlying assumption of Eq. 2.50 is that at the onset of the switching intervals, \mathcal{R}_A and \mathcal{R}_B begin in their respective equilibrium states. As such, the extent to which Eq. 2.50 is satisfied provides an indication of whether equilibrium sampling is achieved during the sampling intervals.

As a consistency check on our simulations we can graphically verify the FT through the overlapping distribution method [13, 41] applied to switching segments of finite duration. Following Bennett, let us define two functions:

$$L_A(w) = \ln \rho_A(w) - w/2, \quad (2.51)$$

$$L_B(w) = \ln \rho_B(-w) + w/2. \quad (2.52)$$

By use of Eq. 2.50, we see that L_A and L_B combined give a free energy estimator, $\Delta f = L_B(w) - L_A(w)$, for every value of w . Furthermore, the flatness of $\Delta L(w) \equiv L_B(w) - L_A(w)$ (over the range of values for which we have good sampling statistics for both ρ_A and ρ_B) is an indication of the extent to which Eq. 2.50 is satisfied.

We demonstrate this consistency check in Fig. 2.2, where we apply the above analysis

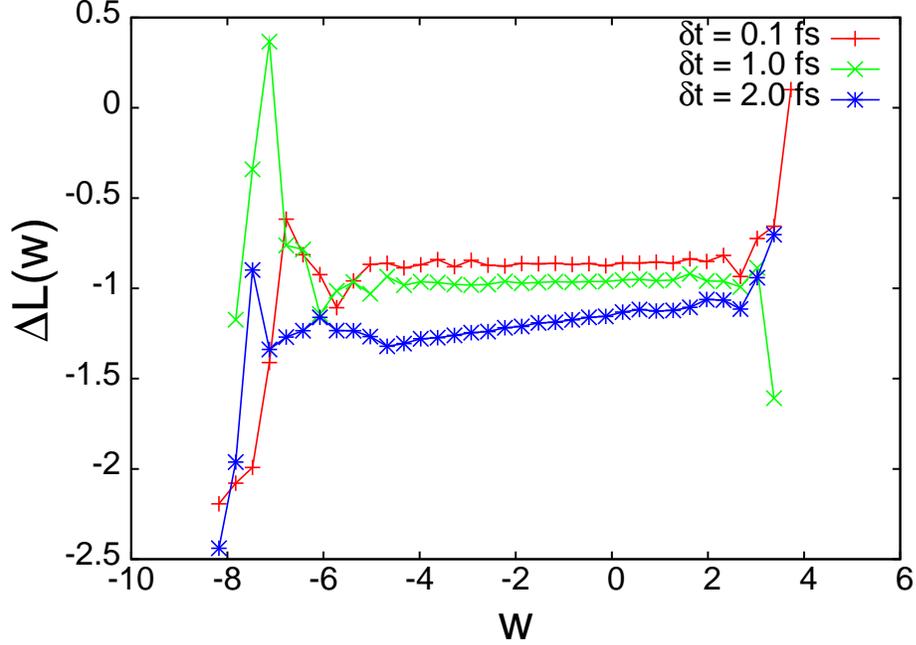


Figure 2.2: Choosing a sufficiently small timestep: $\Delta L(w)$ is plotted for REM simulations of alanine dipeptide with various timesteps. A flat line indicates Crooks’s fluctuation theorem is satisfied (see Eqs. 2.50 and 2.51), suggesting data is sampled from equilibrium.

on RENS simulations of alanine dipeptide (see Sec. 3.3 below). In these simulations the system evolved under Langevin dynamics, which formally only satisfies detailed balance in the limit the integration timestep $\delta t \rightarrow 0$ [111]. To choose a small enough timestep to sample equilibrium, we calculate $\Delta L(w)$ for REM simulations with timesteps of 2.0, 1.0 and 0.1 fs, displayed in Fig. 2.2. L_A, L_B , and ΔL were calculated by constructing histograms of work values in \mathcal{R}_A and \mathcal{R}_B . While for $\delta t = 2.0$ fs we see a systematic slope appear, the smaller timesteps are flat within statistical errors, suggesting that $\delta t = 1.0$ fs is sufficiently small. However, because we observe a small offset in the plots between the two smaller timesteps, we chose the more conservative $\delta t = 0.1$ fs in our simulations below.

2.6 Summary

In this chapter we introduced our method of Replica Exchange with Nonequilibrium Switches in detail. While RENS is valid for a very general class of dynamics, its proper implementation depends upon the evolution used to model the system. We described and derived our method for systems evolving under deterministic and stochastic equations of motion. We further discussed novel dynamical schemes which can be used in conjunction with RENS to enhance replica overlap, as well as a consistency check that can be used to verify equilibrium sampling. In the following chapter we proceed to test RENS on model systems, demonstrating the sampling gains our method can bring.

Replica Exchange with Nonequilibrium Switches: numerical studies

¹Systems that are amenable to replica exchange are typically complex, containing rugged energy landscapes with deep wells and high barriers. Under standard Molecular Dynamics, these barriers give rise to infrequent transitions between metastable states, hindering full phase space exploration and supplying a bottleneck to system equilibration on long timescales. Replica exchange provides an additional channel for transitions between metastable states, realized by configurational swaps between replicas. An effective implementation will significantly lower transition times such that system equilibration is achieved much more rapidly.

In the previous chapter we showed that RENS provides a means to sample equilibrium distributions, and described a variety of dynamical schemes that can be used to

¹This chapter is based on the papers “Replica exchange with nonequilibrium switches“, A. J. Ballard and C. Jarzynski, *Proc. Natl. Acad. Sci. U.S.A.*, 106 (30):12224 (2009) and “Replica Exchange with Nonequilibrium Switches: Enhancing Equilibrium Sampling by Increasing Replica Overlap“, A. J. Ballard and C. Jarzynski, *J. Chem. Phys.*, 136 194101 (2012).

perform RENS simulations. In this chapter we test our method on model systems, with the goal of understanding when exactly it can be advantageous to use RENS over REM. The two systems we consider below are model complex systems, containing rugged energy landscapes. Our first example, presented in Sec. 3.2, is a system of independent particles moving in a 1D potential containing deep wells separated by high energetic barriers. While this example is uninteresting as a physical system, it allows us to check our empirical results against numerically-exact values as well as perform benchmarking to a high statistical accuracy. Our second example, an alanine dipeptide molecule, is more relevant physically and a prototypical example of a complex biomolecule containing long-lived metastable states. We present this system in Sec. 3.3.

Before moving to the numerical studies, we first describe in Sec. 3.1 the intuition we have for when RENS can be advantageous, as well as the methods we will use to numerically test this intuition on our simulation results.

3.1 Using RENS to enhance simulation efficiency

The switching simulations in RENS provide an added flexibility to replica exchange which can be used to enhance simulation efficiency. The choice of the switching time τ , the duration of the switching segment, ultimately determines the overlap shared between replicas and the sampling quality of the simulation in general. In building up an intuition for RENS, it is instructive to look at a couple limiting cases for the switching simulations: For $\tau = 0$, our method reduces to REM, where the overlap is determined solely by the equilibrium distributions of A and B . If there is little overlap to begin with, we expect that REM would be inefficient, and we would do better to increase τ . Taking the opposite limit, $\tau \rightarrow \infty$, the switching between h_A and h_B occurs infinitely slowly, and,

if properly thermostated, the system evolves reversibly between equilibrium distributions corresponding to A and B . In this case [67]

$$w_A = \Delta f = f_B - f_A = -w_B, \quad (3.1)$$

where $f_A = -\ln \int dx e^{-h_A}$ is the reduced free energy difference of \mathcal{R}_A ; hence $P_{\text{acc}} = 1$. However, in this limit a very large amount of time is spent on the switching segments, which will also be inefficient. In any case, we intuit that by increasing τ , we can also increase the acceptance rate of replica exchanges, at the expense of the added simulation time τ . Because an increased acceptance rate comes at the cost of longer switching simulations (which can't directly be used for equilibrium sampling), this implies a computational tradeoff. What is the optimal τ ?

One of the goals in our examples below will be to characterize the effect of τ on the sampling quality and overall simulation efficiency, which will be done as follows. In our analysis we will consider M replicas, and assume for specificity that we are mainly interested in sampling from one of them, which we denote the *primary replica*. The term *output trajectory* will denote the trajectory obtained by concatenating the sampling intervals generated in the primary replica, after discarding the switching intervals. The output trajectory samples the equilibrium distribution of interest, and is in effect the end product of our method. Now, the sampling quality will be quantified through a correlation time t_c associated with the long-time system relaxation observed from the output trajectory. This correlation time determines the statistical independence of the equilibrium samples and hence provides a measure for the quality of sampling. The simulation efficiency will be determined through calculation of a *sample cost* t^* , defined as the total CPU time needed to generate one output trajectory of length t_c . As such, the sample cost t^* determines the total CPU time needed to generate one statistically-independent configuration in the

primary replica (see Eqs. 3.4 and 3.8 below). The optimal RENS implementation will minimize t^* . We now proceed to our examples, where we will analyze the performance of RENS using the above quantities.

We note that in the following simulations we have not implemented the inversion of momenta upon rejection of a trial move (see Sec. 2.1). Our data suggest that this inversion has little effect on the work distributions, and therefore on the applicability of our method. Preliminary efficiency calculations including momentum inversion show trends that remain qualitatively similar to our numerics reported below.

3.2 Example 1: 1D system

Our first example is a model system of $n_p = 10$ particles, moving independently in the potential shown in Fig. 3.1. For our simulations we chose $M = 2$ replicas, at temperatures $T_A = 0.30$ and $T_B = 2.0$ (arbitrary units). In the primary replica at $T_A = 0.30$, sampling is hindered by the barriers separating the local minima of $U(x)$; whereas at $T_B = 2.0$ the particles are able to jump from well to well. We performed Molecular Dynamics simulations on this system using the velocity scaling dynamics described above (see Sec. 2.4) in combination with an Andersen thermostat [3] (see Appendix A).

While REM performs well for a single particle in this potential [41], when $n_p = 10$ it encounters difficulties due to poor phase space overlap, as can be understood using an argument due to Kofke [64]. Typically, in \mathcal{R}_A each particle is found near a local minimum of $U(x)$, while in \mathcal{R}_B they are distributed more uniformly. Thus a configuration swap is likely to be accepted only if all particles in \mathcal{R}_B are found very near the minima of $U(x)$, which is unlikely when $n_p \gg 1$. With RENS, the switching simulation in \mathcal{R}_B increases the swap acceptability by shepherding the particles closer to the minima of $U(x)$.

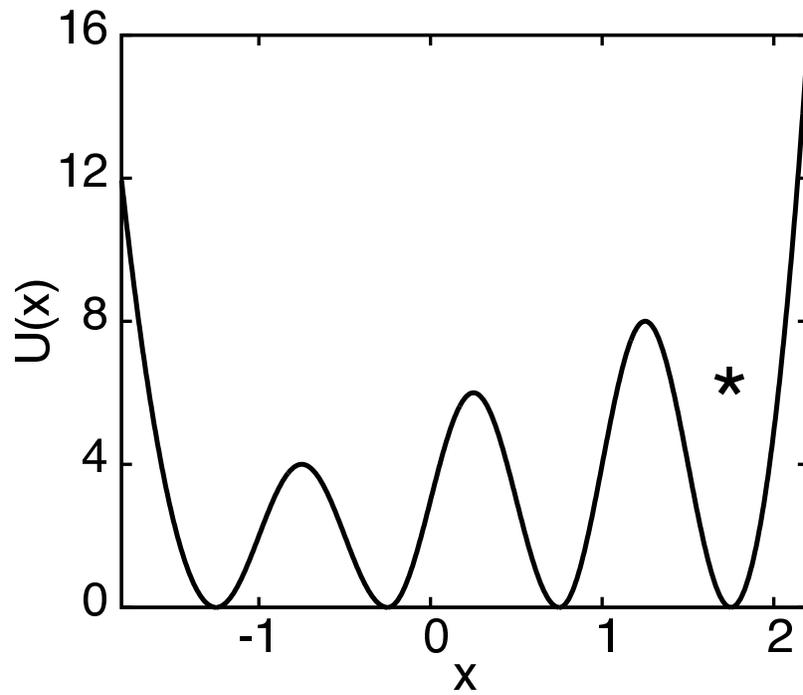


Figure 3.1: Mock-up of a rough potential energy landscape, adapted from Ref. [41], Chapter 14. An asterisk marks the fourth well ($x \geq 1.25$). Ordinary replica exchange works well for one particle, but encounters difficulties when $n_p = 10$.

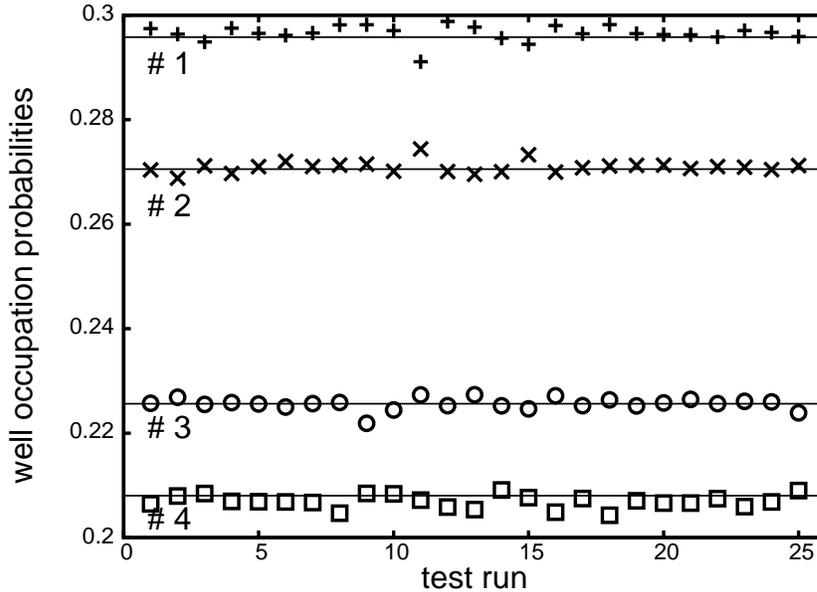


Figure 3.2: Observed occupations of wells 1 - 4, obtained by following a single tagged particle in the output trajectory, for each of the 25 test runs. The solid horizontal lines are exact values determined by integration of the single-particle Boltzmann distribution.

When simulating this system using RENS, the replicas “toggle” between sampling and switching intervals (Fig. 2.1). We implemented this as follows. During an interval of sampling, a switching simulation was initiated at random, with an attempt rate $r = 0.166$. Once initiated, the work interval lasted for the prescribed switching time τ , after which the replicas reverted to sampling, and so on. Thus the average duration of a sampling interval was $\bar{\tau}_{\text{eq}} = 1/r \approx 6.0$, which is roughly three times the relaxation rate *within* one of the local wells of U .

With these parameters, we performed twenty-five test runs, with τ ranging from 0 to 100. To establish proof of principle, we tabulated empirical occupation probabilities for the four wells by following a tagged particle in the output trajectory. For each test run,

we found that the relative amount of time the particle spent in each well was in agreement with the equilibrium distribution, within statistical error (Fig. 3.2).

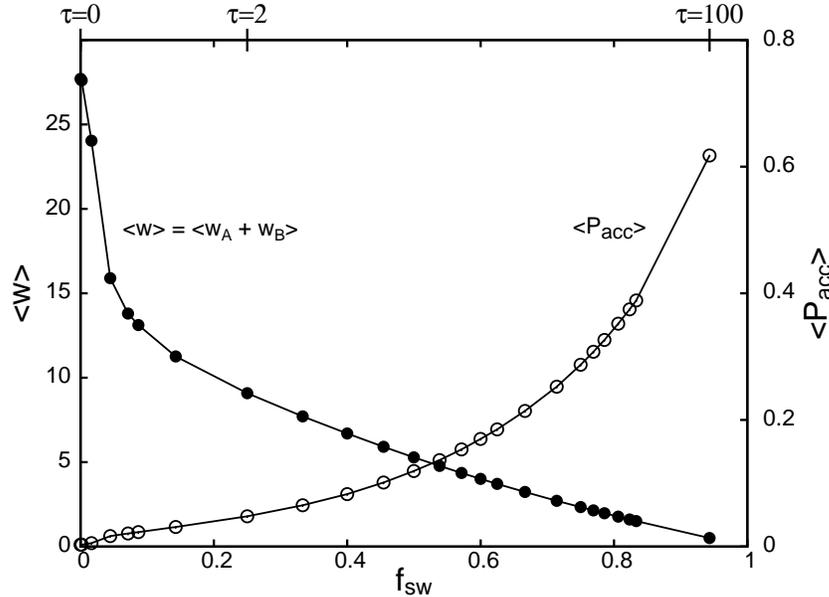


Figure 3.3: Average reduced work (filled circles) and observed acceptance frequency (open circles), as functions of the fraction of simulation time devoted to switching intervals. $\langle P_{\text{acc}} \rangle$ is defined as the fraction of attempted replica exchanges that were accepted, in a given test run. At $f_{\text{sw}} = 0$, corresponding to ordinary (instantaneous) replica exchange, $\langle P_{\text{acc}} \rangle \approx 0.003$.

For the same set of test runs, in Fig. 3.3 we plot the observed swap acceptance frequency and average reduced work, as functions of the fraction of simulation time devoted to the switching intervals,

$$f_{\text{sw}} = \frac{\tau}{\bar{\tau}_{\text{eq}} + \tau}. \quad (3.2)$$

As anticipated, with increasing f_{sw} (or τ) we approach the reversible limit of $w = 0$ and $P_{\text{acc}} = 1$ (Eq. 3.1).

To illustrate the accelerated sampling achieved with RENS, we considered $n_4(t)$, the

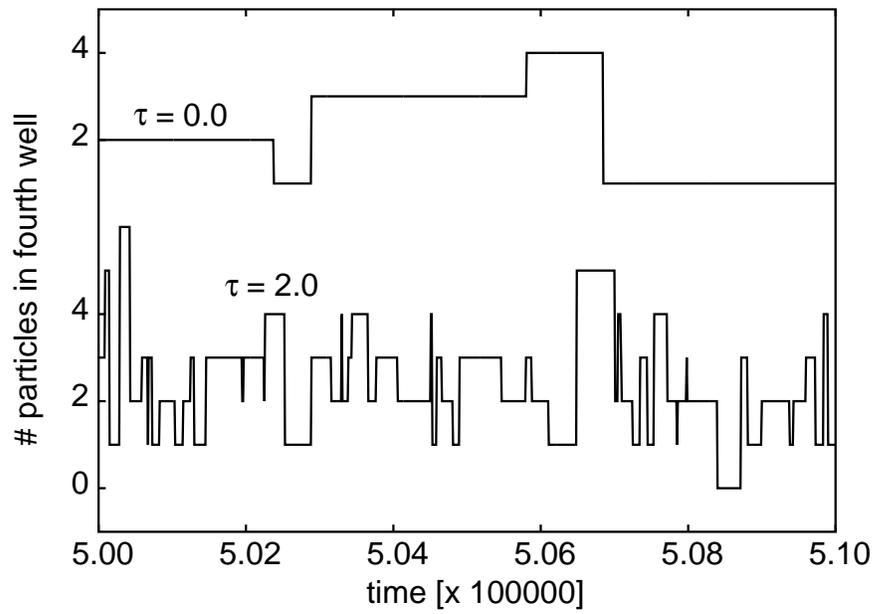


Figure 3.4: $n_4(t)$ is plotted over an interval of time, for output trajectories obtained with REM (upper trace, shifted for clarity) and RENS (lower trace). The two traces represent roughly the same number of attempted replica exchanges (≈ 1700), but reveal substantially different acceptance rates.

number of particles found in the fourth well of $U(x)$ at time t of the output trajectory. Fig. 3.4 shows $n_4(t)$ for segments of the $\tau = 0.0$ and $\tau = 2.0$ test runs. For the relatively modest cost of setting aside 25% of the simulation time to the work intervals, transitions into and out of the fourth well are greatly facilitated.

The long-time relaxation of our system is described by the random variable $n_4(t)$. To evaluate this, we calculate for each test run a correlation time

$$t_c = \frac{1}{\sigma^2} \int_{-\infty}^{+\infty} dt C(t), \quad (3.3)$$

where σ^2 and $C(t)$ are the variance and auto-correlation of $n_4(t)$. The correlation time was calculated using block-averaging [39]. As expected, t_c decreases monotonically with f_{sw} (data not shown).

To measure simulation efficiency we define the sample cost t^* , introduced in Sec. 3.1 above, as

$$t^* = (1 + X)Mt_c, \quad (3.4)$$

with

$$X \equiv \tau/\bar{\tau}_{\text{eq}}. \quad (3.5)$$

Here t^* is a measure of the total computational cost, *summed over all M replicas* [129], of producing a single, statistically independent sample in the primary replica. The factor $(1 + X)$ accounts for the overhead cost of the switching intervals: for every unit of sampling time, X units of time were devoted to the discarded switching simulations. The sample cost provides a figure of merit; the smaller the value of t^* , the more efficiently we are using the computational resources. While the correlation time t_c generally decreases with increasing M or τ – through the randomizing effect of successful replica exchanges – in Eq. 3.4 this trend competes with the overhead factors M and $1 + X$.

In Fig. 3.5 we plot the sample cost t^* , for each test run (empty squares). At $f_{\text{sw}} = 0$ (that is, when using REM), this cost is high, $t^* > 4000$; few swaps are accepted, and particles are trapped in the fourth well for long times. As we increase f_{sw} the sample cost drops significantly, reaching a broad minimum $t^* \sim 450 - 500$ for $f_{\text{sw}} \sim 0.2 - 0.6$; here the allocation of CPU time to switching intervals delivers a clear benefit. For $f_{\text{sw}} > 0.6$ we enter a regime of diminishing returns: $\langle P_{\text{acc}} \rangle$ continues to increase with τ (Fig. 3.3), but not enough to justify the expense of increasingly long switching simulations.

For this efficiency calculation we have neglected the computational cost of the acceptance / rejection step itself, as well as that of the possible subsequent exchange of configurations (which typically involves communication between different processors). Moreover, we have assumed identical costs, per unit simulation time, for the switching and the sampling intervals. It is easy enough to drop the latter assumption: we replace X by αX in Eqs. 3.4, where α is the observed CPU cost of generating a switching simulation, relative to that of a sampling trajectory of equal duration. In our test runs we found $\alpha = 2.9$, and the points shown as filled squares in Fig. 3.5 have been adjusted for this value.

Whether or not we make the adjustment to account for $\alpha \neq 1$, Fig. 3.5 clearly shows that for a fixed set of replicas, it can be highly advantageous to use nonequilibrium switching simulations to generate attempted configuration swaps. The benefits of increased acceptance substantially outweigh the overhead cost of generating the trial configurations.

