Accelerating Ring-Polymer Molecular Dynamics using Hyperdynamics

by

Luke Clyde Adams

A thesis submitted in partial fulfillment of the requirements
for graduation with Honors in Physics.

Whitman College
2018

Certificate of Approval

This is to certify that the accompanying thesis by Luke Clyde Adams has been accepted in partial fulfillment of the requirements for graduation with Honors in Physics.

_____

Frederick G. Moore, PhD

Whitman College
May 9, 2018

Abstract

Molecular Dynamics (MD) is an incredibly powerful tool in the computational study of atomistic systems. However it is also very computationally costly, which severely limits the types of events that can be studied using MD. Additionally, traditional MD methods are purely classical, and thus they do not incorporate quantum zero-point energy and tunneling effects. At low temperatures in particular, this omission of quantum effects produces results that may be incorrect by many orders of magnitude.

In this thesis, we present a new method, hyperRPMD, which both incorporates quantum effects, and also reduces the computational cost of simulating state-to-state dynamics of an atomistic system. As the name suggests, hyperRPMD combines the power of ring-polymer MD to incorporate quantum effects into the simulation, and the power of hyperdynamics to accelerate the state-to-state dynamics of the system. We validate hyperRPMD by simulating a model one-dimensional system.

Acknowledgments

This work was conducted at the Los Alamos National Laboratory during the summers of 2016 and 2017 under the mentorship of Dr. Arthur Voter and Dr. Danny Perez. They have allowed me the independence to learn on my own, while also offering helpful guidance to prevent me from becoming too mired down in the details. I am happy to report that I will be returning to Los Alamos for another year to continue this work.

I am also indebted to the Whitman College Physics and Math Departments, and in particular, Professor Fred Moore for his helpful advice while writing this thesis. By my count, I have taken more classes with Professor Moore than any other professor at Whitman.

Additionally, my girlfriend of three years, Amanda, provided valuable feedback on this thesis.

Lastly, I am grateful for the enduring support of my parents, Greg and Toni, who have always supported me in my various endeavors.

Thank you!

Contents

List of Figures

Chapter 1

Introduction

In recent years, the exponential increase in the availability of computing power has made the direct computational study of atomic systems feasible. These computational studies allow physicists to refine physical models by testing them at scale on large systems, and to make predictions that would not be possible using purely analytical means.

Many methods have been developed to simulate atomic systems; each makes particular assumptions about the dynamics of the system. Some of these assumptions are necessary for the system to be simulated at all. Other assumptions make the simulation more computationally tractable, and thus faster and easier to simulate.

If the assumptions are well-founded, then the quicker methods are valid, and preferable to less efficient methods. However, human intuition consistently fails to