With RENS we improve efficiency by tuning the switching time, τ , as in Fig. 3.5. With REM, one can instead vary the number of replicas. To compare these two options, we performed test runs of REM ($\tau = 0$) at $M = 2, 3, \dots, 11$. (In each run, we set $T_1 = 0.30$ and $T_M = 2.0$, with intermediate replicas spaced evenly in T^{-1} .) Among these runs, the smallest sample cost, $t^* = 706$, was achieved with $M = 4$ replicas, and is shown as a

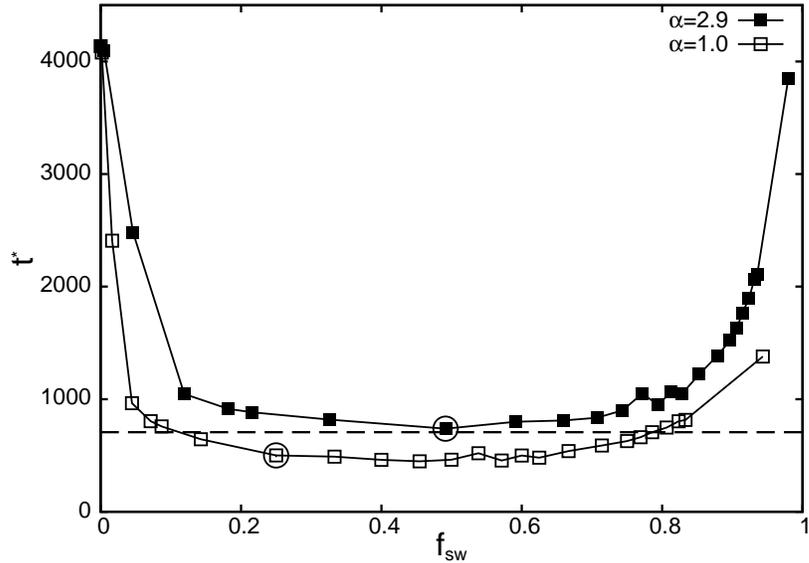


Figure 3.5: Sample cost t^* plotted against f_{sw} , both with and without adjustment for increased relative cost of work simulations (see text). The circles identify the run at $\tau = 2.0$. The dashed line is the sample cost of REM with $M = 4$ replicas.

straight line in Fig. 3.5. This value is comparable to the optimal sample cost achieved with RENS using $M = 2$. Thus for this simple system, RENS is able to match the efficiency of REM, with fewer replicas.

The simple model we have borrowed [41] is well suited as an initial test case of our method: it exhibits the difficulties faced by REM for “large” ($n_p = 10$) systems, its equilibrium properties can be evaluated exactly (Fig. 3.2), and its efficiency can be computed with high statistical accuracy for many values of τ (Fig. 3.5).

3.3 Example 2: Alanine Dipeptide

In our second example, we apply RENS to a more relevant complex system in biology, the terminally-blocked alanine dipeptide (ACE ALA NME) in implicit solvent. Alanine dipeptide is also an appropriate test case for study with temperature RENS: While at lower temperatures AD tends to get stuck in locally-stable states, replica exchange attempts with configurations at higher temperatures can greatly enhance sampling. The alanine dipeptide (AD) system was simulated using the AMBER ff96 force field [65] with a GBSA implicit solvent model [97,114]. The evolution was governed by Langevin dynamics, with collision rate $\gamma_p = 1 \text{ ps}^{-1}$. For the Langevin integrator we chose a conservative timestep of 0.1 fs (see Sec. 2.5). Our simulations of AD were performed using the OpenMM molecular dynamics library [42]. The replica exchange implementation was facilitated through PyOpenMM ², the python wrapper for OpenMM, with modifications made to the existing replica exchange routine to incorporate features specific to RENS. In Appendix C we discuss these code modifications in more detail.

We proceed first by characterizing the metastable states of AD, from which we calculate across a range of RENS simulations a characteristic timescale associated with long-time configuration-space exploration of the system. From this we analyze the efficiency of these RENS runs.

3.3.1 Metastable states and correlation times

In biomolecules, fast femtosecond motion is typically dominated by bond vibrations, whereas large-scale changes such as rotation of one amino acid with respect to another exist on much longer timescales [94]. The metastable states of alanine dipeptide are cap-

²<http://simtk.org/home/pyopenmm>

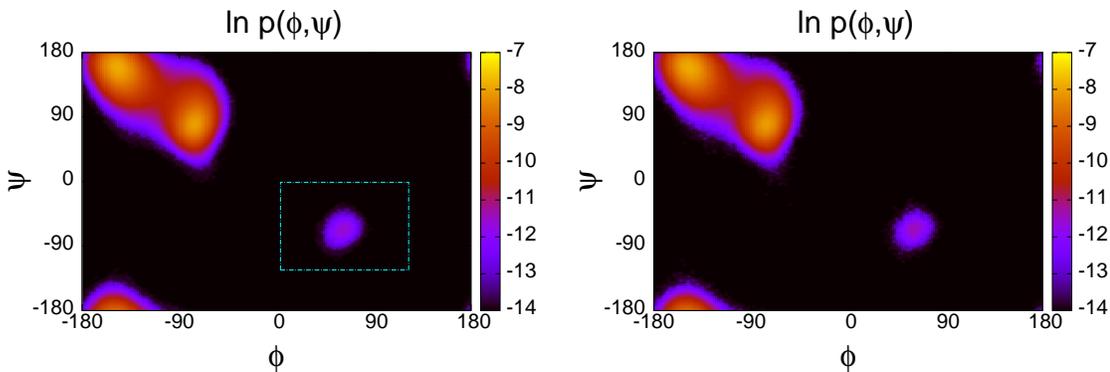


Figure 3.6: Ramachandran plot of dihedral angles ϕ and ψ for alanine dipeptide at $T_A = 300$ K. The plots were calculated from REM (left panel) and RENS (right panel) simulations, and are shown on a log scale. The boxed region in the left panel indicates conformer $C7_{ax}$.

tured by the the backbone torsional angles (ϕ, ψ) . We display in Fig. 3.6 Ramachandran plots of these angles for equilibrium data at $T = 300$ K. In the two plots generated from independent REM and RENS runs we find identical sampling – a consistency check that the two methods converge to the same result. From the Ramachandran plot we identify three primary regions: Two conformers, centered around $(-150, 150)$ and $(-75, 75)$, contain relatively equal populations and are dominant regions sampled. We identify a third state, conformer $C7_{ax}$, as all configurations for which $0 \leq \phi \leq 120$ and $-130 \leq \psi \leq 0$ (boxed region in Fig. 3.6). At 300 K, conformer $C7_{ax}$ has a small population, $p \approx 0.01$. Furthermore, under equilibrium Langevin dynamics, the system readily jumps between the first two regions approximately every 10 ps, whereas transitions into $C7_{ax}$ are observed much more rarely, on 100 ns timescales. Transitions between $C7_{ax}$ set the largest timescale in our system and provide the bottleneck to full system equilibration.

To assess the enhanced sampling of our method, we investigate the timescale associated with the $C7_{ax}$ transition in the lowest temperature replica, under various RENS

conditions. The output trajectory in this replica, consisting of the stitched-together equilibrium segments (solid red lines in Fig. 2.1) after having discarded the switching segments, samples the desired equilibrium distribution, and evolves under effective dynamics that are modified by the replica swap moves. To find a timescale associated with $C7_{\text{ax}}$, we first define an indicator function Y :

$$Y(x) = \begin{cases} 1 & \text{if in } C7_{\text{ax}} \\ 0 & \text{otherwise.} \end{cases} \quad (3.6)$$

Our output trajectory provides us with a time series of Y , from which we calculate a correlation time as

$$t_c = \frac{1}{\sigma^2} \int_0^\infty dt C(t). \quad (3.7)$$

Here σ^2 and $C(t)$ are the variance and time correlation function of Y calculated from the output trajectory, as in our previous example. The value of t_c quantifies the long-time correlations and the quality of equilibrium sampling generated by RENS. Below we proceed to characterize this across a number of RENS simulations.

In all of our RENS simulations of AD we consider $M = 2$ replicas, A and B , with $T_A < T_B$ and a common Hamiltonian H . We fix the highest temperature replica to $T_B = 600$ K, which is hot enough for the system to easily cross free-energy barriers. A number of different T_A values are considered, ranging from 160K to 300K. We considered three different switching times, $\tau = 0, 10$ and 100 fs, and implemented steplike switching protocols (see Fig. 3.10 and Sec. 3.3.3 below) in conjunction with the velocity rescaling procedure (Sec. 2.4). Between the switching intervals each replica sampled its equilibrium state for $\tau_{\text{eq}} = 100$ fs. In Fig. 3.7 we display t_c as a function of T_A , for the various values of τ . At a coarse level, we see that for large values of T_A , where the replicas begin with a good degree of overlap, RENS provides minimal enhancement over REM. With decreasing T_A ,

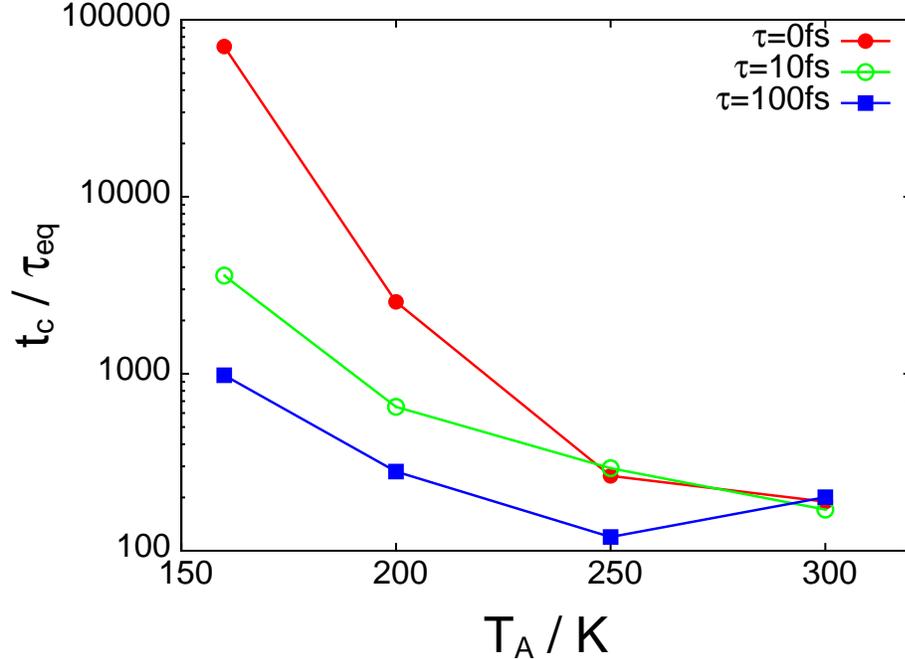


Figure 3.7: Calculated values of t_c , in units of swap attempts, for various values of T_A . The red, green and blue curves correspond, respectively, to switching protocols with $\tau = 0$, 10, and 100 fs.

the replicas move further apart, and the sampling enhancement increases with increasing τ .

To ensure good statistics, each value of t_c reported in Fig. 3.7 was calculated as an average over t_c estimates from many (50 to 75) independent RENS production runs, with the t_c values from each run estimated from Eq. 3.7 using the block averaging technique [39]. From the results of these individual runs, we found sampling to be overwhelmingly difficult for the simulations with $T_A = 160$ K, due to the very little overlap shared between \mathcal{R}_A and \mathcal{R}_B : In the $T_A = 160$ K replica only 40% of the 75 REM ($\tau = 0$) runs transitioned into $C7_{ax}$ region over the course of the simulation. When the switching time is increased to $\tau = 10$ fs, 75% of the simulations observed at least one transition, and at $\tau = 100$ fs, all simulations transitioned into $C7_{ax}$. Because it is impossible to estimate t_c if no transitions

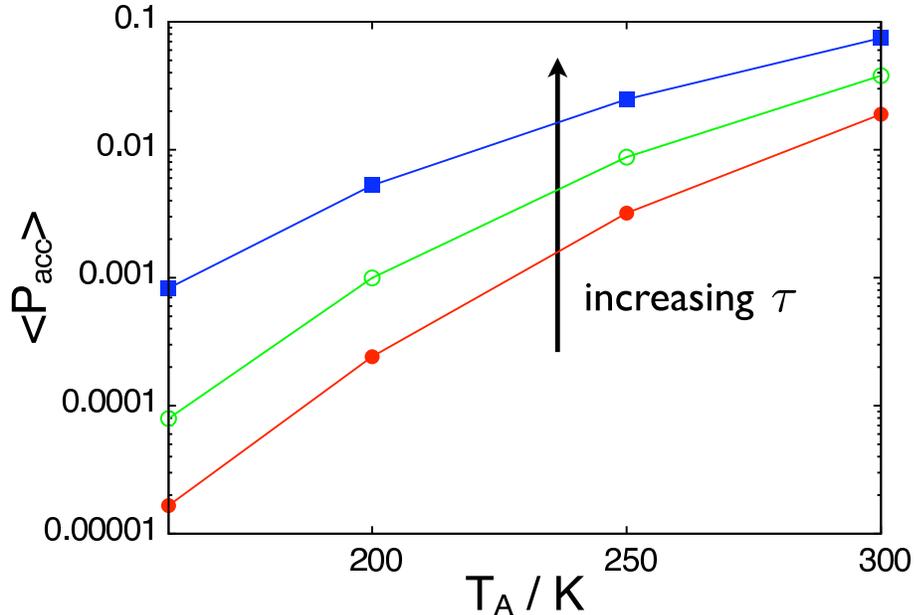


Figure 3.8: Average replica exchange acceptance probability as a function of T_A for various values of the switching time: $\tau = 0$ (filled red circles), 10 fs (open green circles) and 100 fs (filled blue squares). Increasing τ systematically increases $\langle P_{\text{acc}} \rangle$.

were observed, these simulations were not included when calculating the average values reported in Fig. 3.7. By neglecting these simulations we certainly underestimate the true value of t_c , and correspondingly also underestimate the sampling gain of RENS over REM in this low overlap regime.

The fluctuations of t_c estimates across the production runs also reveal interesting behavior. We observed at $T_A = 250 \text{ K}$ and 300 K a few outlying t_c values as large as three orders of magnitude that of typical values. We believe this long timescale results from a very rare situation in which replica A is found in region $C7_{\text{ax}}$ and replica B is simultaneously sampling a low probability region in the vicinity of $(-190, -120) \leq (\phi, \psi) \leq (-20, 0)$. In this scenario the replica exchanges are overwhelmingly unsuccessful because this region sampled by \mathcal{R}_B is extremely unfavorable in \mathcal{R}_A – causing A to be stuck

in $C7_{ax}$ until B exits the region. These outlying t_c values are not representative of the typical relaxation time during the vast majority of the simulation: For a given simulation, the probability of this scenario occurring, conditioned on A being in $C7_{ax}$, is roughly 0.05 on average, and occurs over the entire simulation with a probability of roughly 10^{-3} . As a result of this rare behavior, we neglect outlying t_c values from 3 simulations in our averages reported in Fig. 3.7.

Another indicator of the sampling quality of RENS is given by $\langle P_{acc} \rangle$, the fraction of accepted swaps, which we plot in Fig. 3.8. We see that for a given replica pair (T_A, T_B) , as τ increases, so too does the fraction of accepted swap attempts. As the equilibrium distributions of A and B move further apart with decreasing T_A , the acceptance rate of RENS relative to REM grows: At $T_A = 300$ K we see $\tau = 100$ fs provides about a ten-fold increase over REM; at $T_A = 160$ K this increase grows to a factor of 100.

3.3.2 Simulation efficiency

RENS provides a means to tune the overlap between neighboring replicas via the switching time. Given this added freedom, what is the optimal switching time? From the previous section we have seen that increasing τ leads to enhanced sampling through decreased correlations in the output trajectory (see Fig. 3.7). This enhanced sampling, however, is only achieved at the cost of the computational time needed to generate the switching trajectories, and to make a fair comparison between implementations, both of these competing factors must be taken into account.

Suppose a given RENS implementation provides a system correlation time t_c as defined above. For this implementation, let T_{sw}^{CPU} represent the CPU time per replica of performing a switching trajectory with the chosen protocol. Similarly, let T_{eq}^{CPU} denote

the per-replica CPU time associated with an equilibrium trajectory. In terms of these quantities, we define an efficiency measure

$$t^* = M \frac{t_c}{\tau_{\text{eq}}} \cdot (T_{\text{eq}}^{\text{CPU}} + T_{\text{sw}}^{\text{CPU}}) \quad (3.8)$$

as the total CPU time, summed over all M replicas, needed to generate an equilibrium trajectory of one correlation time in the primary replica. Although in Eq. 3.8 we assume uniform $T_{\text{sw}}^{\text{CPU}}$ across all replica pairs, a more general definition for t^* can also be made that relaxes this restriction. The measure t^* defined in Eq. 3.8 is analogous to the sample cost defined in Eq. 3.4 in our previous numerical example. In Table 3.1 we list the empirical values of $T_{\text{sw}}^{\text{CPU}}$ for all τ considered. We note that there is a finite contribution to the switching cost even for $\tau = 0$, due to the acceptance/rejection and velocity rescaling steps. With a fixed τ_{eq} across all RENS runs, the CPU contribution to the equilibrium sampling was calculated as $T_{\text{eq}}^{\text{CPU}} = 0.07$ s.

We are after a RENS protocol which minimizes t^* . Generally, an optimal protocol will depend upon many factors such as the physical system under consideration, replica spacings and other input parameters. As we have discussed before, when there is very little overlap between the equilibrium distributions of replicas A and B , REM ($\tau = 0$), there will be very infrequent swaps between replicas, and t^* will be very large due to a large t_c . As we devote more time to the switching segments, we expect a reduction in t_c due to an increase in swap acceptances, and a corresponding increase in efficiency. Beyond a certain point, however, the gain in sampling due to increased replica overlap is outweighed by the large overhead cost, $T_{\text{sw}}^{\text{CPU}}$, associated with generating the swap configurations, and t^* rises. A protocol which minimizes t^* should lie somewhere in between these two regimes.

We are interested in finding the regimes in which it is advantageous to use RENS. In Fig. 3.9 we plot t^* as a function of T_A for the same RENS simulations as in Fig. 3.7. At

τ (fs)	T_{sw}^{CPU} (clock seconds)
0	0.048
10	0.093
100	0.501

Table 3.1: Empirical CPU times needed for switching trajectories.

large values of T_A we find REM ($\tau = 0$) to give the lowest value of t^* . As T_A decreases, the gain in RENS grows relative to REM, with a crossover occurring between 250 and 200 K. At 160 K REM performs poorly compared to RENS. As in Sec. 3.3.1, we expect the efficiency gain of RENS to be underestimated in the low overlap regime, since the $\tau = 0$ and $\tau = 10$ fs results for $T_A = 160$ K reported in Figs. 3.7 and 3.9 contain estimates from only 40% and 75% of the runs, respectively, due to a number of simulations failing to produce a single transition within the simulation time. Interestingly, the protocol that seems to be most efficient overall is the intermediate run of $\tau = 10$ fs.

We note that although our simulations were performed with a conservative timestep of $\delta t = 0.1$ fs, had we used a standard $\delta t = 1$ fs our efficiency profile would to a good approximation remain unchanged (apart from an overall scaling factor). This is because the vast majority of T_{sw}^{CPU} is due to processes associated with updating λ (see Sec. 3.3.3 below), which is a property of the switching protocol alone and independent of the timestep.

As shown in Fig. 3.9, the advantage of RENS over REM increases with decreasing overlap (decreasing T_A). For the case of very little overlap, REM would dictate adding more replicas to enhance sampling. To test this against our method, we calculated t^* for a REM simulation with 5 replicas (with a geometric progression of temperatures [64]), with

the lowest temperature replica $T_A = 160$ K. Comparing this result to the corresponding RENS run containing 2 replicas at $\tau = 100$ fs, we find REM outperforms RENS by roughly a factor of 80. This result, which is contrary to our previous one-dimensional example (Sec. 3.2) where RENS matched the efficiency of REM, is due to the fact that the software used for our AD simulations severely restricted the type of switching protocols performed. We expect that RENS could become much more efficient if the software was more amenable to generation of nonequilibrium trajectories. We proceed to explain this in more detail below.

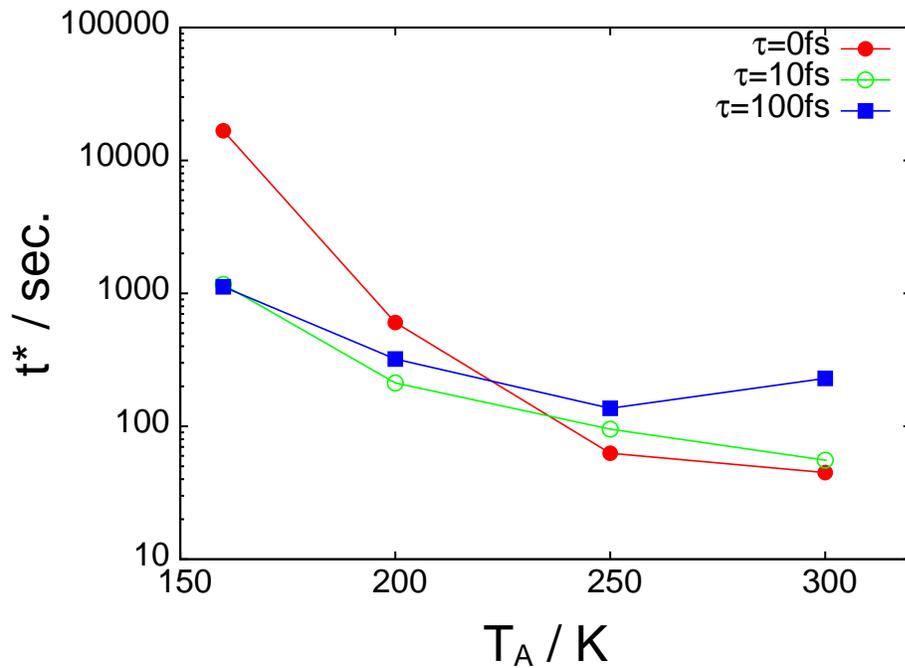


Figure 3.9: The CPU cost, t^* , for the simulations of Fig. 3.7.

3.3.3 Switching protocols and CPU costs

The overhead cost T_{sw}^{CPU} associated with the switching simulations can be quite expensive depending on the particular implementation of the switching protocol $\lambda(t)$. For

the software used for our AD simulations, updates of the control parameter $\lambda \rightarrow \lambda + d\lambda$ are a particularly costly step. Because we do not have easy access to the base code, each incremental λ update must be implemented at an outer code level, requiring the system integrator to be reinitialized every time h_λ is updated, a costly step computationally. Another expensive contributor to the λ updates were the velocity rescaling steps (introduced in Sec. 2.4), as without direct access to the integrator the velocities need to be called, modified and read back in at an outer level of the code multiple times during the course of a single switching trajectory, adding much overhead. Because of these overhead costs, it was necessary to implement steplike protocols, illustrated by Fig. 3.10, in which updates of λ are separated by relaxation intervals of duration Δt at fixed λ , with $\Delta t \gg \delta t$. Choosing $\Delta t = 10$ fs was sufficiently large such that the number of λ updates, $n = \tau/\Delta t + 1$, remained small for all τ considered and kept the overhead costs $T_{\text{sw}}^{\text{CPU}}$ relatively small. A switching protocol containing a λ update after every integration timestep ($\Delta t = \delta t$) – standard in free energy simulations [56, 68, 128] – is for our AD simulations computationally unfeasible. In Appendix C we discuss in more detail the implementation of switching simulations in the OpenMM code.

While the steplike protocol circumvents a large switching overhead, this comes at the cost of reduced replica overlap. Large jumps of the parameter λ disturb the system fairly strongly compared to linear protocols, leading to larger dissipation and reduced number of accepted replica exchange moves. Had we been able to implement linear protocols, we expect that RENS would experience greater overall acceptance rates and would perform as efficiently as REM, as we saw in our previous test system.

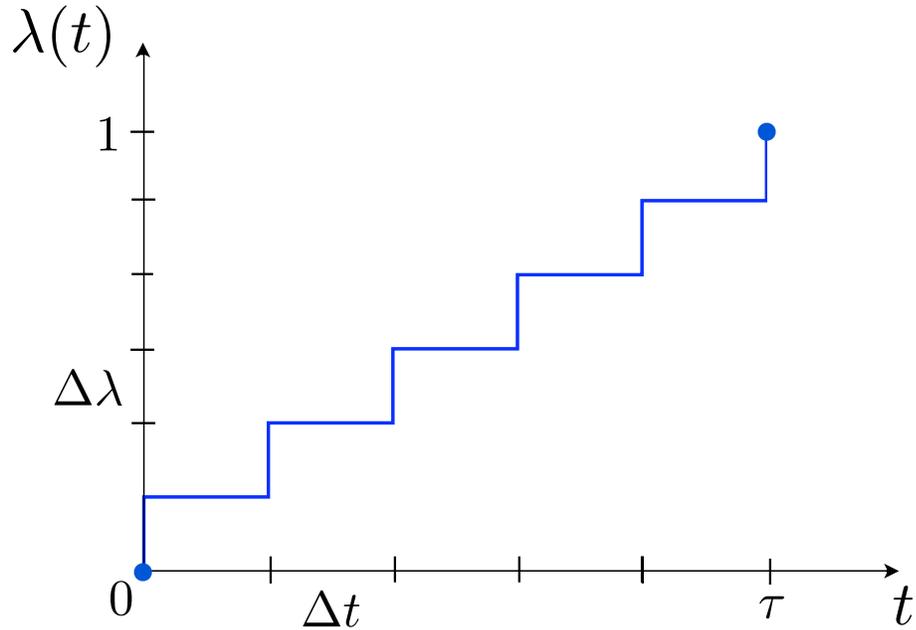


Figure 3.10: Discretized switching protocol: The evolution of the system out of equilibrium proceeds via instantaneous changes $\Delta\lambda$ in the control parameter, followed by relaxation steps of duration Δt . Here Δt is much larger than the integration timestep δt .

3.4 Discussion

Our method of Replica Exchange with Nonequilibrium Switches enhances equilibrium sampling through switching simulations which guide neighboring replicas towards each other in phase space. By devoting extra computer resources to the switching simulations, we can increase an effective overlap between replicas, increasing the average swap acceptance rate and foregoing the need for additional replicas (see Fig. 1.2).

In this chapter we tested RENS on two model systems, a simple 1D example, and alanine dipeptide. Our results confirm that increasing the switching time 1) increases the average swap acceptance between replicas, and 2) provides equilibrium data that is less correlated. We calculated the efficiency for RENS simulations that takes into account the

CPU time associated with the switching simulations. For simulations with a restricted number of replicas, M , we find that the benefit of RENS over REM grows with shrinking overlap.

We also tested RENS against REM simulations that were optimized by allowing M to vary. We found for the simple 1D simulations that RENS can match the efficiency of REM, while using fewer replicas; the simulations of alanine dipeptide, however, were unable to perform as well as optimized REM. This is due to the fact that the code alterations that performed the switching simulations were quite costly computationally, due to limited access to the base level code. Had this software been more amenable to generation of nonequilibrium trajectories, the efficiency of RENS relative to REM would be closer to what we observed for the 1D case. This remains to be tested and will be the subject of future studies.

As discussed in Chapter 1, when applying REM to a problem of interest, the phase space overlap requirement dictates a minimum number of replicas, M^* , needed to achieve a reasonable swap acceptance frequency. With RENS, the switching simulations have the effect of increasing phase space overlap, thus allowing for fewer replicas, $M < M^*$. There are several reasons why one might wish to exploit this flexibility.

(i) Most obviously, if we perform simulations using a cluster of P processors, then RENS allows us to assign one replica per processor – the easiest and most natural (and traditional) allocation – even if $P < M^*$.

(ii) It is often useful to picture replica exchange as a diffusion process in which trajectories hop randomly along the chain $\mathcal{R}_1, \dots, \mathcal{R}_M$. In this picture, $\sim M^2$ successful swaps are needed for a given trajectory to accomplish an entire transit between \mathcal{R}_1 to \mathcal{R}_M . Thus using fewer replicas (with RENS) can significantly reduce the cost of inter-processor

communication associated with attempted configuration swaps.

(iii) REM is often implemented synchronously: swaps are attempted only after every replica completes a pre-determined duration of equilibrium sampling. With one replica per processor, this can be highly inefficient, limited by the speed of the slowest processor. RENS lends itself naturally to asynchronous implementation. A master process initiates work simulations in a randomly chosen replica pair, while the remaining replicas, unaffected, continue sampling.

(iv) With any replica exchange strategy, there are parameters we adjust to optimize efficiency, such as the number of replicas, M , and the choice of intermediate temperatures or Hamiltonians. It is potentially very useful to improve efficiency *adaptively*, during the actual production run [117]. RENS offers a relatively painless way to accomplish this, namely by adjusting the durations of the switching simulations. E.g. if it is observed that a low P_{acc} between \mathcal{R}_n and \mathcal{R}_{n+1} poses a bottleneck for efficient sampling, then the switching time for that replica pair, $\tau_{n,n+1}$, can be increased.

(v) To this point we have treated the data generated during the switching simulations as “junk”, to be discarded after the attempted configuration swap. However, by a trick of statistical reweighting one can scavenge equilibrium information from such nonequilibrium trajectories (see Eq. 4 of Ref. [53]). This suggests that the switching simulations themselves can contribute to the equilibrium sampling in each replica, thus increasing the efficiency of RENS. For Monte Carlo sampling, Frenkel [40] has developed an analogous, thrifty algorithm that relies on the “waste-recycling” of otherwise rejected trial moves.

With the development of RENS and similar methods [6, 18, 83–85, 89, 113], it is important to further characterize the efficiency gains nonequilibrium simulations can bring to equilibrium sampling. As we have seen from our two examples, the efficiency gains

RENS can offer is dependent upon how efficiently the switching simulations can be implemented in the computer code. We suspect our method could become much more beneficial if one has predefined routines to perform the switching simulations which update time-dependent parameters at the integration level in the code. Such preprogrammed routines would significantly decrease the CPU time needed to perform switching simulations, and in doing so provide more natural implementation of nonequilibrium sampling methods such as RENS in standard simulation packages.

Small molecule solvation

4.1 Introduction

The solvation free energy, the reversible work required to move a solute from the gas phase into solution, is a key concept in thermodynamics and important for the understanding of thermodynamic systems in chemistry and biology. Research efforts in computational thermodynamics have long strived for efficient calculation of solvation free energies, with applications ranging from phase coexistence [41] to hydrophobicity [25]. To begin we consider a system composed of a single solute immersed in an explicit solvent environment. The interaction energy which describes this system

$$U(x, y) = U_u(x) + U_v(y) + V_{\text{int}}(x, y) \quad (4.1)$$

is composed of solute-solute interactions, U_u , solvent-solvent interactions U_v , and interactions between the solute and solvent, V_{int} . Here x denotes the configuration of the solute, and y the configuration of the solvent atoms, such that $z = (x, y)$ is a point in the configuration space of the complete system. The solvation free energy of our solute, ΔF_{solv} , is the free energy change associated with “turning on” the the interactions V_{int} between the

solute and solvent. In the standard definition of ΔF_{solv} , we have

$$e^{-\beta\Delta F_{\text{solv}}} = \frac{Z}{Z_0} = \frac{\int dx \int dy e^{-\beta[U_u(x)+U_v(y)+V_{\text{int}}(x,y)]}}{\int dx \int dy e^{-\beta[U_u(x)+U_v(y)]}}, \quad (4.2)$$

where Z is the partition function of the fully interacting system, and Z_0 is the partition function of the system in which the solute and solvent are free to move but do not interact with each other. We can alternatively conceive of a solvation free energy *where the solute is frozen* in a configuration x :

$$e^{-\beta\Delta F_{\text{solv}}(x)} = \frac{Z(x)}{Z_0^v} = \frac{\int dy e^{-\beta[U_v(y)+V_{\text{int}}(x,y)]}}{\int dy e^{-\beta U_v(y)}}. \quad (4.3)$$

Here $\Delta F_{\text{solv}}(x)$ is defined as the reversible work required to move the solvent molecule – constrained to position x – from the gas phase into the solution. In Eq. 4.3 the partition function $Z(x)$ corresponds to the frozen solute interacting with the solvent, and Z_0^v is the partition function of the solvent in absence of the solute. While the solvation free energy of the unconstrained solute (Eq. 4.2) is more relevant for describing the thermodynamic properties of the solute as a whole, the free energy of the constrained solute (Eq. 4.3) can provide insight into the microscopic details of the solute which promote or hinder solvation (such as hydrophobicity of the solute surface). Such configuration-dependent solvation free energies are a central concept for implicit solvation theories [105], and accurate determination of $\Delta F_{\text{solv}}(x)$ could assist in the development of such models.¹ In this study we will focus on the calculation of configuration-dependent solvation free energies described by Eq. 4.3. For notational convenience we will suppress the dependence of $\Delta F_{\text{solv}}(x)$ on x in the discussion below.

In this chapter we propose a RENS method for calculation of solvation free energies.

In contrast to our previous studies, where RENS was used primarily to overcome poor

¹We note that one can also conceive of constrained solvation studies in which some collective variable of the solute (for instance a dihedral angle) is frozen instead of an entire configuration.

sampling in complex systems, here method is designed for the efficient estimation of ΔF_{solv} . Because we will be considering solvation of frozen solutes, effective sampling of the solute degrees of freedom is not a concern. However, the RENS solvation simulations we describe below can also be combined with other RENS methods (such as temperature RENS) such that estimation of ΔF_{solv} of an unconstrained solute can be achieved in a reliable and efficient manner.

To illustrate how RENS simulations can be used for solvation free energy calculation, we first consider a Hamiltonian RENS simulation, whereby replicas are defined by *unique Hamiltonians* and a uniform temperature. In a naïve implementation with two replicas, let replica A be comprised of the noninteracting solute and solvent, $U_A = U_u + U_v$, and replica B the fully-interacting system, $U_B = U_u + U_v + V_{\text{int}}$. Here the switching simulations between A and B are performed by switching U between U_A and U_B , which in replica A amounts to switching on the solute-solvent interactions V_{int} . In replica B the switching simulations correspond to turning off V_{int} . The work values performed during these switching simulations provide an estimate of $\Delta F_{\text{solv}} = F_B - F_A$ through the nonequilibrium work relation discussed in Chapter 1: from the work values w_A generated in replica A , for instance, we have

$$e^{-\beta\Delta F_{\text{solv}}} \approx \frac{1}{n} \sum_{i=1}^n e^{-w_A^i}. \quad (4.4a)$$

Similarly, the set of w_B values gives us an estimator

$$e^{+\beta\Delta F_{\text{solv}}} \approx \frac{1}{n} \sum_{i=1}^n e^{-w_B^i}. \quad (4.4b)$$

If our switching simulations are performed instantaneously, $\tau = 0$, the free energy estimators from the nonequilibrium work relation coincide with Widom's particle insertion method [127]. In the opposite limit as $\tau \rightarrow \infty$, our estimator resembles thermodynamic

integration [63].

While this simple RENS implementation, together with standard free energy estimators, can in principle be used to estimate ΔF_{solv} , the convergence of this estimate can be somewhat problematic. If during the switching simulations in replica A the interactions are turned on very quickly, we expect the work values to be quite large on average, since typical configurations sampled from A will have the solute sitting right on top of the solvent. The dominant contributors to the exponential average, however, will come from the small work values which are very rare, and can lead to large errors associated with ΔF_{solv} estimates. However, if the interactions are turned on more slowly this dissipation should decrease, as the surrounding solvent has more time to equilibrate to the presence of the solute.

In general we can address large dissipation which plagues free energy estimation by ensuring good replica overlap. Indeed, Vaikuntanathan and Jarzynski have shown that this dissipation is related to a lack of overlap between the current state of the system and its corresponding equilibrium state [120]. This implies that by increasing effective overlap between our replicas – through altering the switching time or by other means – we can reduce dissipation and hence reduce the number of samples needed for convergence of ΔF_{solv} .

The method we propose for estimation of ΔF_{solv} directly addresses this overlap problem by embedding solute solvation within a larger thermodynamic cycle involving cavity formation. The cycle, depicted in Fig. 4.1, includes a progression of replicas representing growth of a cavity in the pure solvent, followed by a turning-on of solvent-solute interactions in presence of the cavity, and culminating in the shrinking of the cavity in presence of the solute-solvent interactions. As the initial and final replicas correspond to the un-

solvated and completely solvated states, the quantity ΔF_{solv} is recovered by summing up the free energy changes between all intermediate replica pairs.

The thermodynamic cycle involving cavity growth used for calculation of ΔF_{solv} is adapted from the ideas of quasichemical theory (QCT) formalism [10] developed by Pratt and others. In QCT, the solvation free energy ΔF_{solv} is decomposed into three contributions related to the solvent statistics in proximity to the solute. Specifically,

$$\Delta F_{\text{solv}} = k_B T \ln x_0 - k_B T \ln p_0 + \Delta F_{\text{int}}. \quad (4.5)$$

Here x_0 and p_0 are probabilities that the system is spontaneously void of solvent within a given radius of the solute. The probability p_0 is calculated in the noninteracting solvent-solute ensemble, and x_0 is calculated in the ensemble where the solute-solvent interactions are present. The contribution $-k_B T \ln p_0$ to Eq. 4.5 can then be interpreted as the free energy associated with forming a hard sphere cavity *in absence of the solute*; similarly, $k_B T \ln x_0$ is the free energy associated with shrinking the hard sphere cavity *in presence of the solute*. The final term ΔF_{int} is the free energy associated with inserting the solute into the solvent-void cavity. (See e.g. Eq. 6 and Sec. II of Ref. [90] for further discussion.) QCT calculations of solvation free energies have estimated the various components in Eq. 4.5 through standard methods: ΔF_{solv} is calculated by Widom insertion [127]; the probabilities p_0 and x_0 are each estimated from solvent statistics obtained from a corresponding MD or MC simulation [5, 90, 108]. In our RENS method we incorporate this thermodynamic cycle into a replica exchange simulation. However, in contrast to QCT, our method involves cavity formation from non hard-sphere potentials (see Eq. 4.10). Because of this, our calculation of ΔF_{solv} , while similar in spirit to the QCT calculation, is not identical (see Eq. 4.12). Our RENS method also naturally incorporates cavity growth / shrinking from nonequilibrium switching simulations into the calculation of the various

free energy differences.

By incorporating this thermodynamic cycle within the framework of RENS, we can use the tools and intuition for increasing replica overlap in order to efficiently estimate free energy differences. In particular, this method addresses the large dissipation associated with switching on solute-solvent interactions, as now the solute is inserted in the center of a cavity void of solvent. For a large enough cavity size R , we expect the solvent density to be largely unaffected by turning on the solute since all interactions are long-ranged – resulting in a good overlap between replicas. On the other hand, for small values of R we expect very little overlap, since the unsolvated system typically contains solvent molecules that protrude into the region occupied by the solute, which is extremely rare in the solvated system. We can then think of the cavity size R as an additional parameter to tune the overlap between the unsolvated and solvated systems. Whereas in the RENS investigations up until now our strategy for enhancing replica overlap has been by altering the switching time, now we take a hybrid approach, where greater overlap is achieved through τ as well as the cavity size R .

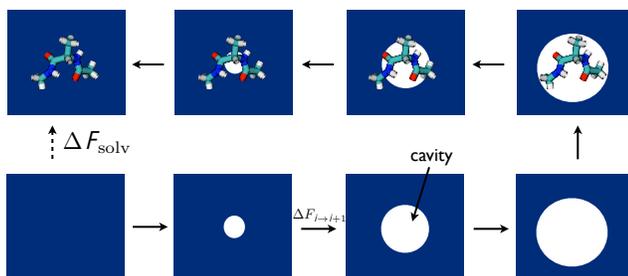


Figure 4.1: RENS solvation free energy calculations assisted by cavity growth: Replicas progress from the fully noninteracting to fully interacting system through a series of solvent cavities of growing size. Switching on solute-solvent interactions (rightmost pair of replicas) becomes easier in the presence of a large cavity.

4.2 RENS implementation

We investigate solute solvation through Hamiltonian RENS consisting of M replicas described in Fig. 4.1. For replicas $i < M/2$, the solute-solvent interactions are turned off (bottom set of replicas in Fig. 4.1), and the interaction energy is

$$U^i(x, y) = U_u(x) + U_v(y) + V_{\text{cav}}(y; R_i). \quad (4.6a)$$

Similarly, replicas $i \geq M/2$ contain the interacting solute-solvent system (top half), with

$$U^i(x, y) = U_u(x) + U_v(y) + V_{\text{int}}(x, y) + V_{\text{cav}}(y; R_i). \quad (4.6b)$$

Here the cavity of replica i is realized through an externally-applied potential $V_{\text{cav}}(y; R_i)$ acting on the solvent configuration y alone (See Eq. 4.10 below), with a size specified by R_i . Although our specific simulations below consider of a frozen solute, where $U_v(x) = \text{const.}$, we keep the x -dependence in Eq. 4.6 above for generality.

Because our replicas now contain unique potentials but the same temperature, $h_i = \beta H_i$, the switching simulations between replica pair $(i, i + 1)$ evolves each replica under a time-dependent potential $U_{\lambda(t)}$ which interpolates between U^i and U^{i+1} through control parameter λ in switching time τ . For the switching simulations within the noninteracting or interacting replicas, the cavity is grown / shrunk (see Fig. 4.2) as $R(t) = R_i + \lambda(t)(R_{i+1} - R_i)$. Because λ only couples to U through R , the reduced work (Eq. 2.28) is given by incremental changes in V_{cav} upon changes in λ . In replica i we have

$$w_{i \rightarrow i+1} = \beta \int_0^\tau dt \dot{\lambda}_{i \rightarrow i+1} \frac{\partial V_{\text{cav}}(y_t; R_i(\lambda_{i \rightarrow i+1}))}{\partial \lambda_{i \rightarrow i+1}} \quad (4.7a)$$

Similarly, in replica $i + 1$, the work is

$$w_{i \leftarrow i+1} = \beta \int_0^\tau dt \dot{\lambda}_{i \leftarrow i+1} \frac{\partial V_{\text{cav}}(y_t; R_i(\lambda_{i \leftarrow i+1}))}{\partial \lambda_{i \leftarrow i+1}}. \quad (4.7b)$$

As usual $\dot{\lambda}_{i \rightarrow i+1} = -\dot{\lambda}_{i \leftarrow i+1}$.

The switching simulations performed between the middle set of replicas involve the turning on/off of the interaction term V_{int} . Here $\lambda(t)$ specifies the time-dependence of interaction potentials between all solute-solvent atom pairs, and the work performed in replica $A = M/2 - 1$ is

$$w_{A \rightarrow B} = \beta \int_0^\tau dt \dot{\lambda}_{A \rightarrow B} \frac{\partial V_{\text{int}}(y_t; \lambda_{A \rightarrow B})}{\partial \lambda_{A \rightarrow B}}, \quad (4.8a)$$

and in replica $B = M/2$, we have

$$w_{B \leftarrow A} = \beta \int_0^\tau dt \dot{\lambda}_{A \leftarrow B} \frac{\partial V_{\text{int}}(y_t; \lambda_{A \leftarrow B})}{\partial \lambda_{A \leftarrow B}}. \quad (4.8b)$$

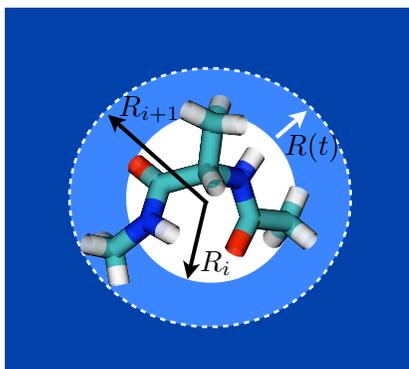


Figure 4.2: Cavity growth simulations: During the switching simulations between replicas i and $i + 1$ the cavity size R is grown / shrunk at a finite rate.

4.3 Solute of study: Monoethanolamine

We test our method on the solvation of monoethanolamine (MEA), an organic solute ($\text{NH}_2\text{-(CH}_2\text{)}_2\text{-OH}$) consisting of a single ethanol attached to an amine group, displayed in Fig. 4.7. MEA and other ethanolamine molecules are commonly used in industry

for the manufacturing of detergents, cosmetic creams, cooling lubricants and industrial cements [4]. More recently MEA has been gaining interest for its potential use in carbon scrubbing, where aqueous solutions of ethanolamines can be used for absorption of CO₂ and other acidic gases [4] released by power plants, making it an important potential cleaning agent for greenhouse gases emissions [1, 98]. Recent computational studies of ethanolamines have investigated thermodynamic properties aqueous solutions of MEA [72] as well as pure MEA liquids [20].

In our simulations we considered a system composed of a single MEA molecule in a solvent environment of $N = 421$ waters. Our simulations were performed in the canonical NVT ensemble with temperature $T = 298K$. The system volume V was chosen such that the solvent density in absence of the solute corresponded to standard temperature and pressure conditions: $\rho_{solv} = 1g/mL$. The solute interactions were determined by the GAFF force field (identical to Ref. [20]), with harmonic bond and angle terms, as well as dihedral contributions. The water molecules were described by the rigid TIP3P model. Intermolecular interaction energies consisted of Coulomb and Lennard-Jones terms,

$$U_{ij}(r_{ij}) = \frac{q_i q_j}{4\pi\epsilon_0 r_{ij}} + 4\epsilon_{ij} \left[\left(\frac{\sigma_{ij}}{r_{ij}} \right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}} \right)^6 \right] \quad (4.9)$$

with charges q_i, q_j , Lennard-Jones parameters σ_{ij} and ϵ_{ij} for atomic pair i, j specified by the GAFF force field.

4.4 Cavity formation

The solvation of MEA with our method requires the formation of a solvent cavity in which MEA is to be inserted. To form the cavity we apply a spherically-symmetric

repulsive potential that acts independently on each solvent oxygen as

$$V_{\text{cav}}(r; R) = E_0 e^{-(r-\alpha R)/\sigma} \quad (4.10)$$

(depicted in Fig. 4.3.) Here r is the distance of the given oxygen relative to the center of mass of the solute, with $\sigma = 4$ nm, $E_0 = 0.238$ kcal/ mol, and $\alpha = 2$. The parameter R controls the size of the cavity whose value varies across replicas. Calculations of solvation free energies based on QCT have also implemented non hard-sphere cavity potentials [102].

In Fig. 4.3 we show solvent profiles obtained from simulations performed in presence of V_{cav} , for various values of R . We plot the solvent density $\rho(r)$ as a function of r , relative to the bulk density $\rho_0(r) = 4\pi r^2 N/V$ of a uniform solvent system (in absence of V_{cav}). Simulations in absence of the solute (lower left panel) show the solvent distributed relatively uniformly for small values of R . For larger values of R we see solvent depletion near the origin, which widens with increasing R . (We note that for small values of R we see some solvent structure for $r < 3$ Å, caused by the presence of the repulsive potential even for $R = 0$ Å.) For simulations in the presence of the solute (lower right panel) the solvent behaves differently: in contrast to the pure solvent case, here the solute prevents the solvent from reaching inside $r \approx 3$ Å. The effect of R is to control the solvent density at contact. For small R we see a pronounced peak around $r = 4$ Å, which is diminished with increasing R . For $R > 6$ Å, however, the cavity boundary is also shifted, as is clearly seen by the rightmost curve at $R = 10$ Å. For reference we plot in the upper panel of Fig. 4.3 the cavity potential $V_{\text{cav}}(r; R)$ for various cavity sizes.

Now that we understand quantitatively the effect of R on the unsolvated and solvated systems, we can choose a prudent value of $R = R^*$ at which to insert the solute. On the one hand, R^* should be large enough such that there is a decent overlap between the solvated and unsolvated replicas; on the other hand, it should be small enough such that

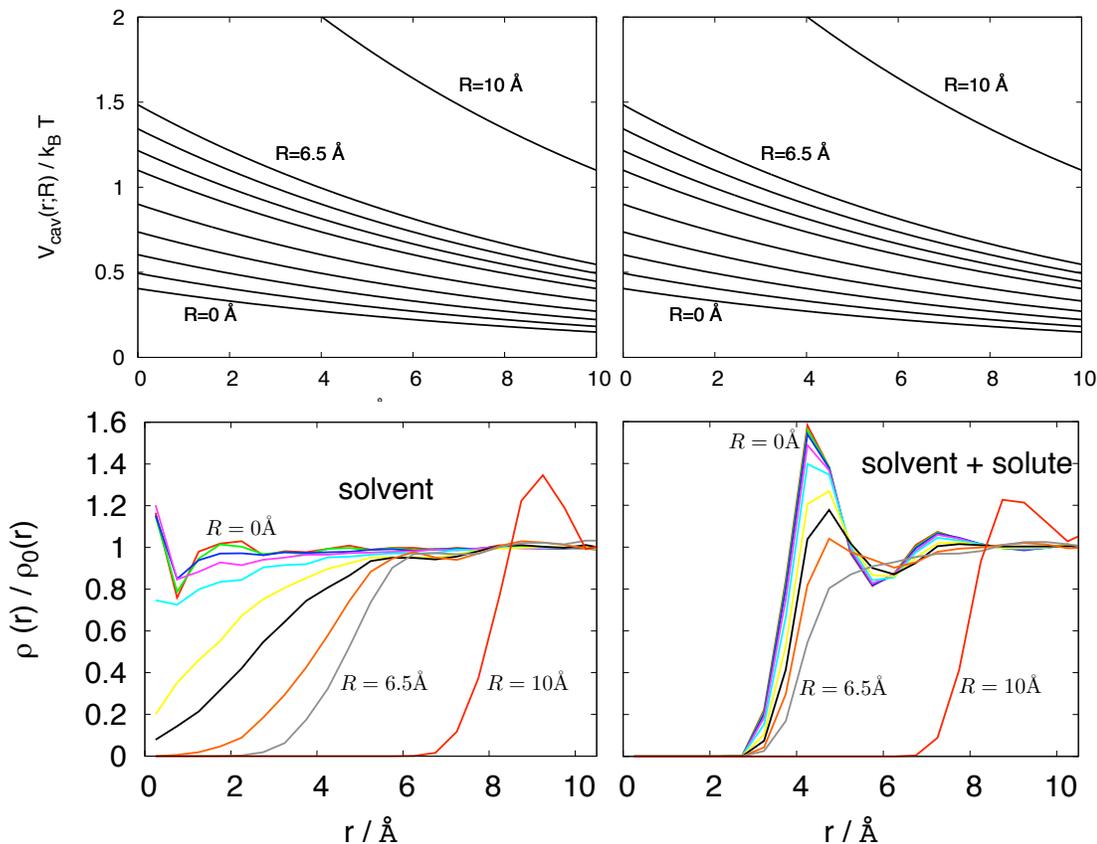


Figure 4.3: The solvent density profile. *Lower Panels:* Solvent density profile relative to bulk, as a function of the distance r from the center of the cavity. The noninteracting (lower left panel) and interacting (lower right panel) solute-solvent systems behave differently for small R . For $R > 6\text{\AA}$, however, they show similar behavior. *Upper Panels:* The external potential $V_{\text{cav}}(r; R)$ (Eq. 4.10) for various cavity sizes R .

we are not simulating an unnecessary number of replicas. Comparison of the upper and lower panels in Fig 4.3 shows that for small R the solvent behaves quite differently when in the presence of the solute. At $R = 6.5\text{\AA}$ (in grey) however, we start to see similar solvent structure between these two distributions, indicating a growing overlap. At $R = 10\text{\AA}$ (rightmost red curve), we see very good agreement between solvated and unsolvated systems. However, a RENS implementation with a cavity so large would require a very large number of replicas and as such would be computationally expensive and unnecessary.

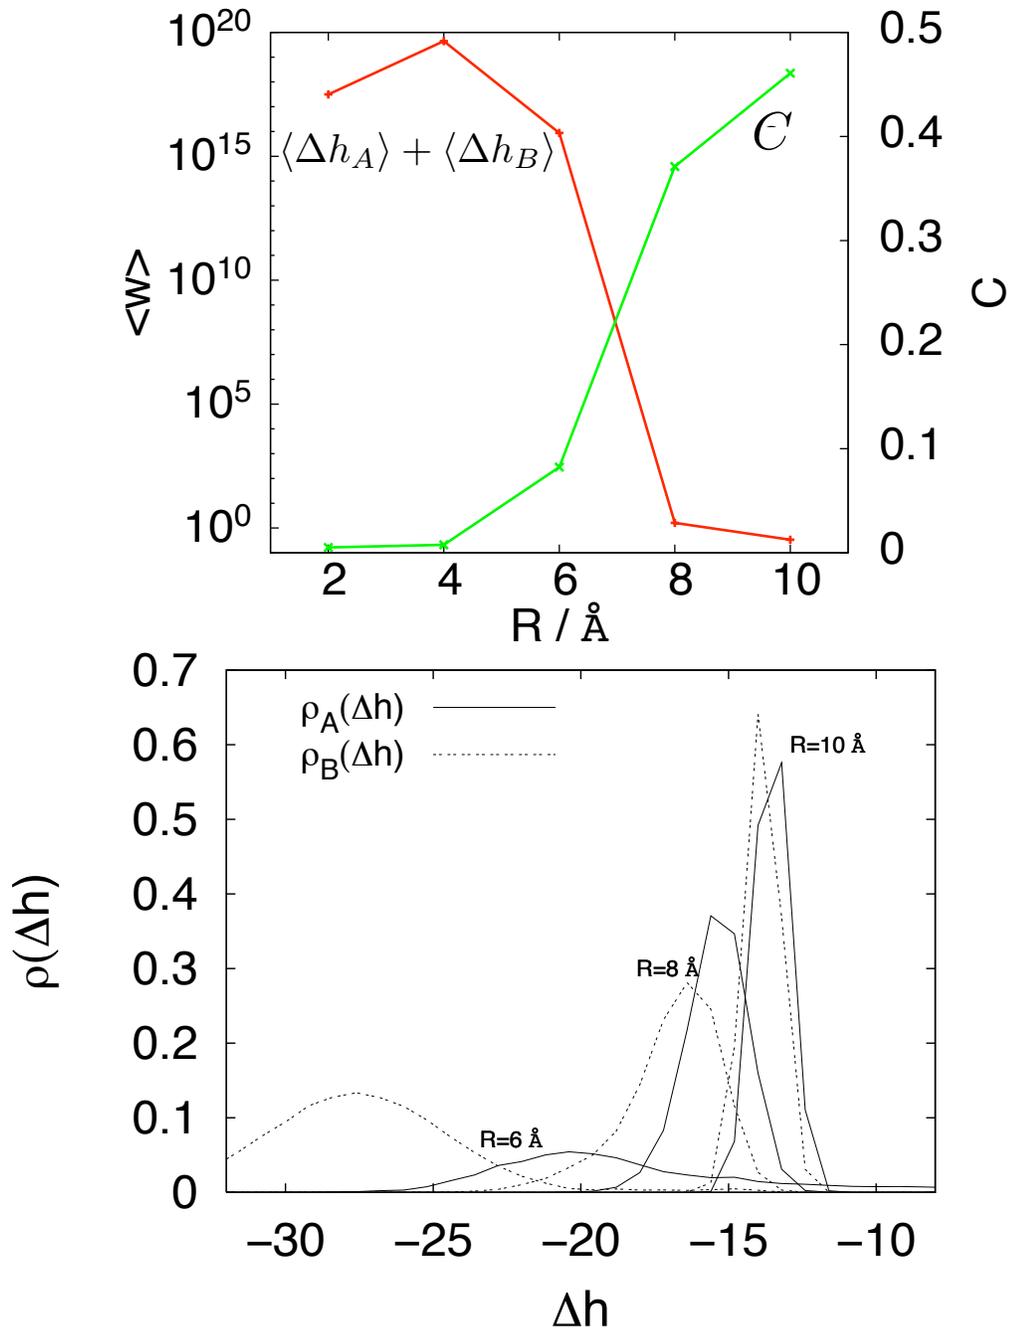


Figure 4.4: Replica overlap associated with instantaneous solute insertion and deletion for various cavity sizes. *Upper Panel:* The average total work and overlap measure C as a function of cavity size R . *Lower Panel:* Work distributions of conjugate insertion (solid lines) and deletion (dashed lines) switching simulations. Replica overlap increases with R .

The phase space overlap requirement between the solvated and unsolvated systems can be probed more directly by investigation of the overlap between distributions of *work values*. If the work distributions $\rho_A(w)$ and $\rho_B(-w)$ associated with solute insertion (A) and deletion (B) share a large overlap, then $\langle w_{tot} \rangle = \langle w_A + w_B \rangle \approx 0$ guarantees a good swap acceptance rate – and good phase space overlap. If, however, there is little overlap, then $\langle w_{tot} \rangle \gg 0$, most swaps are rejected and there is little phase space overlap. To find a satisfactory value of R^* we performed RENS simulations on pairs of replicas involving solute insertion / deletion, for various values of R . The work values associated with solute insertion / deletion were calculated using Eq. 4.8. To measure the overlap we calculate

$$C = \int dw \frac{\rho_A(w)\rho_B(-w)}{\rho_A(w) + \rho_B(-w)}. \quad (4.11)$$

The quantity $C \geq 0$ is a statistical measure of the distinguishability between distributions ρ_A and ρ_B [13, 121]. When ρ_A and ρ_B share no overlap, $C = 0$; in the limit where the distributions are identical, C obtains its maximal value of $1/2$. In Fig. 4.4 we plot both C and $\langle w_{tot} \rangle$ as a function of R for RENS simulations involving solute insertion / deletion. The turning on / off of the cavity potential V_{cav} was performed instantaneously ($\tau = 0$) such that $w_A(x) = \Delta h_A(x) = h_B(x) - h_A(x)$ and $w_B(y) = \Delta h_B(y) = h_A(y) - h_B(y)$. For small cavity sizes, we see that $\langle w_{tot} \rangle \gg 0$ and $C \approx 0$, indicating very little overlap between A and B . With increasing R the value of C increases towards its maximum of $1/2$ and $\langle w_{tot} \rangle \approx 0$. In the lower panel of Fig. 4.4, plots of the distributions $\rho_A(w)$ and $\rho_B(-w)$ visually confirm an increasing overlap with increasing R . In light of Figs. 4.3 and 4.4 we choose a cavity size $R^* = 6.5 \text{ \AA}$ for solute insertion, roughly the point at which C takes half its maximal value.

4.5 RENS simulation setup

Our RENS production runs for calculation of ΔF_{solv} consisted of $M = 19$ replicas, 9 of which containing the solvent alone, and 10 containing the solvent-solute system. The R values were chosen roughly uniformly between 0 and R^* , for both the noninteracting and interacting replicas.

The switching simulations, initiated every $\tau_{\text{eq}} = 100$ fs, were implemented simultaneously on multiple replica replica pairs, as follows. At the onset of a switching simulation, neighboring replicas were paired together either with an odd low replica, $(1, 2), (3, 4), \dots$, or an even low one $(2, 3), (4, 5) \dots$, with equal probability. While these replicas were switching, the remaining unpaired replicas (replicas 1 and/or M) underwent equilibrium sampling. For the switching simulations involving cavity growth we chose a switching time $\tau = 10$ fs. The solute insertion switching simulations were performed instantaneously, $\tau = 0$ fs. During the switching simulations the system evolved under Langevin dynamics with a switching rate $\dot{\lambda} = 1/\tau$ and a damping rate $\gamma = 50 \text{ ps}^{-1}$. The simulation timestep was set at $\delta t = 1$ fs.

All RENS simulations were performed using the OpenMM simulation package [42] in conjunction with the python wrapper PyOpenMM². The integration of the switching dynamics was implemented using the CustomIntegrator class of OpenMM, with an Langevin integration scheme proposed by Vanden-Eijnden and Ciccotti [123] developed for systems with holonomic constraints and accurate to δt^2 .

As we described above, our calculations below involve the solvation of MEA in a *frozen configuration* x . The solvation free energy, then, is calculated with respect to this constraint, and $\Delta F_{\text{solv}} = \Delta F_{\text{solv}}(x)$ (see Eq. 4.3 above). Because our RENS setup is

²<http://simtk.org/home/pyopenmm>

tailored for solvation calculations and not efficient phase space sampling, the sampling of the unconstrained solute would be challenging in the current implementation. While we present work in this direction in Sec. 4.7 through analysis of the metastable states of MEA, we leave other solvation calculations of the unconstrained molecule for future studies that will potentially combine cavity growth with temperature replica exchange.

4.6 Solvation free energy calculation

The solvation free energy of the solute is calculated from the thermodynamic cycle formed by our sequence of replicas (see Fig. 4.5). In analogy to QCT calculation (Eq. 4.5), ΔF_{solv} can be split up into three contributions,

$$\Delta F_{\text{solv}} = \Delta F_{\text{cav}}^0 - \Delta F_{\text{cav}} + \Delta F_{\text{int}} \quad (4.12)$$

where ΔF_{cav}^0 is the free energy change associated with growing out the cavity in absence of the solute, and ΔF_{cav} is the free energy change associated with growing out the cavity in presence of the solute. Finally, ΔF_{int} is the free energy cost associated with turning on the solute-solvent interactions. We note that while ΔF_{cav} and ΔF_{int} depend on the details of the solute, ΔF_{cav}^0 depends solely on the solvent state, and need only be calculated once when one is interested in many solutes. In the case that our cavity potential is modeled by a hard sphere, Eq. 4.12 is analogous to the QCT decomposition, Eq. 4.5, above.

In Fig. 4.6 we plot the free energy change associated with cavity growth as a function of cavity size R , both in the pure solvent ΔF_{cav}^0 (red) and in presence of the solute ΔF_{cav} (green). These quantities are given as a sum over the free energy differences $\Delta F_{i \rightarrow i+1}$ between neighboring replicas i and $i+1$, which represent the cost associated with growing the cavity from R_i to R_{i+1} . In absence of the solute, cavity growth proceeds from replicas

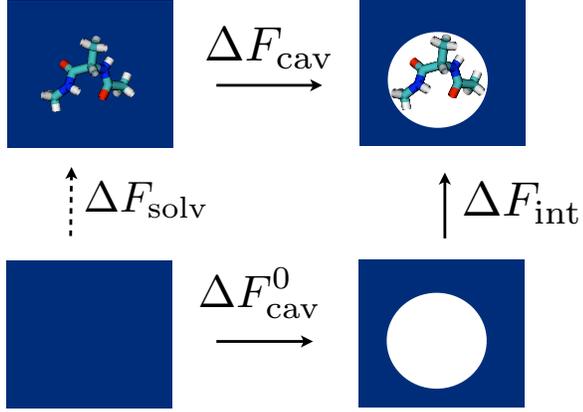


Figure 4.5: Calculation of ΔF_{solv} from a thermodynamic cycle with cavity growth.

1 to $M/2 - 1$, and for a given R

$$\Delta F_{\text{cav}}^0(R) = \sum_{i=1}^n \Delta F_{i \rightarrow i+1} \quad (4.13)$$

with $R_1 = 0$ and $R_n = R$. Similarly, cavity growth *in presence of the solute* is performed from replica M to replicas $M/2$, and

$$\Delta F_{\text{cav}}(R) = \sum_{i=m}^{M-1} \Delta F_{i \leftarrow i+1} \quad (4.14)$$

where $R_M = 0$ $R_m = R$ (see Fig. 4.1). The $\Delta F_{i \rightarrow i+1}$ values between neighboring replicas were each estimated with Bennett's acceptance ratio [13], generalized to nonequilibrium processes [28]. Finally, ΔF_{int} was calculated from the the solute insertion step performed between replicas $M/2 - 1$ and $M/2$ at $R = R^*$ (rightmost pair of replicas in Fig 4.1), calculated to be $\Delta F_{\text{int}} = -21.9k_B T$.

We can now obtain the solvation free energy ΔF_{solv} from Eq. 4.12 by using our estimate for ΔF_{int} , as well as ΔF_{cav}^0 and ΔF_{cav} evaluated at $R = R^*$. However, because the cavity potential is still non-zero for $R = 0$ (see Eq. 4.10), there is an additional

contribution to ΔF_{solv} given by the free energy cost associated with turning off V_{cav} at $R = 0$ in presence of the solute, relative to that in absence of the solute. Because V_{cav} is relatively weak and slowly decaying for $R = 0$, this contribution, calculated to be $-0.8k_B T$, is small relative to the above calculation. Using Eq. 4.12, along with this small correction, our final estimate is $\Delta F_{\text{solv}} = -12.8k_B T$.

We note that our estimate for ΔF_{solv} is relatively small compared to the cost associated with growing out the cavity, potentially contributing a large uncertainty in the estimation of ΔF_{solv} . This issue remains to be investigated more fully.

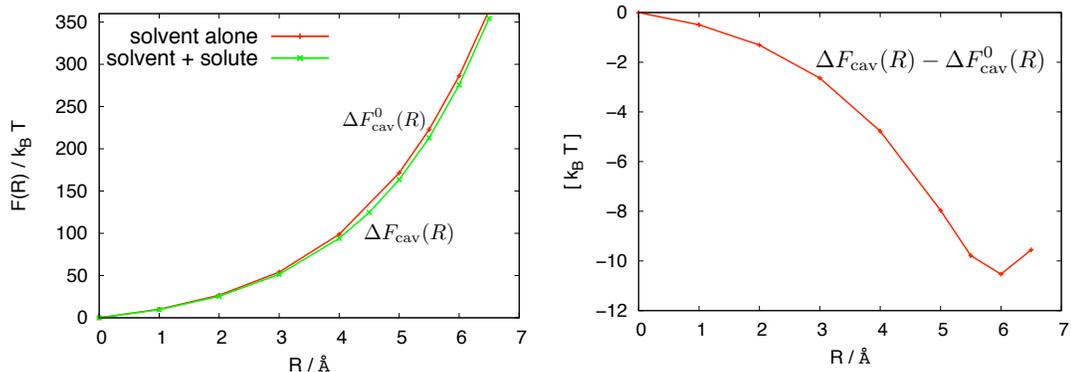


Figure 4.6: *Left:* Free energy cost of cavity growth as a function of cavity size R for replicas in absence of solute-solvent interactions (red) and in presence of solute-solvent interactions (green). *Right:* Free energy of cavity growth in presence of solute, relative to that in absence of solute

4.7 Metastable states

In working towards a complete understanding of the solvation of MEA, we analyze the metastable states arising from our molecule. The conformationally-stable states of MEA can be described by the torsion angle ψ defined by the backbone atom sequence N-C-C-O (see Fig. 4.7). Because the stable states of MEA can be described by a one-

dimensional known variable, we need not use temperature RENS to probe the equilibrium distribution of MEA. Again we performed Hamiltonian RENS on the fully interacting solvent-solute system (in absence of the cavity potential), where the potential of replica i was augmented by a harmonic bias on ψ about center ψ_i

$$U^i(x) = U(x) + \frac{1}{2}k_\psi [\psi(x) - \psi_i]^2. \quad (4.15)$$

We performed a RENS simulation with $M = 20$ replicas, with ψ_i values equally spaced from -162° to $+162^\circ$ and a spring constant $k_\psi = 1.2 \times 10^{-2}k_B T/\text{degree}^2$. The switching time of $\tau = 10$ fs was chosen between all replica pairs, with a total simulation run time of 1.25 ns which provided sufficient sampling within each replica.

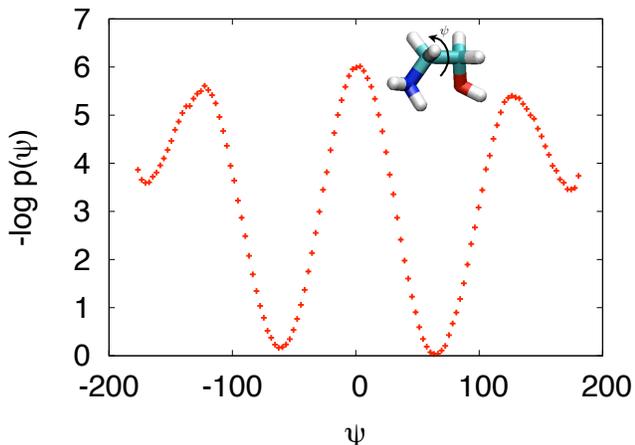


Figure 4.7: Potential of mean force of MEA along dihedral angle ψ .

The potential of mean force $PMF(\psi) \equiv -k_B T \ln p(\psi)$ of the unbiased system was recovered by using the WHAM method [38, 66] on the histograms of ψ for each of the 20 replicas. The potential of mean force in Fig. 4.7 reveals the *gauche* state at $\psi = \pm 60^\circ$ to be the most stable. This is in accordance with previous findings [29, 72], and presumably

influenced by the presence of strong intramolecular hydrogen bonding between ends of the molecule, as previous studies have found [19, 23, 24, 62]. MEA also exhibits a slightly less populated *trans* state at $\psi = 180^\circ$. The symmetry of $PMF(\psi)$ about $\psi = 0$ is expected from the underlying symmetries of the force field.

4.8 Summary

In this chapter we developed a RENS implementation for the calculation of solvation free energies small solutes. While traditional solvation free energy estimates involving insertion of the solute into the solvent are typically plagued by slow convergence due to poor overlap, with our method insertion was assisted by the formation of a solvent cavity to hinder solvent-solute overlaps. We tested our method on a frozen configuration of the solute monoethanolamine in explicit TIP3P water. In working towards a complete picture of MEA solvation we also characterized the metastable states arising through a series of umbrella sampling windows used in conjunction with RENS, finding stable *gauche* and *trans* configurations.

In the future we plan to calculate the free energy of solvation of the unfrozen MEA molecule by combination of the above RENS solvation method with efficient sampling strategies such as temperature replica exchange.

Dynamics of complex systems: finding the mechanism of ionic dissociation in water

5.1 Introduction

¹In the chapters up until now we have focused on efficient methods for sampling equilibrium distributions of complex systems. While novel sampling methods such as RENS can be very helpful for understanding thermodynamic properties such as the free energy or the classification of metastable states, a complex system is not completely understood without an adequate description of the dynamics that govern transitions between these states. In this chapter we focus on the dynamics of rare events which bring a system from one state to another.

Characterizing the dynamics of these rare events is vital to the understanding of large-scale structural changes that occur in many complex systems in nature. Rare events

¹This chapter is based upon work performed with Christoph Dellago at the University of Vienna.

typically involve many cooperative parts that act together in a complicated way along a transition from one long-lived state to another, and the determination of the collective variables responsible for the reaction dynamics can provide much insight into the physical mechanism underlying a transition. In the best circumstance a *reaction coordinate* is found, a function of these collective variables that alone is sufficient to track the progress of a reaction. Unfortunately finding an adequate description of the reaction in terms of a reaction coordinate or even only identifying the collective variables is a challenging task: not only is sampling computationally demanding due to the rare nature of the transition, but the reaction also proceeds through a high-dimensional phase space and so it is often difficult to discern which variables promote the transition. Despite many novel techniques for reaction analysis (see e.g. [74,93]), finding a good reaction coordinate remains a challenge for many processes occurring in complex systems.

In this chapter we investigate the kinetic pathways leading to ionic dissociation, in particular the dissociation of Na^+Cl^- in water. Microscopically this system contains metastable associated and dissociated states, separated by a free energy barrier preventing frequent transitions. Along a reaction in which the ion pair transitions between associated and dissociated states, a number of system rearrangements must take place which crucially involve the surrounding solvent molecules. The first simulations of this system were performed by McCammon et al. [12,14] and Rey et al. [100] who used umbrella sampling and constrained solute simulations, respectively, to investigate solvent structure and thermodynamic properties as the interionic distance $r_{\text{ion}} = |\mathbf{r}_{\text{Na}^+} - \mathbf{r}_{\text{Cl}^-}|$ is varied. More recent work by Geissler et al. [43] employed transition path sampling to study the reaction, showing under careful statistical analysis that r_{ion} alone is a poor reaction coordinate in describing dissociation – and that the surrounding solvent must be taken into account in

a good reaction coordinate. Despite this work and others, a complete description of the solvent motion leading to dissociation is missing. While the ultimate goal is to find a reaction coordinate for the event, even a complete set of solvent variables that are responsible for dissociation is still unknown, and hence further investigation is needed.

In this study we shed some more light on water’s unique role by investigating the thermodynamic and dynamical properties of the dissociation reaction. Our main results are organized as follows. After describing our model in Sec. 5.2, we present in Sec. 5.3 a thermodynamic description of the reaction in terms of competing thermodynamic driving forces, showing that dissociation is an energetically favorable but entropically unfavorable process. We argue that the decrease in solvent entropy upon dissociation is due to an increasingly large number of highly-coordinated solvent molecules in the first hydration shell as the ions move apart. In Sec. 5.4 we investigate the relative importance of various system variables in promoting dissociation. Following Ref. [43], we employ statistical analysis of dissociation (committor) probabilities on data from various constrained ensembles: For data with constrained r_{ion} we verify in Sec 5.4.1 that r_{ion} is indeed important in the reaction, but does not capture the entire mechanism, in confirmation with Ref. [43]. In Sec. 5.4.2 various solvent degrees of freedom are constrained to pinpoint the range over which the solvent influences the dissociation event. In Secs. 5.5 and 5.6 we investigate various dynamical aspects of dissociation, highlighting in Sec. 5.5 timescales under which solvent rearrangements occur which drive dissociation, and in Sec. 5.6 the importance of inertial effects near the transition state. Finally, our results are summarized and discussed in Sec. 5.7.

5.2 Model

The system we studied consists of one Na^+ ion and one Cl^- ion immersed in a bath of $N_w = 216$ water molecules. The ion pair and ion-water interactions were modelled using the OPLS force field [60], which includes short-ranged Lennard-Jones and long-ranged Coulomb terms. More specifically, the ion-ion interaction is given by

$$V(r_{\text{ion}}) = 4\epsilon \left[\left(\frac{\sigma_{\text{ion}}}{r_{\text{ion}}} \right)^{12} - \left(\frac{\sigma_{\text{ion}}}{r_{\text{ion}}} \right)^6 \right] - \frac{e^2}{4\pi\epsilon_0 r_{\text{ion}}}, \quad (5.1)$$

where e is the elementary charge, ϵ_0 is the permittivity of free space, and Lennard-Jones parameters for the ion pair are $\sigma_{\text{ion}} = 3.8355 \text{ \AA}$ and $\epsilon = 0.0756034201 \text{ kJ/mol}$. The water molecules interact via the rigid TIP4P model [61]. Calculation of the long-ranged electrostatic forces from periodic boundary conditions were handled with particle mesh Ewald summation. The simulations were performed at a constant temperature $T = 300 \text{ K}$ and constant volume $V = (18.64 \text{ \AA})^3$, which was chosen from an equilibrated constant-pressure simulation at ambient conditions. To sample the NVT ensemble the system evolved under Langevin dynamics with a friction coefficient corresponding to a timescale of 0.1 ps . Our simulations were performed using Gromacs [11], with a timestep $\delta t = 2 \text{ fs}$ and integration performed via a stochastic leap-frog algorithm [45].

To begin our study we generated 10 trajectories sampling the canonical ensemble, totaling 80 ns , that involved 227 transitions between associated and dissociated states. Initial conditions for each of the 10 trajectories were taken from a previous simulation run using same Langevin dynamics. Each of the 10 points were separated by 400 ps from their neighbors, such that they can be considered statistically independent. We calculate the free energy along the interionic distance r_{ion} as

$$F(r_{\text{ion}}) = -k_{\text{B}}T \ln p(r_{\text{ion}}), \quad (5.2)$$

by histogramming r_{ion} from the concatenated trajectories. Here k_{B} is the Boltzmann constant. From Fig. 5.1 we see that $F(r_{\text{ion}})$ contains a metastable associated state with a corresponding free energy minimum at $r_{\text{ion}} = 2.7 \text{ \AA}$. This minimum is separated from the solvent-separated state, centered around 5 \AA , by a barrier of $5 k_{\text{B}}T$. For future reference we identify the associated state as all configurations for which $r_{\text{ion}} < 3.2 \text{ \AA}$, the dissociated state as $r_{\text{ion}} > 4.4 \text{ \AA}$, and the transition region to be $3.2 \text{ \AA} < r_{\text{ion}} < 4.4 \text{ \AA}$.

5.3 Thermodynamics of ionic dissociation

To gain an understanding of ionic dissociation, we first investigate the thermodynamics of the process along the order parameter r_{ion} . In the NVT ensemble, the Helmholtz free energy F contains energetic and entropic contributions, which we calculate as a function of the ion pair separation (see Fig. 5.1). The energy profile is computed from a number of simulations, in each of which r_{ion} is constrained to a value between 2.47 \AA and 7.51 \AA . For each simulation the potential energy E was averaged over a 100 ns-long trajectory. Plotted in green is $U(r_{\text{ion}}) = \langle E \rangle_{r_{\text{ion}}} - \langle E \rangle_{\infty}$ the average energy, after subtracting the asymptotic value. The entropy S , plotted in blue, is identified from

$$F(r_{\text{ion}}) = U(r_{\text{ion}}) - TS(r_{\text{ion}}). \quad (5.3)$$

Note that the errors on U and S are due to the large energy fluctuations of the many solvent-solvent interactions in the bulk. We see, in Fig. 5.1, that the associated state is stabilized energetically, with a $3 k_{\text{B}}T$ barrier to overcome before energetically-favorable dissociation occurs. The entropy S leads to an attractive contribution to the free energy opposing dissociation in the range $r_{\text{ion}} < 4.0 \text{ \AA}$, a behavior familiar from entropy-driven hydrophobic association [112]. Thus, the energy and entropy shown in Fig. 5.1 show markedly different behavior than in the implicit solvent case where the solvent is mod-

eled simply by a dielectric constant. In this case, the energy has only one minimum at the associated state, and the driving force to dissociation is entirely entropic, due to an available configuration space that grows as r_{ion}^2 . This confirms that the solvent plays a non-trivial role in the dissociation process. We note that this thermodynamic picture is contrary to the behavior of a model protein-ligand complexes in water as found in recent simulations studies by McCammon [9,107]. For oppositely charged protein and ligands, the dissociation is an enthalpically-unfavorable and entropically favorable process.

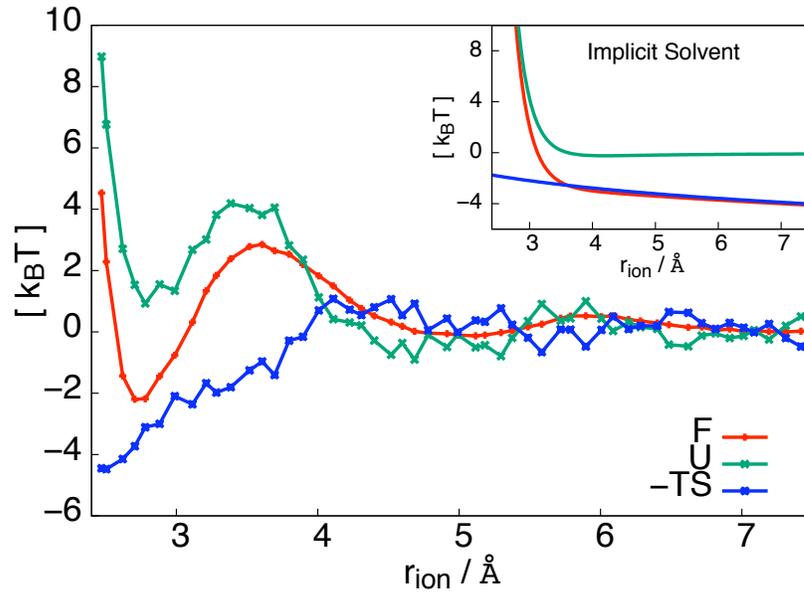


Figure 5.1: Thermodynamics of ionic dissociation. The free energy (red) as a function of r_{ion} displays a stable associated state at $r_{\text{ion}} = 2.7 \text{ \AA}$, separated from the dissociated state by a free energy barrier of $5 k_B T$. Also plotted are the average energy (green) and negative entropy (blue) as a function of r_{ion} . Dissociation is driven energetically and opposed by entropy. The inset shows the free energy, the energy and the entropic contribution for an implicit solvent model, in which the electrostatic interaction between the two ions is reduced by a factor $\epsilon = 80$.

To further investigate the influence of the solvent on the system entropy, we plot

in Fig. 5.2 the average numbers n_{Na} and n_{Cl} of water molecules within the first solvation shell of Na^+ and Cl^- , respectively. [The solvation shell radii for each ion correspond to the respective minimum in the ion-oxygen radial distribution function, 3.34 Å for Na^+ and 3.74 Å for Cl^- (data not shown).] During the dissociation process, the solvation numbers of the Na^+ and Cl^- ions increase by about 2 and 1, respectively. In Fig. 5.2 we also plot in blue the average number of water molecules simultaneously in the solvation shells of both ions. While for the associated state there is one shared molecule, the number of such molecules starts to increase at about $r_{\text{ion}} = 3.3$ Å and reaches 2 at $r_{\text{ion}} = 4$ Å, where the solvation number the ions saturate. The number of shared water molecules then falls to 1 as the solvent-separated state is reached around $r_{\text{ion}} = 5$ Å and finally to 0 around $r_{\text{ion}} = 6$ Å. Since the number of shared water molecules is constant for $r_{\text{ion}} \lesssim 3.3$ Å, the increase in the average solvation numbers n_{Na} and n_{Cl} is due to additional water molecules entering the respective solvation shells from the bulk. In the range 3.3 Å $\lesssim r_{\text{ion}} \lesssim 4.0$ Å, however, the total number of water molecules in the combined solvations shells of Na^+ and Cl^- grows only slowly, while the solvation numbers in the individual shells increase by a total of about 1 water molecule due to a solvent reorganization that creates an additional shared water molecule. As the interionic distance grows further, the solvation numbers of the individual ions stays roughly constant while the number of shared water molecules decreases to 0 for sufficiently separated ions. During this final stage of the dissociation, 2 water molecules enter the solvation shell of the ions from the bulk to compensate for the loss of shared water molecules. During the entire dissociation process, the total number of water molecules in the combined solvations shells of Na^+ and Cl^- increases by about 4 in the average.

As shown in Fig. 5.2, the solvation numbers n_{Na} and n_{Cl} show the same roughly

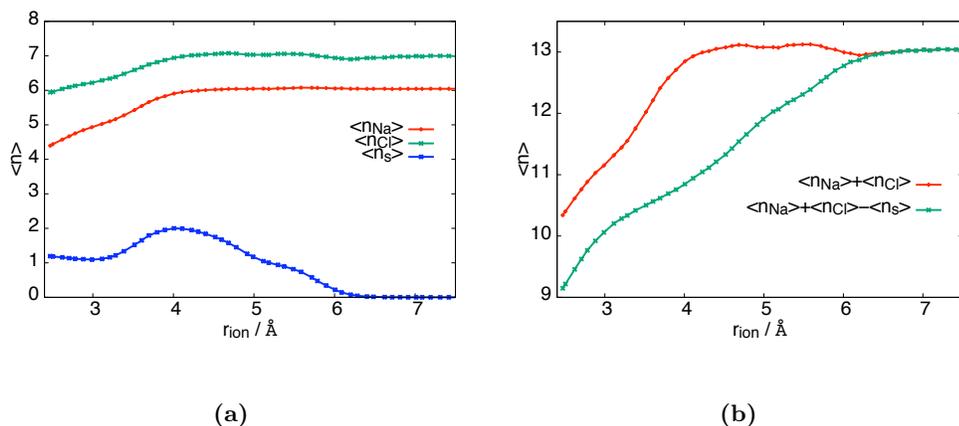


Figure 5.2: *Top:* Average number $\langle n_{\text{Na}} \rangle$ and $\langle n_{\text{Cl}} \rangle$ of water molecules in the first solvation shell of Na⁺ and Cl⁻. Shown in blue is the number $\langle n_s \rangle$ of waters common to the solvation shells of both ions. *Bottom:* Sum $\langle n \rangle = \langle n_{\text{Na}} \rangle + \langle n_{\text{Cl}} \rangle$ of the number of water molecules in the first solvation shells of Na⁺ and Cl⁻ and total number of water molecules $\langle m \rangle = \langle n_{\text{Na}} \rangle + \langle n_{\text{Cl}} \rangle - \langle n_s \rangle$ in the combined first solvation shell.

linear increase with interionic distance as the entropy (see Fig. 5.1), suggesting that the entropy change during dissociation is due to the reduced freedom of motion of water molecules tightly bound to the ions. Indeed, water molecules in the first solvation shell of ions have been observed in simulations to be orientationally highly restricted [35], reducing the configurational space available to the molecules compared to the bulk. The orientational restraints acting on first solvation shell molecules are particularly pronounced for water molecules shared by both ions. As the ions dissociate, the number of such “low entropy” solvent molecules increases, leading to a net entropy decrease in the system. For interionic distances $r_{\text{ion}} \approx 4$ Å and larger, the entropy is approximately constant even though the number of total water molecules in the combined first solvation shells of the two ions continues to increase (see green line in the bottom panel of Fig. 5.2). In this

regime, the sum n_{Na} and n_{Cl} , which double counts shared water molecules and equals the number close contacts of water molecules with one of the ions, remains constant. This indicates that the entropy is related to the number of such close contacts rather than the total number of solvating water molecules. This conclusion is confirmed by the linear behavior of the entropy in the range $3.3 \text{ \AA} \lesssim r_{\text{ion}} \lesssim 4.0 \text{ \AA}$, where the total number of solvating water molecules grows only slowly, but the number of close contacts increases due to the increase of shared water molecules. For $r_{\text{ion}} \lesssim 3.3 \text{ \AA}$, on the other hand, the number of close contacts increases due to water molecules entering the first solvation shells of the ions from the bulk. Thus, as shown in Fig. 5.3, overall there is a roughly linear relationship between the number of close contacts and the entropy, where each close contact contributes an entropy decrease of $\Delta s \approx 1.9k_{\text{B}}$.

The above analysis, however, provides only a partial picture of ionic dissociation: As shown in Ref. [43], the interionic distance r_{ion} can be used as an *order parameter* to distinguish between the associated and dissociated states, but fails in describing the progress of the dissociation. In other words, r_{ion} is a poor *reaction coordinate* implying that solvent degrees of freedom must be explicitly taken into account. In the next section we corroborate this finding and carry out a new type of statistical analysis to identify the range within which solvent degrees of freedom affect the dissociation process.

5.4 Transition path analysis

One of the major goals of characterizing a reaction pathway is the determination of the system variables that are important for the reaction to proceed. This set of variables, if known, provide a basis for understanding the physical mechanism underlying a complex transition. A good reaction coordinate r will in general be a function of a number of

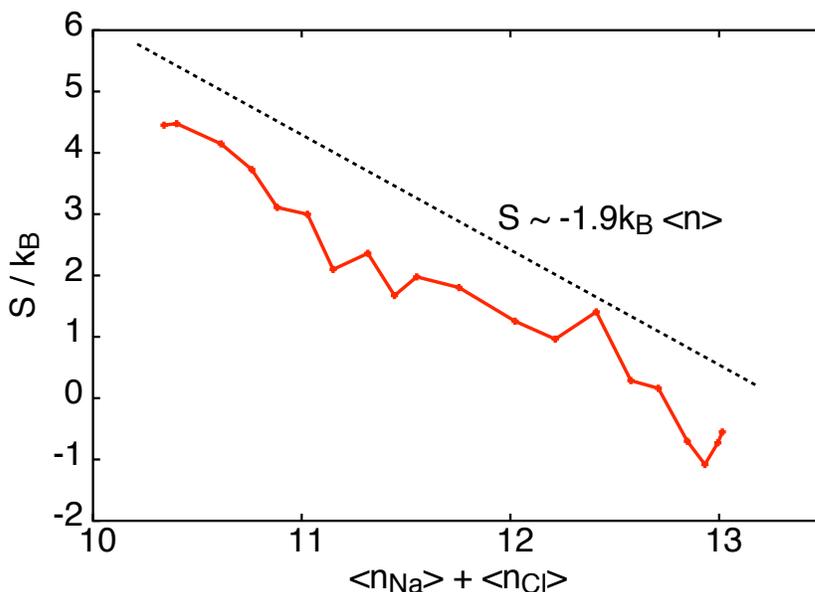


Figure 5.3: Entropy S as a function the number $\langle n \rangle = \langle n_{\text{Na}} \rangle + \langle n_{\text{Cl}} \rangle$ of close contacts between the ions and water molecules in the first solvation shell. The dotted line denotes a linear fit to the data, displaced vertically for better visibility.

such collective variables, which together completely specify the progress of a reaction. In this section we employ committor analysis first to test the quality of r_{ion} as a reaction coordinate, confirming the results of Ref. [43], and then to examine the influence of various solvent degrees of freedom on the dissociation process.

In searching for the important collective variables, or optimally a reaction coordinate, we ultimately seek a projection of phase space that preserves the dynamical information pertaining to the reaction. This dynamical information is captured by the committor probability, $p_B(x)$, a key tool in determining these collective variables. For a system containing two long-lived stable states, labeled A and B , $p_B(x)$ is defined as the probability that a trajectory initiated from configuration x will relax to state B before reaching state

A. As such, p_B is a statistical measure of the progress of a reaction. In particular, configurations with $p_B = 1/2$ can be considered to be transition states as they have equal probability to relax in to A or B .

While the perfect reaction coordinate is the committor itself [32, 52], $r^\dagger(x) = p_B(x)$, a *good* reaction coordinate $r(x)$ will to a good accuracy specify $p_B(x)$, in the sense that the committor can be written as $p_B(x) \approx p_B[r(x)]$ and the reaction information is contained in the variables in $r(x)$. Thus, for a good reaction coordinate, the distribution of p_B values for configurations restricted to a particular value of r should be sharply peaked around some characteristic value.

The quality of a trial reaction coordinate can then be investigated by probing this distribution of committor values. For a trial reaction coordinate \hat{r} , one calculates p_B values for configurations in the constrained equilibrium ensemble with $\hat{r}(x) = \text{const}$. A distribution $P_{\hat{r}}(p_B)$ is estimated by histogramming the p_B values of the constrained ensemble, and the quality of \hat{r} is assessed from the shape of $P_{\hat{r}}(p_B)$: if $P_{\hat{r}}(p_B)$ is a sharply peaked function of p_B , then the degrees of freedom specified by $\hat{r}(x)$ determine to a good approximation the fate of the reaction. If however $P_{\hat{r}}(p_B)$ is not sharply peaked, then other degrees of freedom not included in $\hat{r}(x)$ play a role in specifying how the reaction will proceed, and thus $\hat{r}(x)$ is an insufficient reaction coordinate.

In the following sections we assess the relative importance of various system variables in Na^+Cl^- dissociation by applying committor analysis to various constrained ensembles. In our solvated Na^+Cl^- system, we define B as the dissociated state, for which $r_{\text{ion}} > 4.4\text{\AA}$, and A as the associated state, $r_{\text{ion}} < 3.2\text{\AA}$. For our committor calculations, $p_B(x)$ is estimated by shooting off N_s independent trajectories starting from x , with initial velocities drawn from the corresponding Maxwell-Boltzmann distribution and evolving

under Hamilton’s equations of motion. The estimate for $p_B(x)$ is given by the fraction of these trajectories that reach the dissociated state before associating, with an error

$$\sigma = \sqrt{\frac{p_B(1 - p_B)}{N_s}}. \quad (5.4)$$

For our calculations $N_s = 100$ shots were performed for each p_B estimate such that $\sigma \leq 0.05$. Note that this statistical error of the estimated committor leads to a broadening of the committor distribution that can be statistically quantified and needs to be taken into account in the interpretation of committor distributions [92].

5.4.1 Constrained interionic distance r_{ion}

In characterizing the kinetic pathways leading to ionic dissociation, we first test the performance of the interionic distance r_{ion} as a reaction coordinate. Such a calculation for r_{ion} has been done previously [43], but we repeat it here, because we use a slightly different force field for the ion-ion interaction and the ion-water interaction. To test whether r_{ion} alone is a good reaction coordinate, we apply committor analysis on configurations with constrained r_{ion} . Committor values were estimated for 665 configurations having $3.45 \text{ \AA} \leq r_{\text{ion}} \leq 3.75 \text{ \AA}$ taken from the equilibrium run used to generate Fig. 5.1, which involved many transitions between associated and dissociated states. The constraint range of r_{ion} was chosen around the position of the top of the free energy barrier (see Fig. 5.1). We plot in Fig. 5.4 the distribution of p_B values on this constrained surface. Because this range of r_{ion} corresponds to the top of the free energy barrier, one would expect for a good reaction coordinate a sharp unimodal distribution centered at $p_B = 0.5$. What one sees, however, is a bimodal distribution peaked at $p_B \approx 0$ and 1 and relatively low population at 0.5. Hence, there are structures with the same interionic distance r_{ion} but very different

²Narrowing the width of r_{ion}^* did not qualitatively change the behavior of our results.

relaxation behavior, indicating that the solvent degrees of freedom are important in the system committing to associate or dissociate. As this behavior was observed previously by Geissler et al. [43] for a different force field, these findings highlight that the solvent’s role in dissociation is robust and of general importance in describing the reaction.

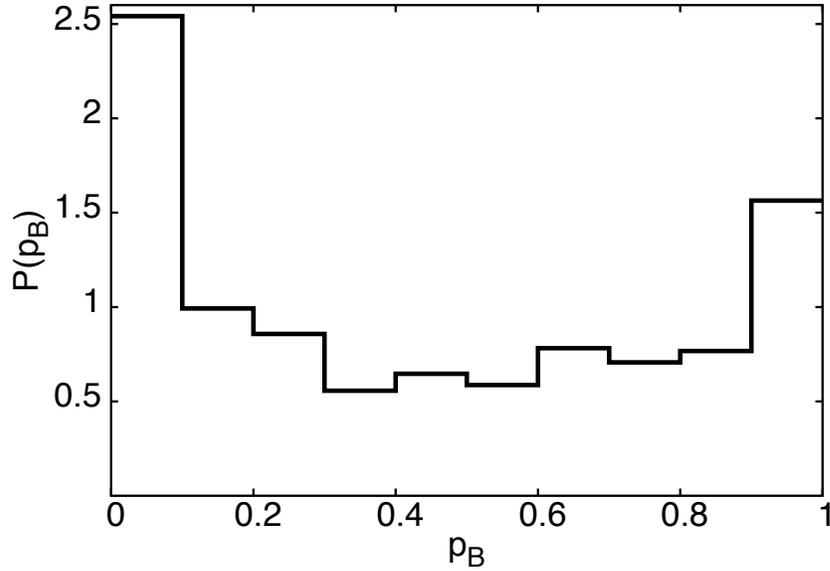


Figure 5.4: Distribution of p_B values for equilibrium configurations x restricted to $r_{\text{ion}}(x) = r_{\text{ion}}^*$, corresponding to the top of the free energy barrier shown in Fig. 5.1. The bimodal behavior indicates that r_{ion} alone is not a sufficient reaction coordinate.

5.4.2 Constrained solvent

Since the interionic distance r_{ion} alone is an insufficient reaction coordinate, the surrounding solvent must play a crucial role in the system committing to associate or dissociate. To investigate the role of the solvent more closely we seek to identify which water molecules are important in the reaction, with the specific goal of finding a length scale

over which the water molecules influence the reaction. Specifically, we perform committor analysis, on a constrained system as above, where in addition to a fixed interionic distance r_{ion} , we also constrain or “freeze” water molecules within a particular probe range of the ion pair. Committor analysis applied to this system in which a part of the water molecules is held at fixed positions will guide us in finding the length scale that determines the range of solvent influence on dissociation: If the distribution of p_{B} values on these constrained configurations is sharply-peaked, then the dissociation event is only sensitive to the frozen molecules within the given probe range; If, however, the remaining unfrozen molecules in the periphery of the simulation box strongly influence the reaction, then the p_{B} distribution will not show a single pronounced peak.

We wish to find the probe range over which the peripheral variable solvent molecules cease to influence p_{B} . In the limit of a very small probe range, only the ions are restrained, and we expect to see a distribution of p_{B} values like Fig. 5.4, where the other unfrozen molecules are clearly influencing the fate of the ion pair. In the opposite limit of a very large probe range, the entire simulation box is frozen and we expect a p_{B} distribution that is very sharply peaked about some characteristic value. We seek the smallest probe range over which the variability of p_{B} becomes small enough that we are confident the molecules within the probe range specify the fate of the reaction. To this end, three separate probe ranges were considered, set by the hydration structure of the ion-pair: either all molecules up through the first, second or third hydration shells were frozen (see Fig. 5.5). [A water molecule was said to be within the n -th solvation shell of the ion pair if the oxygen resides within the n -th hydration shell of either Na or Cl. The solvation shell distances for each ion were set by the respective minimum in the ion-O radial distribution function (data not shown).]

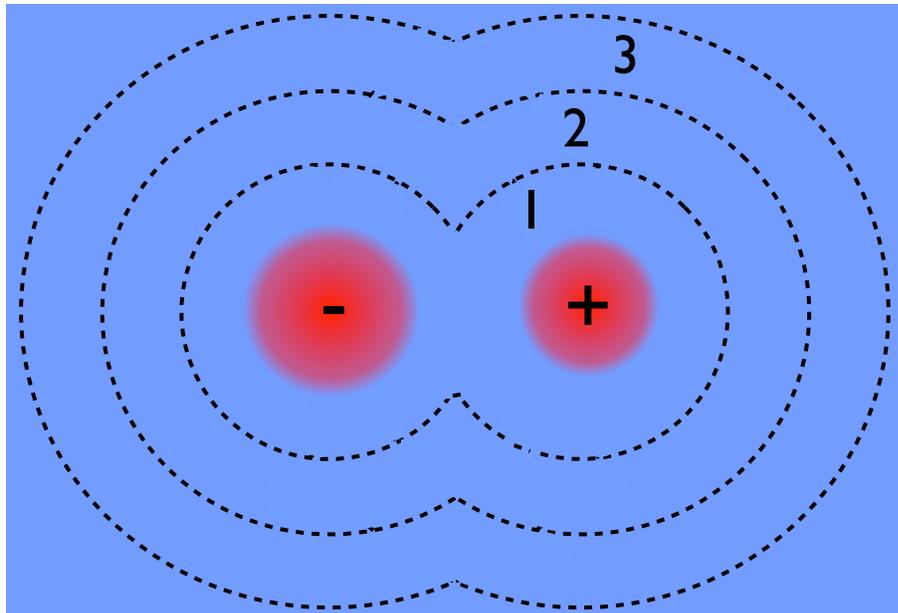


Figure 5.5: Depiction of the first three solvation shells of the ion pair, which were selectively constrained to investigate solvent influence on ionic dissociation.

To begin the analysis, one of the three probe ranges is chosen, centered around the ion pair, within which all waters are constrained. A representative initial configuration is selected, from which 125 new configurations are generated by evolving the system dynamically, each with identical solvent positions within the probe range but variable positions outside. We enforced these constraints on the dynamics by simply preventing the positions of the relevant waters to be updated during the integration of the Langevin equation of motion. This dynamical scheme samples a constrained equilibrium state that is equivalent to a reduced system (the solvent outside the probe region), in the presence of a static external field (imposed by the frozen solvent molecules and ion pair within the probe region). Committor analysis is then performed on this constrained state by calculating p_B values for each of the 125 configurations and histogramming the obtained

committor values. Note that in the trajectories generated for the committor calculation all constraints used to prepare the initial conditions were released.

The results of the committor analysis are shown in Fig. 5.6, where each subfigure corresponds to a given probe range. Plotted within a given probe range are three distributions, colored red, green, and blue, which correspond to three distinct sets of 125 configurations with frozen solvent having p_B values near 0, 0.5 and 1 respectively. For the smallest probe range, where the solvent is constrained only in the first hydration shell (subfigure (a)), there is a very wide distribution of p_B values for each of the three configuration sets. In subfigure (b), where the first two solvation shells are constrained, the distributions are not as broad as (a), but still show rather large p_B -variability, implying that molecules farther out are of importance. Finally, when all three solvation shells are constrained, in subfigure (c), we see a tight distribution of p_B values, which suggests that the commitment to associate or dissociate is more or less determined by the solvent molecules within the first three solvation shells. These results are consistent with Ref. [43], who found that the dissociation couples to solvent motion between the second and third solvation shells. Interestingly, the $p_B = 0.5$ configuration set (green) shows the broadest distribution for all three probe ranges, implying that p_B is particularly sensitive to long-ranged solvent motion near the transition state. This is also consistent with Ref. [43], who came to the same conclusion from studies of the mean solvent force on the ion pair at the transition state.

The waters within the first three hydration shells seems to be sufficient to specify the outcome of the reaction. Can this picture be refined further? Specifically, what is the relative importance of steric forces to electrostatics? Because of the long-ranged influence of solvent on p_B , we would expect electrostatics to play an important role. To test this

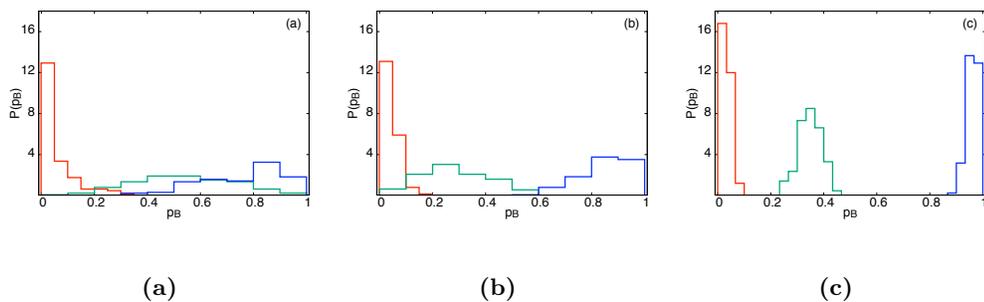


Figure 5.6: Committor analysis applied to configurations containing “frozen” (identical) solvent coordinates in (a) the first, (b) first two, and (c) first three solvation shells. In each figure, the colors distinguish between different sets of frozen solvent, chosen near the associated state (red), transition region (green), and dissociated state (blue). Note that the variance of the committor distributions shown in panel (c) is consistent with sharp p_B values in the respective sets of configurations and should just be due to the statistical uncertainty in the p_B calculations [92].

conjecture, we generated a set of 125 configurations having identical oxygen positions within the first three solvation shells, but variable hydrogen and dummy atom positions for our 4-point TIP4P model. Because the Lennard-Jones forces are specified by the oxygen position alone, the only variable forces within the three solvation shells are due to electrostatic interactions. The resulting committor distributions, shown in Fig. 5.7, are still somewhat broad, indicating that the charge distribution of the waters are of general importance, and that the p_B is determined by a combination of steric and electrostatic effects.

5.5 Timescales of p_B fluctuations

In this section we investigate the time-fluctuations of p_B , with the goal of finding the relevant timescales under which the solvent rearranges itself to promote dissociation.

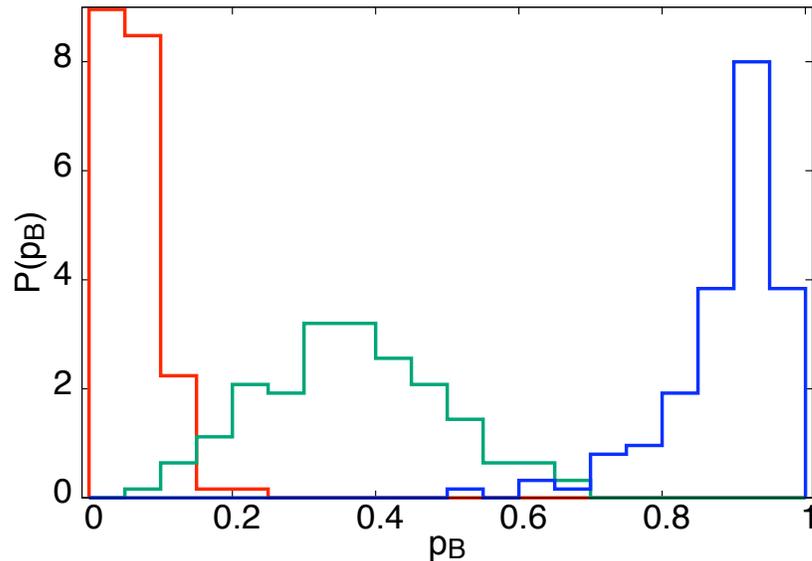


Figure 5.7: The effect of water orientations on p_B . All configurations within each color contain identical oxygen coordinates but have varying orientations of the water molecules in the first three solvation shells. We analyzed three sets of oxygen positions, chosen near the associated (red), transition (green) and dissociated (blue) states.

We capture the dynamics of the entire solvent by calculating p_B along a trajectory with constraint $r_{\text{ion}} = 3.73 \text{ \AA}$, near the peak of the free energy barrier, which we plot in Fig. 5.8a. This is contrasted with Fig. 5.8b, where our trajectory contains a constrained first solvation shell as well as constrained $r_{\text{ion}} = 3.73 \text{ \AA}$. We see qualitatively that the p_B fluctuations are somewhat suppressed when the first solvation shell is fixed in addition to r_{ion} , consistent with committer analysis of previous sections (compare Figs. 5.4 and 5.6a). In Fig. 5.8a where all solvent is free, we observe two timescales: on a large timescale of roughly 5ps we observe large p_B fluctuations between 0 and 1, and on a shorter timescale of roughly 0.5ps we see oscillatory-like fluctuations of a much smaller magnitude. Because this smaller

timescale persists in Fig. 5.8b, when we freeze the first solvation shell, it is tempting to conclude that the smaller fluctuations are due to solvent rearrangements outside the first solvation shell, and the larger p_B fluctuations are due to the water rearrangements in the first solvation shell, which occur under timescales ten times as large.

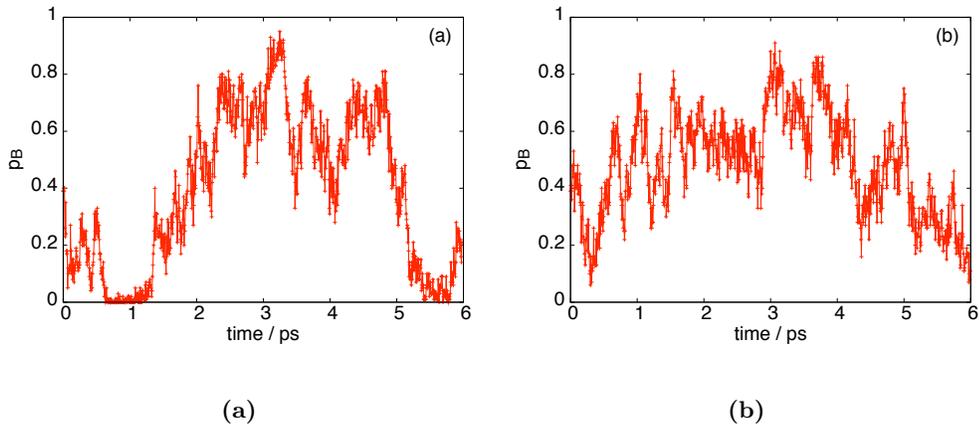


Figure 5.8: Time-dependence of $p_B[x(t)]$, for trajectories generated from dynamics with (a) constrained r_{ion} and (b) constrained r_{ion} and first solvation shell.

5.6 Inertial effects

The reaction pathways characterizing rare events will ultimately depend upon the system dynamics. Observables such as kinetic rate constants, committor values and transition probabilities are generally properties of the underlying dynamics governing the time evolution of the system. By analyzing transition pathways for various types of dynamics, one can then learn something about the relative importance of certain dynamical features in facilitating a rare transition. In this section, we compare our calculations for Hamiltonian dynamics to analytic results for diffusive dynamics, highlighting the importance of

inertial effects in enhancing reaction probability.

To track the differences that arise in these two dynamical regimes, we compare the committor probability p_B to the transition path probability p_{TP} . While $p_B(x)$ is the probability that a trajectory passing through x relaxes into B , $p_{\text{TP}}(x)$ quantifies the probability that a trajectory passing through x is a transition pathway. For a given configuration x , $p_{\text{TP}}(x)$ is estimated by generating $N_{\text{TP}} = 100$ trajectories from x by integrating the equations of motion (Hamilton’s equations) forward and backwards in time with initial velocities sampled from the Maxwell-Boltzmann distribution. In the diffusive regime, Hummer [52] showed that p_{TP} is determined solely by p_B :

$$p_{\text{TP}}(x) = 2p_B(x) [1 - p_B(x)] \tag{5.5}$$

We compare this analytic result to correlations we observe between p_{TP} and p_B when using Hamiltonian dynamics.

In Fig. 5.9 we display a scatter plot of p_{TP} vs p_B for equilibrium configurations constrained to $r_{\text{ion}} = r_{\text{ion}}^*$ (data from Fig. 5.4), plotted against the analytic result, Eq. 5.5, for diffusive dynamics. While for configurations close to $p_B = 0$ and 1 we see similar behavior between the two regimes, near the transition state p_{TP} is enhanced relative to diffusive behavior. Hence for Hamiltonian dynamics, inertial effects enhance the reaction probability near the transition state by up to 40-50%. This is intuitive: Under Hamiltonian dynamics, when the system evolves from A and up to the transition state, the probability to complete the transition by moving down to the reactant region B will be influenced by the instantaneous value of the momenta at the top of the free energy barrier. However, under diffusive dynamics this enhancement is not present simply because the momenta are equilibrated instantaneously, providing no means to help push the system to the other side.

It is known that the rate of a given reaction is dependent upon the isotopic composition of the reactants and products [2, 15], which ultimately quantifies the influence of molecular mass on a given reaction. We suspect that the importance of inertia on ionic dissociation can be further refined by investigating the effect of isotopic substitution of Na^+ and Cl^- on transition probabilities.

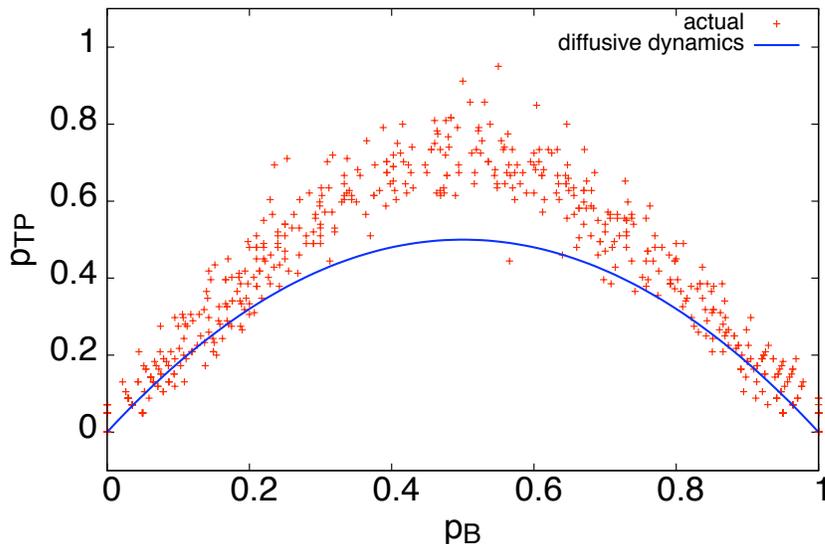


Figure 5.9: Inertial effects near the transition state. The calculated transition path probability p_{TP} is plotted against the committor probability p_{B} for configurations constrained to $r_{\text{ion}} = r_{\text{ion}}^*$. Our results under Hamiltonian dynamics, shown in red, show a deviation of the observed p_{TP} from the analytic result under diffusive dynamics [52]. These inertial effects enhance p_{TP} near the transition state.

5.7 Discussion and conclusions

In this chapter we investigated the thermodynamics and dynamics leading to dissociation Na^+Cl^- in water. We showed that the thermodynamics of dissociation is driven

energetically, and opposed entropically, with the loss of entropy explained by an increasing number of highly-coordinated solvent molecules in the first solvation shell as the ions separate. By performing committor analysis on the system with various constraints, we showed that a) the interionic distance is an insufficient reaction coordinate, in accordance with previous findings, b) the influence of the solvent on ionic dissociation is long-ranged, extending out into the third solvation shell, and c) both steric effects and electrostatics contribute to the system’s commitment to dissociation. We also highlighted the timescales under which solvent fluctuations influence dissociation, as well as the importance of inertial effects near the transition state.

In characterizing the kinetic pathway to dissociation the ultimate goal is to find a good reaction coordinate for the system. Despite the seeming simplicity of the solute, two atoms, finding an accurate description of the solvent is a difficult task. Previous attempts have shown correlations between Na^+Cl^- dissociation and solvation numbers [43] and other orientational indicators of local solvent density [43, 76, 79]. Geissler *et al.* [43] have suggested a mechanism whereby dissociation is accompanied by insertion of a water molecule from the bulk into the first solvation shell, preceded by a buildup of water density in the second solvation shell and a depletion between the second and third shells at the transition state. This picture is consistent with our findings that dissociation is sensitive to solvent rearrangements into the third solvation shell.

We have applied, with little success, a maximum likelihood approach akin to Ref. [93] to find an optimal reaction coordinate that depends on these and other solvent variables sensitive to solvent density rearrangements. We have also found weak correlations between p_B and a) the net solvent dipole along the interionic axis as well as b) the net solvent force along the interionic axis. The microscopic mechanism leading to ionic dissociation,

however, is still not completely known, and more study is needed.

Summary and outlook

Despite the recent growth of computer power and resources, efficient and effective sampling of complex systems remains a challenge. In the first part of this thesis we introduced a simulation method which can efficiently sample equilibrium distributions by incorporating nonequilibrium trajectories. Our method draws upon precise quantitative relations between processes driven out of equilibrium and their corresponding equilibrium states. The method which we developed, Replica Exchange with Nonequilibrium Switches (RENS), is a replica exchange strategy in which trial swap configurations of neighboring replicas are generated by driving the replicas out of equilibrium through finite-time switching simulations. Although these trial configurations are representative of ensembles out of equilibrium, with a work-based acceptance criteria the sampling within each replica preserves the equilibrium distribution.

With any replica exchange strategy, an effective implementation demands that neighboring replicas share a decent overlap in phase space. Our method is able to directly address this overlap requirement which often plagues replica exchange simulations. by devoting time to driving the replicas out of equilibrium, we can increase an effec-

tive phase space overlap, leading to an increased exchange acceptance rate and enhanced sampling quality of our output trajectory. Furthermore, the duration of our switching simulations, τ , is a parameter of the simulation, which can be used to tune a desired overlap. Whereas in ordinary replica exchange (REM) this overlap is achieved by adding replicas, with RENS replica overlap is enhanced through nonequilibrium switching simulations. In Chapter 3 we tested RENS on two model systems, finding it to be beneficial in the low overlap regime and able to match the efficiency of REM using fewer replicas.

In Chapters 2 and 3 we focused on the development of RENS as an efficient method to sample complex systems. In Chapter 4 we used RENS to study a different problem, namely the calculation of the solvation free energy ΔF_{solv} of a small molecule. Here we embed solute solvation within a larger thermodynamic cycle which involves cavity growth [10]. Along this cycle a progression of replicas represents a) cavity growth in absence of the solute, b) insertion of the solute in the cavity, and c) shrinking of the cavity in presence of the solute. Because the solute is now inserted in a cavity, the calculation of the insertion free energy, typically slowly convergent for simple Widom insertion [127], now becomes more manageable. Whereas in our previous studies the replicas were heated and cooled, here the switching simulations involve growing/shrinking of a solvent cavity. We used this method to calculate ΔF_{solv} associated with the small organic molecule monoethanolamine in water.

There are a number of future projects involving development of RENS as a simulation method as well as its applications on various systems. We briefly outline these three projects below: 1. We have demonstrated numerically that our method can bring efficiency gains when a switching protocol is chosen prudently. In general, however, it is impossible to know a priori what the most efficient protocol is. To this end it would be

useful to develop an adaptive version of RENS in which the switching time can be tuned on the fly. If this is possible, one could in principle achieve an optimal protocol by simply adjusting τ in response to the current simulation performance. 2. Along similar lines, it may be advantageous to develop a version of RENS with feedback control, in which a switching protocol is chosen from a few candidates based upon the current state of the system. If, for instance, during the course of a switching simulation the accumulated work is large, a longer switching time would be chosen for the remainder of the switching simulation in order to reduce net dissipation. This hypothetical method has a theoretical justification, as recently fluctuation theorems have been discovered for systems evolving under feedback control [49]. 3. As we have discussed in Sec. 2.4, RENS can be used in conjunction with escorted dynamics [119,121] in which the system evolution is coupled to changes in the control parameter λ . We described one such example in which switching simulations involving heating / cooling protocols are accompanied by scaling of velocities. It would be useful to implement such escorted dynamics within the context of the solvation free energy study of Chapter 4. In particular, as the solvent cavity size increases during the switching simulations, the water molecules near the surface could simultaneously be escorted outside the cavity. Previous studies have shown that this technique can provide accurate and efficient estimates of solvation free energies for isotropic solutes [58,121]. Refined dynamical schemes could be developed for more complicated solutes which take into account non-uniform solvent density near the surface of the solute.

In the second part of the thesis, distinct from the above studies, we investigated the dynamics of rare events. We studied the dissociation of Na^+Cl^- in water, with the specific aim of determining the role of the surrounding solvent in promoting the dissociation of the ion pair. We first studied the thermodynamics of dissociation, finding the process

to be energetically-favorable and opposed entropically; we argued that the decrease of entropy upon dissociation can be attributed to an increasingly large number of restricted solvent molecules which join the first solvation shell. Finally, we used committor analysis to investigate various dynamical aspects related to dissociation. Our study revealed that the solvent molecules can influence ionic dissociation at long ranges, and that sterics as well as electrostatics are important for the dissociation event. Ultimately, we hope to find a reaction coordinate that can describe this dissociation process. Future studies will focus on the determination of collective variables of the solvent which are responsible for this reaction.

Derivation of RENS for deterministic dynamics with Andersen thermostat

Here we extend our derivation of RENS to allow for stochastic updates in the system's evolution. For specificity, we assume that our dynamics are given by a combination of deterministic equations of motion (e.g. Hamiltonian dynamics), in conjunction with the Andersen thermostat, in a manner described below. We assume that the deterministic equations satisfy the assumptions in Sec. 2.2, namely that they are symmetric under time-reversal. For these dynamics we define reduced work by Eq. A-8 below, and we argue that with this definition detailed balance is satisfied in the sense described by Eq. 2.15, which, as we have shown in the above derivations implies that \mathcal{R}_A and \mathcal{R}_B sample their respective equilibrium distributions.

In the Andersen thermostating scheme, the momenta of randomly selected particles are re-assigned from a Maxwell-Boltzmann distribution. To keep our analysis simple, we will at first assume a single Andersen update at time $t = s$ during the switching simulation,

and we will use the notation

$$\gamma_A : x_0 \longrightarrow x_s \Longrightarrow x'_s \longrightarrow x_\tau \quad (\text{A-1})$$

to denote a switching simulation in replica A . The notation is interpreted as follows. The trajectory evolves deterministically from $t = 0$ to $t = s$; this is denoted by the first arrow on the right side of Eq. A-1. Then the momentum of the i 'th particle (selected randomly) is replaced by a new momentum sampled from the Maxwell-Boltzmann distribution: $\mathbf{p}_i \Rightarrow \mathbf{p}'_i$, where $p_{\text{MB}}(\mathbf{p}'_i; \lambda_s^A) \propto \exp[-p_i'^2/2m_i k_B T \lambda_s^A]$. Finally, deterministic evolution continues from $t = s$ to $t = \tau$. In replica B we generate a trajectory

$$\gamma_B : y_\tau \longleftarrow y'_{\tau-s} \longleftarrow y_{\tau-s} \longleftarrow y_0, \quad (\text{A-2})$$

with a similar interpretation. Since the protocols for varying λ in the two replicas are related by time-reversal, $\lambda_t^A = \lambda_{\tau-t}^B$, the Andersen updates in Eqs. A-1 and A-2 occur at the same value of work parameter, which we will denote λ^* .

The probability to generate a given trajectory γ_A in \mathcal{R}_A , conditioned on the initial state x_0 , can be written as

$$\pi_A[\gamma_A] = \pi_A(x_0 \rightarrow x_s) \cdot \pi_A(x_s \Rightarrow x'_s) \cdot \pi_A(x'_s \rightarrow x_\tau). \quad (\text{A-3})$$

The first and third factors are analogous to Eq. 7a, and describe evolution during the two deterministic intervals; and

$$\pi_A(x_s \Rightarrow x'_s) \propto \exp(-p_i'^2/2m_i k_B T \lambda^*). \quad (\text{A-4})$$

Similar expressions hold for $\pi_B[\gamma_B]$.

Now consider a trajectory γ_A in \mathcal{R}_A (Eq. A-1), along with the time-reversed version of that trajectory, which we denote

$$\tilde{\gamma}_B : \bar{x}_0 \longleftarrow \bar{x}_s \longleftarrow \bar{x}'_s \longleftarrow \bar{x}_\tau. \quad (\text{A-5})$$

The conditional probabilities for these two trajectories are related as follows:

$$\pi_A[\gamma_A] = \frac{\pi_B[\tilde{\gamma}_B]}{J_A} \frac{\pi_A(x_s \Rightarrow x'_s)}{\pi_B(\bar{x}_s \Leftarrow \bar{x}'_s)} = \pi_B[\tilde{\gamma}_B] e^{-q_A[\gamma_A]}, \quad (\text{A-6})$$

where

$$q_A[\gamma_A] = h(x'_s; \lambda^*) - h(x_s; \lambda^*) + \ln J_A \quad (\text{A-7})$$

and $J_A = (T_B/T_A)^{N_d/2}$. In Eq. A-6, the factor $\pi_B[\tilde{\gamma}_B]/J_A$ was obtained as in Eq. 10a; and the remaining factor was evaluated using Eq. A-4. Here N_d denotes the number of degrees of freedom in the system, not the number of particles.

The quantity $h(x'_s; \lambda^*) - h(x_s; \lambda^*)$ in Eq. A-7 is the change in h that accompanies the Andersen update $\mathbf{p}_i \Rightarrow \mathbf{p}'_i$. Interpreting $q_A[\gamma_A]$ as the reduced heat absorbed by the system during the trajectory γ_A , we define the reduced work:

$$w_A[\gamma_A] = h(x_\tau; \lambda_\tau) - h(x_0; \lambda_0) - q_A[\gamma_A]. \quad (\text{A-8})$$

Eqs. A-6 and A-8 are analogues of Eqs. 10 and 4 of the derivation for deterministic dynamics, and Eq. A-6 is equivalent to Crooks's result, Eq. 9 of Ref. [18]. With these equations – and similar ones for a trajectory γ_B generated in \mathcal{R}_B (Eq. A-2) – we now establish detailed balance for the accepted transitions $(x, y) \rightarrow (y', x')$, in a manner analogous to Eq. 2.15 of the main text:

$$\begin{aligned} K_{acc}(y', x'|x, y) &= \pi_A[\gamma_A] \pi_B[\gamma_B] \alpha[\gamma_A, \gamma_B] \\ &= \pi_B[\tilde{\gamma}_B] \pi_A[\tilde{\gamma}_A] e^{-q_A[\gamma_A] - q_B[\gamma_B]} \alpha[\tilde{\gamma}_A, \tilde{\gamma}_B] e^{-w_A[\gamma_A] - w_B[\gamma_B]} \\ &= K_{acc}(\bar{x}, \bar{y}|\bar{y}', \bar{x}') e^{-h_A(y') - h_B(x') + h_A(x) + h_B(y)}. \end{aligned} \quad (\text{A-9})$$

Here, γ_A is the unique trajectory that starts at x_0 and ends at x_τ , and similarly for γ_B ; and $\alpha = \min\{1, e^{-w}\}$ where $w = w_A + w_B$.

In obtaining this result we have made the simplifying assumption that only a single Andersen update is made during the work simulation. We now sketch a derivation for the more general case in which K Andersen updates are performed, at times s_1, \dots, s_K , and the evolution between these updates is deterministic and time-reversal symmetric (in the sense discussed in the text), but otherwise quite general. Under these assumptions we can obtain the following expression for the relative probabilities of generating γ_A in \mathcal{R}_A and its time-reverse $\tilde{\gamma}_B$ in \mathcal{R}_B :

$$\pi_A[\gamma_A] = \pi_B[\tilde{\gamma}_B] e^{-q_A[\gamma_A]}, \quad (\text{A-10})$$

where

$$q_A[\gamma_A] = \ln J_A[\gamma_A] + \sum_{k=1}^K [h(x'_{s_k}; \lambda_{s_k}^A) - h(x_{s_k}; \lambda_{s_k}^A)] \quad (\text{A-11})$$

and J_A is the product of Jacobians along the deterministic intervals of the trajectory:

$$J_A[\gamma_A] = \left| \frac{\partial x_\tau}{\partial x'_{s_K}} \right| \cdot \left| \frac{\partial x_{s_K}}{\partial x'_{s_{K-1}}} \right| \cdots \left| \frac{\partial x_{s_1}}{\partial x_0} \right|. \quad (\text{A-12})$$

(As with Eqs. 10 and A-6, Eq. A-10 corresponds to Eq. 9 of Ref. [18].) We then define reduced work as in Eq. A-8.

Since there now exist multiple trajectories that connect given initial and final points, the acceptance kernel K_{acc} will be given by an integral over all intermediate points:

$$P(y', x' | x, y) = \int \mathcal{D}'\gamma_A \int \mathcal{D}'\gamma_B \pi_A[\gamma_A] \pi_B[\gamma_B] \alpha[\gamma_A, \gamma_B] \quad (\text{A-13})$$

where $\int \mathcal{D}'\gamma_A = \int dx_{s_1} \int dx'_{s_1} \cdots \int dx_{s_K} \int dx'_{s_K}$, and γ_A begins and ends in the fixed configurations $x_0 = x$ and $x_\tau = x'$. Analogous comments apply to $\int \mathcal{D}'\gamma_B$.

Eq. A-13 is a path-integral expression for the transition probability for the trial move $(x, y) \rightarrow (y', x')$. Writing the corresponding expression for the reverse trial move,

$(\bar{x}, \bar{y}) \leftarrow (\bar{y}', \bar{x}')$, we get

$$\begin{aligned}
K_{acc}(\bar{x}, \bar{y} | \bar{y}', \bar{x}') &= \int \mathcal{D}' \bar{\gamma}_A \int \mathcal{D}' \bar{\gamma}_B \pi_A[\bar{\gamma}_A] \pi_B[\bar{\gamma}_B] \alpha[\bar{\gamma}_A, \bar{\gamma}_B] \\
&= \int \mathcal{D}' \gamma_A \int \mathcal{D}' \gamma_B \pi_B[\gamma_B] \pi_A[\gamma_A] e^{q_A[\gamma_A] + q_B[\gamma_B]} \alpha[\gamma_A, \gamma_B] e^{w_A[\gamma_A] + w_B[\gamma_B]} \\
&= e^{h_A(y') + h_B(x') - h_A(x) - h_B(y)} \int \mathcal{D}' \gamma_A \int \mathcal{D}' \gamma_B \pi_B[\gamma_B] \pi_A[\gamma_A] \alpha[\gamma_A, \gamma_B] \\
&= e^{h_A(y') + h_B(x') - h_A(x) - h_B(y)} K_{acc}(y', x' | x, y)
\end{aligned} \tag{A-14}$$

Here we have used Eqs. A-8 and A-10, along with the correspondence between a trajectory and its time-reversed twin ($\mathcal{D}' \bar{\gamma}_A = \mathcal{D}' \gamma_B$, etc.).

This derivation can be generalized further by replacing the Andersen updates with any stochastic moves that satisfy detailed balance (for instance Metropolis Monte Carlo moves). Eqs. A-10 - A-14 remain unchanged under these quite general conditions, and therefore RENS remains valid.

Appendix B

RENS implementation with the Nosé-Hoover thermostat

The Nosé-Hoover (NH) thermostat [48, 87] has become a popular algorithm for simulating thermostated systems. In this Appendix we describe how to implement RENS with systems evolving under NH equations of motion. Here the influence of the bath on the system is modeled by an extra dynamical degree of freedom ζ , a “bath variable” that interacts with the system variables. Let $x = (\{\mathbf{q}_i\}, \{\mathbf{p}_i\})$ (and later y) represent a point in the phase space of the physical system. Then $z = (x, \zeta)$ is a point in the expanded phase space consisting of both system and bath variables. In this expanded space the system evolves deterministically under the following dynamics:

$$\dot{\mathbf{q}}_i = \frac{\partial H_\lambda}{\partial \mathbf{p}_i}, \quad \dot{\mathbf{p}}_i = -\frac{\partial H_\lambda}{\partial \mathbf{q}_i} - \zeta \mathbf{p}_i \quad (\text{B-1a})$$

$$\dot{\zeta} = \frac{1}{\tau_{\text{NH}}^2} \left(\frac{K}{K_\lambda^0} - 1 \right). \quad (\text{B-1b})$$

Here $K = \sum p_i^2/2m_i$ is the kinetic energy of the system and $K_\lambda^0 = N_p k_B T_\lambda/2$ is the equilibrium average of K and N_p is the number of degrees of freedom. The parameter τ_{NH} is a relaxation time associated with the coupling of the system and the heat bath.

Intuitively, ζ can be viewed as a friction coefficient whose evolution depends upon the extent to which the kinetic energy deviates from its equilibrium average. If, for instance, the system has a kinetic energy $K > K_\lambda^0$, the variable ζ increases, intensifying the damping of the momenta, and ultimately restoring the system's kinetic energy towards its thermal value.

It can be shown that the distribution

$$p^{\text{eq}}(x, \zeta; \lambda) \propto \exp[-h_\lambda(x, \zeta)], \quad (\text{B-2})$$

is stationary under the NH dynamics (Eq. B-1), where the reduced Hamiltonian in this expanded phase space

$$h_\lambda(x, \zeta) = h_\lambda^{\text{sys}}(x) + \frac{1}{2}N\tau_{\text{NH}}^2\zeta^2. \quad (\text{B-3})$$

contains $h_\lambda^{\text{sys}}(x)$, the reduced Hamiltonian of the physical system, as well as a term that depends solely on the bath variable ζ . In the space of system variables x , then, a trajectory samples the canonical distribution.

The NH equations are deterministic and time-reversible in our expanded phase space, satisfying our assumptions for dynamical schemes for RENS described in Sec. 2.2. Let us use these dynamics to drive our replicas out of equilibrium: During the course of a switching simulation, as λ is switched from 0 to 1 in \mathcal{R}_A , the Hamiltonian H and/or the temperature T of \mathcal{R}_A are switched from H_A, T_A to H_B, T_B . This generates a trajectory

$$\gamma_A : (x_0, \zeta_0^A) \rightarrow (x_\tau, \zeta_\tau^A). \quad (\text{B-4a})$$

Similarly, in \mathcal{R}_B a trajectory γ_B is generated as λ is switched from 1 to 0:

$$\gamma_B : (y_0, \zeta_0^B) \leftarrow (y_\tau, \zeta_\tau^B). \quad (\text{B-4b})$$

The dynamics that generate these trajectories do not preserve phase space volume:

$$\nabla \cdot (\dot{x}, \dot{\zeta}) = \sum_i \frac{\partial}{\partial \mathbf{q}_i} \dot{\mathbf{q}}_i + \sum_i \frac{\partial}{\partial \mathbf{p}_i} \dot{\mathbf{p}}_i + \frac{\partial}{\partial \zeta} \dot{\zeta} = -N_p \zeta \neq 0. \quad (\text{B-5})$$

The Jacobian associated with \mathcal{R}_A , for instance, is then

$$J_A(\gamma_A) = e^{-N_p \int_0^\tau dt \zeta_t^A}. \quad (\text{B-6})$$

The work performed, given by Eq. 2.5, is

$$w_A(\gamma_A) = h_B(x_\tau, \zeta_\tau^A) - h_A(x_0, \zeta_0^A) + N_p \int_0^\tau dt \zeta_t^A. \quad (\text{B-7a})$$

Similarly, for γ_B in \mathcal{R}_B the work is

$$w_B(\gamma_B) = h_A(y_\tau, \zeta_\tau^B) - h_B(y_0, \zeta_0^B) + N_p \int_0^\tau dt \zeta_t^B. \quad (\text{B-7b})$$

In the NH scheme above the system variables x are thermostated through a single bath variable. It is often desirable to implement an entire *chain* of variables: ζ_1 is thermostated by another bath variable ζ_2 , which in turn is coupled to ζ_3 and so on. This method of Nosé-Hoover chains [77] provides a way to overcome cumbersome non-ergodic sampling issues often found with normal NH dynamics [41]. Although not discussed here in detail, it is straightforward to follow the steps above to identify a reduced hamiltonian $h(x, \zeta_1, \dots, \zeta_C)$ and Jacobian J_A for NH dynamics with chains, from which the work w_A can be defined.

Appendix C

Code development

A considerable amount of coding was involved in order to test RENS on alanine dipeptide and monoethanolamine. In this Appendix we describe the major steps involved in producing simulation software capable of RENS simulations.

The majority of our RENS simulations were performed with the OpenMM package, a library of classes and functions pertaining to molecular simulation. The RENS code we developed is largely based upon an existing replica exchange routine written in python, and used in conjunction with OpenMM via the python wrapper for OpenMM, PyOpenMM. Shortly we plan to release these OpenMM code alterations, described in the next paragraphs, to the simulation community for general use.

Two major modifications to the existing replica exchange routine were needed to allow for RENS simulations. The first modification was parallelization of the simulation code. The original development, used mostly for testing purposes, had involved serial integration of replicas on a single processor, where evolution of replica $i + 1$ could only proceed after replica i had evolved. This was impractical for implementations involving larger numbers of replicas and larger system sizes. In order to take full advantage of

multiple-processor CPUs, we parallelized the computer code such that each replica was devoted its own processor, as is standard in replica exchange programs. The parallelization was implemented with `mpi4py`¹, a MPI (Message Passing Interface) package for python allowing for initialization of parallel routines and communication between processors.

The second and most involved component of RENS code development was the incorporation of the switching simulations into the replica exchange code. While practically every simulation package provides “equilibrium integration” schemes for Hamiltonian (Molecular) dynamics, Langevin dynamics and other dynamics, the incorporation of dynamical schemes with time-dependent parameters is largely absent in standard simulation packages. To perform the RENS switching simulations, which involve time-dependent parameters defined through the control parameter $\lambda(t)$ (such as temperature), we incorporated a “switching routine” into the outer (python) code level, consisting of a loop through the following steps:

1. integrate the system using an equilibrium integrator initialized at a current value of λ
2. update λ and all λ -dependent parameters: $\lambda \rightarrow \lambda + \Delta\lambda$
3. reinitialize the integrator to the new λ parameters for the next loop iteration.

Hence the switching trajectory was achieved by integrating the system through a series of equilibrium integrators with increasing λ values. Because step 3 is particularly costly computationally, these switching simulations were limited (in a practical sense) to a relatively small number $n = 1/\Delta\lambda$ of iterations of the above steps, which restricted us to switching protocols involving larger jumps ($\Delta\lambda = \mathcal{O}(0.1)$). As discussed in Sec. 3.3.3, this

¹<http://mpi4py.scipy.org>

limited the gains RENS can bring to the simulation efficiency.

Nonequilibrium methods such as RENS will only be accessible to a broader audience once standard simulations packages (such as Gromacs, LAMMPS, and OpenMM) incorporate “switching integrators” which are initialized once with an input protocol $\lambda(t)$, from which all updates of λ are performed at a base code level². With growing interest in the use of nonequilibrium methods in calculation of thermodynamic properties, incorporation of such routines will become increasingly important. (In our opinion, much research could be done even in the development of accurate integration schemes for systems evolving with time-dependent parameters.)

Another coding project that was part of my graduate research (but not contained in the main thesis text) was the development of the LAMMPS simulation package [95], specifically an expansion of the Nosé-Hoover (NH) thermostating routine to include Nosé-Hoover chains. The NH equations, introduced in Appendix B, are a deterministic set of dynamics in which the system is coupled to an auxiliary bath variable ζ . It is known that the NH thermostat in its original implementation does not sample the desired canonical ensemble for systems that are not sufficiently chaotic [48]. This problem can be rectified by coupling the system to a series or “chain” of bath variables instead of only one, which together are able to achieve system ergodicity [77]. The project consisted of expanding the LAMMPS NH thermostat routine to allow for chains, to ensure proper sampling of the desired NVT ensemble. We used a discretization scheme for the dynamics developed by Martyna et al [78].

²The OpenMM package is taking a step in this direction: Their CustomIntegrator class, currently in development, allows for the building of custom integration routines from a set of instructions provided by the user.

Bibliography

- [1] M.R.M Abu-Zahra, L.H.J. Schneiders, J.P.M. Niederer, P.H.M. Feron, and G.F. Versteeg. Co₂ capture from power plants: Part i. a parametric study of the technical performance based on monoethanolamine. *International Journal of Greenhouse Gas Control*, 1:37–46, 2007.
- [2] M.H. Alexander. Chemical kinetics under test. *Science*, 331:411–412, 2011.
- [3] H.C. Andersen. Molecular dynamics simulations at constant pressure and/or temperature. *J. Chem. Phys.*, 72:2384–2393, 1980.
- [4] H-J. Arpe. *Industrial Organic Chemistry*. Wiley, Weinheim, Germany, 2010.
- [5] D. Asthagiri, S. Merchant, and L.R. Pratt. Role of attractive methane-water interactions in the potential of mean force between methane molecules in water. *J. Chem. Phys.*, 128:244512, 2008.
- [6] M. Athènes. Computation of a chemical potential using a residence weight algorithm. *Phys. Rev. E*, 66:046705, 2002.
- [7] A.J. Ballard and C. Jarzynski. Replica exchange with nonequilibrium switches. *Proc. Natl. Acad. Sci. U.S.A.*, 106:12224–12229, 2009.
- [8] A.J. Ballard and C. Jarzynski. Replica exchange with nonequilibrium switches: Enhancing equilibrium sampling by increasing replica overlap. *J. Chem. Phys.*, 136:194101, 2012.
- [9] R. Baron, P. Setny, and J.A. McCammon. Water in cavity-ligand recognition. *J. Am. Chem. Soc.*, 132:12091–12097, 2010.
- [10] T.L. Beck, M.E. Paulaitis, and L.R. Pratt. *The Potential Distribution Theorem and Models of Molecular Solutions*. Cambridge University Press, Cambridge, UK, 2006.
- [11] H. Bekker, H.J.C. Berendsen, E.J. Dijkstra, S. Achterop, R. van Drunnen, D. van der Spoel, A. Sijbers, H. Keegstra, B. Reitsma, and M.K.R. Renardus. *Gromacs: A parallel computer for molecular dynamics simulations*. World Scientific, Singapore, 1993.
- [12] A.C. Belch, M. Berkowitz, and J.A. McCammon. Solvation structure of a sodium chloride ion pair in water. *J. Am. Chem. Soc.*, 108:1755–1761, 1986.
- [13] C.H. Bennett. Efficient estimation of free energy differences from monte carlo data. *J. Comput. Phys.*, 22:245, 1976.

- [14] M. Berkowitz, O.A. Karim, J.A. McCammon, and P.J. Rossky. Sodium chloride ion pair interaction in water: Computer simulation. *Chem. Phys. Lett.*, 105:577–580, 1984.
- [15] J. Bigeleisen and M.G. Mayer. Calculation of equilibrium constants for isotopic exchange reactions. *J. Chem. Phys.*, 15:261, 1947.
- [16] K. Binder and D.W. Heermann. *Monte Carlo Simulation in Statistical Physics*. Springer, Berlin, 2002.
- [17] P.G. Bolhuis, D. Chandler, C. Dellago, and P.L. Geissler. Transition path sampling: Throwing ropes over rough mountain passes, in the dark. *Annu. Rev. Phys. Chem.*, 53:291–318, 2002.
- [18] S. Brown and T. Head-Gordon. Cool-walking: A new markov chain monte carlo sampling method. *J. Computat. Chem.*, 24:68–76, 2003.
- [19] G. Buemi. Conformational analysis and rotation barriers of 2-aminoethanethiol and 2-aminoethanol: An ab initio study. *Int. J. Quantum Chem.*, 59:227–237, 1996.
- [20] C. Caleman, P. J. van Maaren, M. Hong, J. S. Hub, L. T. Costa, and D. van der Spoel. Force field benchmark of organic liquids: Density, enthalpy of vaporization, heat capacities, surface tension, isothermal compressibility, volumetric expansion coefficient, and dielectric constant. *J. Chem. Theory Comput.*, 8:61–74, 2011.
- [21] F. Calvo and J.P.K. Doyle. Entropic tempering: A method for overcoming quasiergodicity in simulation. *Phys. Rev. E*, 63:010902, 2000.
- [22] D. Chandler. Interfaces and the driving force of hydrophobic assembly. *Nature*, 437:640–647, 2005.
- [23] Y.-P. Chang and T.-M. Su. Global conformational analysis and the anomeric interactions of methanediol, methanediamine and aminomethanol. *J. Mol. Struct. (THEOCHEM)*, 365:183, 1996.
- [24] Y.-P. Chang, T.-M. Su, T.-W. Li, and I. Chao. Intramolecular hydrogen bonding, gauche interactions, and thermodynamic functions of 1,2-ethanediamine, 1,2-ethanediol, and 2-aminoethanol: A global conformational analysis. *J. Phys. Chem. A*, 101:6107, 1997.
- [25] C. Chipot and A. Pohorille. *Free Energy Calculations: Theory and Applications in Chemistry and Biology*. Springer, Berlin, 2007.
- [26] G. E. Crooks. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Phys. Rev. E*, 60:2721–2726, 1999.
- [27] G.E. Crooks. Nonequilibrium measurements of free energy differences for microscopically reversible markovian systems. *J. Stat. Phys.*, 90:1481–1487, 1998.
- [28] G.E. Crooks. Path-ensemble averages in systems driven far from equilibrium. *Phys. Rev. E*, 61:2361–2366, 2000.

- [29] E.F. da Silva and H.F. Svendsen. Prediction of the pK_a values of amines using ab initio methods and free-energy perturbations. *Ind. Eng. Chem. Res.*, 42:4414–4421, 2003.
- [30] R.M. Dirks, H. Xu, and D. E. Shaw. Improving sampling by exchanging hamiltonians with efficiently configured nonequilibrium simulations. *J. Chem. Theory Comput.*, 8:162–171, 2012.
- [31] C.M. Dobson. Protein folding and misfolding. *Nature*, 426:884–890, 2003.
- [32] W. E, W. Ren, and E. Vanden-Eijnden. Transition pathways in complex systems: Reaction coordinates, isocommittor surfaces, and transition tubes. *Chem. Phys. Lett.*, 413:242, 2005.
- [33] D.J. Earl and M. W. Deem. Parallel tempering: Theory, applications, and new perspectives. *Phys. Chem. Chem. Phys.*, 7:3910–3916, 2005.
- [34] R. Faller, Q. Yan, and J.J. de Pablo. Multicanonical parallel tempering. *J. Chem. Phys.*, 116:5419–5423, 2002.
- [35] C. F. Fenell, A. Bizjak, V. Vlachy, and K. A. Dill. Ion pairing in molecular simulations of aqueous alkali halide solutions. *J. Phys. Chem. B*, 113:6782, 2009.
- [36] M.K. Fenwick and F.A. Escobedo. Expanded ensemble and replica exchange methods for simulation of protein-like systems. *J. Chem. Phys.*, 119:11998–12010, 2003.
- [37] E. Fermi. *Thermodynamics*. Dover, New York, 1956.
- [38] Alan M. Ferrenberg and Robert H. Swendsen. Optimized monte carlo data analysis. *Phys. Rev. Lett.*, 63:1195–1198, 1989.
- [39] H. Flyvbjerg and H.G. Petersen. Error estimates on averages of correlated data. *J. Chem. Phys.*, 91:461–466, 1989.
- [40] D. Frenkel. Speed-up of monte carlo simulations by sampling of rejected states. *Proc. Natl. Acad. Sci. U.S.A.*, 101:17571–17575, 2004.
- [41] D. Frenkel and B. Smit. *Understanding Molecular Simulation*. Academic Press, San Diego, 2002.
- [42] M.S. Friedrichs, P. Eastman, V. Vaidyanathan, M. Houston, S. LeGrand, A.L. Berberg, D.L. Ensign, C.M. Burns, and V.S. Pande. Accelerating molecular dynamic simulation on graphics processing units. *J. Comp. Chem.*, 30:864–872, 2009.
- [43] P.L. Geissler, C. Dellago, and D. Chandler. Kinetic pathways of ion pair dissociation in water. *J. Phys. Chem. B*, 103:3706–3710, 1999.
- [44] S. Gnanakran, H. Nymeyer, J. Portman, K.Y. Sanbonmatsu, and A.E. Garcia. Peptide folding simulations. *Curr Opin Struc Biol*, 13:168–174, 2003.
- [45] W. F. Van Gunsteren and H. J. C. Berendsen. A leap-frog algorithm for stochastic dynamics. *Mol. Simulat.*, 1:173–185, 1988.

- [46] A. N. Gupta, A. Vincent, K. Neupane, H. Yu, F. Wang, and M.T. Woodside. Experimental validation of free-energy-landscape reconstruction from non-equilibrium single-molecule force spectroscopy measurements. *Nature Physics*, 7:631–634, 2011.
- [47] U.H.E. Hansmann. Parallel tempering algorithm for conformational studies of biological molecules. *Chem. Phys. Lett.*, 281:140–150, 1997.
- [48] W.G. Hoover. Canonical dynamics: Equilibrium phase-space distributions. *Phys. Rev. A*, 31:1695–1697, 1984.
- [49] J.M. Horowitz and S. Vaikuntanathan. Nonequilibrium detailed fluctuation theorem for repeated discrete feedback. *Phys. Rev. E*, 82:061120, 2010.
- [50] X. H. Huang, M. Hagan, B. Kim, R. A. Friesner, R. H. Zhou, and B. J. Berne. Replica exchange with solute tempering: Efficiency in large scale systems. *J. Phys. Chem. B*, 115:5405–5410, 2007.
- [51] K. Hukushima and K. Nemoto. Exchange monte carlo method and application to spin glass simulations. *J. Phys. Soc. Jpn.*, 65:1604–1608, 1996.
- [52] G. Hummer. From transition paths to transition states and rate coefficients. *J. Chem. Phys.*, 120:516–523, 2004.
- [53] G. Hummer and A. Szabo. Free energy reconstruction from nonequilibrium single-molecule pulling experiments. *Proc. Natl. Acad. Sci. U.S.A.*, 98:3658–3661, 2001.
- [54] S. Jang, S. Shin, and Y. Pak. Replica-exchange method using the generalized effective potential. *Phys. Rev. Lett.*, 91:058305, 2003.
- [55] C. Jarzynski. Equilibrium free-energy differences from nonequilibrium measurements: A master-equation approach. *Phys. Rev. E*, 56:5018–5035, 1997.
- [56] C. Jarzynski. Nonequilibrium equality for free energy differences. *Phys. Rev. Lett.*, 78:2690–2693, 1997.
- [57] C. Jarzynski. Microscopic analysis of clausius-duhem processes. *J. Stat. Phys.*, 96:415–427, 1999.
- [58] C. Jarzynski. Targeted free energy perturbation. *Phys. Rev. E*, 65:1–5, 2002.
- [59] C. Jarzynski. Equalities and inequalities: Irreversibility and the second law of thermodynamics at the nanoscale. *Annu. Rev. Condens. Matter Phys.*, 2:329–351, 2010.
- [60] W.L. Jorgensen, D.S. Maxwell, and J. Tirado-Rives. Development and testing of the opls all-atom force field on conformational energetics and properties of organic liquids. *J. Am. Chem. Soc.*, 118:11225–11236, 1996.
- [61] W.L. Jorgenson, J. Chandrasenkar, R.W. Impey, and M.L. Klein. Comparison of simple potential functions for simulating liquid water. *J. Chem. Phys.*, 79:926–936, 1983.
- [62] A.-M. Kelterer, M. Ramek, F. F. Regina, M. Cao, and L. Schafer. Basis set influence in ab initio calculations: the case of 2-aminoethanol and n-formylproline amide. *J. Mol. Struct. (THEOCHEM)*, 310:45, 1994.

- [63] J. G. Kirkwood. Statistical mechanics of fluid mixtures. *J. Chem. Phys.*, 3:300, 1935.
- [64] D.A. Kofke. On the acceptance probability of replica-exchange monte carlo trials. *J. Chem. Phys.*, 117:6911–6914, 2002.
- [65] P. A. Kollman, R. Dixon, W. Cornell, T. Fox, C. Chipot, and A. Pohorille. The development/application of a ‘minimalist’ organic/biochemical molecular mechanic force field using a combination of ab initio calculations and experimental data. In W.F. van Gunsteren, P.K. Weiner, and A.J. Wilkinson, editors, *Computer simulation of biomolecular systems: theoretical and experimental applications*. Kluwer Academic Publishers, Dordrecht, The Netherlands, 1997.
- [66] S. Kumar, J.M. Rosenberg, D. Bouzida, R.H. Swendsen, and P.A. Kollman. The weighted histogram analysis method for free-energy calculations on biomolecules. i. the method. *J. Comput. Chem.*, 13:1011–1021, 1992.
- [67] L.D. Landau and E.M. Lifshitz. *Statistical Physics*. Butterworth-Heinemann, Oxford, 2000.
- [68] W. Lechner, H. Oberhofer, C. Dellago, and P.L. Geissler. Equilibrium free energies from fast-switching trajectories with large time steps. *J. Chem. Phys.*, 124:044113, 2006.
- [69] T. Lelièvre, M. Rousset, and G. Stoltz. *Free Energy Computations: A Mathematical Perspective*. Imperial College Press, London, 2010.
- [70] H. Li, M. Fajer, and W. Yang. Simulated scaling method for localized enhanced sampling and simultaneous alchemical free energy simulations: A general method for molecular mechanical, quantum mechanical, and quantum mechanical/molecular mechanical simulations. *J. Chem. Phys.*, 126:024106, 2007.
- [71] P. Liu, B. Kim, R.A. Friesner, and B.J. Berne. Replica exchange with solute tempering: A method for sampling biological systems in explicit water. *Proc. Natl. Acad. Sci. U.S.A.*, 102:13749–13754, 2005.
- [72] R. Lòpez-Rendòn, M.A. Mora, J. Alejandre, and M.E. Tuckerman. Molecular dynamics simulations of aqueous solutions of ethanolamines. *J. Phys. Chem. B*, 110:14652–14658, 2006.
- [73] A.P. Lyubartsev, A.A. Martsinovski, S.V. Shevkunov, and P.N. Vorontsovveliaminov. New approach to monte carlo calculation of the free energy: Method of expanded ensembles. *J. Chem. Phys.*, 96:1775–1783, 1992.
- [74] A. Ma and A.R. Dinner. Automatic method for identifying reaction coordinates in complex systems. *J. Phys. Chem. B*, 109:6769–6779, 2005.
- [75] E. Marinari and G. Parisi. Simulated tempering: a new monte carlo scheme. *Europhys. Lett.*, 19:451–458, 1992.
- [76] J. Marti, F.S. Csajka, and D. Chandler. Stochastic transition pathways in the aqueous sodium chloride dissociation process. *Chem. Phys. Lett.*, 328:169–176, 2000.

- [77] G. Martyna, M.L. Klein, and M.E. Tuckerman. Nosè-hoover chains: The canonical ensemble via continuous dynamics. *J. Chem. Phys.*, 97:2635–2643, 1992.
- [78] G. Martyna, M.E. Tuckerman, D.J. Tobias, and M.L. Klein. Explicit reversible integrators for extended systems dynamics. *Mol. Phys.*, 87:1117–1157, 1996.
- [79] T.A. McCormick and D. Chandler. Grid-flux method for learning the solvent contribution to the mechanisms of reactions. *J. Phys. Chem. B*, 107:2796–2801, 2003.
- [80] N. Metropolis, A.W. Rosenbluth, M. N. Rosenbluth, A. N. Teller, and E. Teller. Equation of state calculations by fast computing machines. *J. Chem. Phys.*, 21:1087–1092, 1953.
- [81] M.A. Miller and W.P. Reinhardt. Efficient free energy calculations by variationally optimized metric scaling: Concepts and applications to the volume dependence of cluster free energies and to solidsolid phase transitions. *J. Chem. Phys.*, 113:7035–7046, 2000.
- [82] A. Mitsutake, Y. Sugita, and Y. Okamoto. Generalized-ensemble algorithms for molecular simulations of biopolymers. *Biopolymers*, 60:96–123, 2001.
- [83] R. M. Neal. Sampling from multimodal distributions using tempered transitions. *Stat. Comput.*, 6:353–366, 1996.
- [84] J. Nilmeier and M. P. Jacobson. Monte carlo sampling with hierarchical move sets: Posh monte carlo. *J. Chem. Theory Comput.*, 5:1968–1984, 2009.
- [85] J. P. Nilmeier, G. E. Crooks, D. D. L. Minh, and J. D. Chodera. Nonequilibrium candidate monte carlo is an efficient tool for equilibrium simulation. *Proc. Natl. Acad. Sci. U.S.A.*, 108:E1009–E1018, 2011.
- [86] F. Noé, C. Schütte, E. Vanden-Eijnden, L. Reich, and T.R. Weigl. Constructing the equilibrium ensemble of folding pathways from short off-equilibrium simulations. *Proc. Natl. Acad. Sci. U.S.A.*, 106:19011–19016, 2009.
- [87] S. Nosé. A unified formulation of the constant temperature molecular dynamics methods. *J. Chem. Phys.*, 81:511–519, 1984.
- [88] H. Oberhofer, C. Dellago, and P.L. Geissler. Biased sampling of nonequilibrium trajectories: can fast switching simulations outperform conventional free energy calculation methods? *J. Phys. Chem. B*, 109:6902–6915, 2005.
- [89] S. Opps and J. Schofield. Extended state-space monte carlo methods. *Phys. Rev. E*, 63:056701, 2001.
- [90] A. Paliwal, D. Asthagiri, L. R. Pratt, H. S. Ashbaugh, and M. E. Paulaitis. An analysis of molecular packing and chemical association in liquid water using quasi-chemical theory. *J. Chem. Phys.*, 124:224502, 2006.
- [91] S. Park. Comparison of the serial and parallel algorithms of generalized ensemble simulations: An analytical approach. *Phys. Rev. E*, 77:016709, 2008.
- [92] B. Peters. Using the histogram test to quantify reaction coordinate error. *J. Chem. Phys.*, 125:241101, 2006.

- [93] B. Peters and B.L. Trout. Obtaining reaction coordinates by likelihood maximization. *J. Chem. Phys.*, 125:054108, 2006.
- [94] R. Phillips, J. Kondev, and J. Thieriot. *Physical Biology of the Cell*. Garland Science, New York, 2009.
- [95] S. Plimpton. Fast parallel algorithms for short-range molecular dynamics. *J. Comp. Phys.*, 117:1–9, 1997.
- [96] A. Pohorille, C. Jarzynski, and C. Chipot. Good practices in free-energy calculations. *J. Phys. Chem. B*, 114:10235–10253, 2010.
- [97] D. Qiu, P. S. Shenkin, F. P. Hollinger, and W. C. Still. The gb/sa continuum model for solvation. a fast analytical method for the calculation of approximate born radii. 101:3005–3014, 1997.
- [98] A. B. Rao and E. S. Rubin. A technical, economic, and environmental assessment of amine-based co₂ capture technology for power plant greenhouse gas control. *Environ. Sci. Technol.*, 36:4467–4475, 2002.
- [99] N. Rathore, Manan Chopra, and J.J. de Pablo. Optimal allocation of replicas in parallel tempering simulations. *J. Chem. Phys.*, 122:024111, 2005.
- [100] R. Rey and E. Guardia. Dynamical aspects of the sodium(1+)-chloride ion pair association in water. *J. Phys. Chem.*, 96:4712–4718, 1992.
- [101] S. W. Rick. Replica exchange with dynamical scaling. *J. Chem. Phys.*, 16:054102, 2007.
- [102] D.M. Rogers and T.L. Beck. Modeling molecular and ionic absolute solvation free energies with quasichemical theory bounds. *J. Chem. Phys.*, 129:134505, 2008.
- [103] E. Rosta and G. Hummer. Error and efficiency of replica exchange molecular dynamics simulations. *J. Chem. Phys.*, 131:165102, 2009.
- [104] E. Rosta and G. Hummer. Error and efficiency of simulated tempering simulations. *J. Chem. Phys.*, 132:034102, 2010.
- [105] B. Rout. Implicit solvent models. In Oren M. Becker, Jr. Alexander D. MacKerell, Benoît Rout, and Masakatsu Watanabe, editors, *Computational Biochemistry and Biophysics*. Marcel Dekker, Inc., New York, 2001.
- [106] A. Schug, T. Herges, and W. Wenzel. All-atom folding of the three-helix hiv accessory protein with an adaptive parallel tempering method. *Proteins*, 57:792–798, 2004.
- [107] P. Setny, R. Baron, and J.A. McCammon. How can hydrophobic association be enthalpy driven? *J. Chem. Theory Comput.*, 6:2866–2871, 2010.
- [108] J. K. Shah, D. Asthagiri, L. R. Pratt, and M. E. Paulaitis. Balancing local order and long-ranged interactions in the molecular theory of liquid water. *J. Chem. Phys.*, 127:144508, 2007.

- [109] D.E. Shaw, M.M. Deneroff, R.O. Dror, J.S. Kuskin, R.H. Larson, J.K. Salmon, C. Young, B. Batson, K. J. Bowers, J.C. Chao, M.P. Eastwood, J. Gagliardo, J.P. Grossman, C.R. Ho, D.J. Ierardi, I. Kolossvry, J.L. Klepeis, T. Layman, C. McLeavey, M.A. Moraes, R. Mueller, E.C. Priest, Y. Shan, J. Spengler, M. Theobald, B. Towles, and S.C. Wang. Anton, a special-purpose machine for molecular dynamics simulation. *Communications of the ACM*, 51:91–97, 2008.
- [110] J.-E. Shea and C.L. Brooks III. From folding theories to folding proteins: A review and assessment of simulation studies of protein folding and unfolding. *Annu. Rev. Phys. Chem.*, 52:499–535, 2001.
- [111] D.A. Sivak, J.D. Chodera, and G.E. Crooks. Driven langevin dynamics: heat, work and pseudo-work. [arXiv:1107.2967v2](https://arxiv.org/abs/1107.2967v2), 2011.
- [112] D. E. Smith and A. D. J. Haymat. Free energy, entropy, and internal energy of hydrophobic interactions: Computer simulations. *J. Chem. Phys.*, 98:6445, 1993.
- [113] H.A. Stern. Molecular simulation with variable protonation states at constant ph . *J. Chem. Phys.*, 126:164112, 2007.
- [114] W. C. Still, A. Tempczyk, R. C. Hawley, and T. Hendrickson. Semianalytical treatment of solvation for molecular mechanics and dynamics. *J. Am. Chem. Soc.*, 112:6127–6129, 1990.
- [115] Y. Sugita and Y. Okamoto. Replica-exchange molecular dynamics method for protein folding. *Chem. Phys. Lett.*, 314:141–151, 1999.
- [116] R.H. Swendsen and J-S Wang. Replica monte-carlo simulation of spin-glasses. *Phys. Rev. Lett.*, 57:2607–2609, 1986.
- [117] S. Trebst, M. Troyer, and U.H.E. Hansmann. Optimized parallel tempering simulations of proteins. *J. Chem. Phys.*, 124:174903, 2006.
- [118] M. Utz, P.G. Debenedetti, and F.H. Stillinger. Atomistic simulation of aging and rejuvenation in glasses. *Phys. Rev. Lett.*, 84:1471–1474, 2000.
- [119] S. Vaikuntanathan and C. Jarzynski. Escorted free energy simulations: Improving convergence by reducing dissipation. *Phys. Rev. Lett.*, 100:190601, 2008.
- [120] S. Vaikuntanathan and C. Jarzynski. Dissipation and lag in irreversible processes. *Europhys. Lett.*, 87:60005, 2009.
- [121] S. Vaikuntanathan and C. Jarzynski. Escorted free energy simulations. *J. Chem. Phys.*, 134:054107, 2011.
- [122] N.G. van Kampen. *Stochastic Processes in Physics and Chemistry*. Elsevier, Amsterdam, 2007.
- [123] E. Vanden-Eijnden and G. Ciccotti. Second-order integrators for langevin equations with holonomic constraints. *Chem. Phys. Lett.*, 429:310–316, 2006.
- [124] L. L. Wang, R. A. Friesner, and B. J. Berne. Replica exchange with solute scaling: A more efficient version of replica exchange with solute tempering (rest2). *J. Phys. Chem. B*, 115:9431–9438, 2011.

- [125] J.D. Weeks, D. Chandler, and H.C. Andersen. Role of repulsive forces in determining the equilibrium structure of simple liquids. *J. Chem. Phys.*, 54:5237–5247, 1971.
- [126] T.W. Whitfield, L. Bu, and J.E. Straub. Generalized parallel sampling. *Physica A*, 305:157–171, 2002.
- [127] B. Widom. Some topics in the theory of fluids. *J. Chem. Phys.*, 39:2808–2812, 1963.
- [128] F.M. Ytreberg and D.M. Zuckerman. Single-ensemble nonequilibrium path-sampling estimates of free energy differences. *J. Chem. Phys.*, 120:10876–10879, 2004.
- [129] D.M. Zuckerman and E. Lyman. A second look at canonical sampling of biomolecules using replica exchange simulation. *J. Chem. Theory. Comput.*, 2:1200–1202, 2006.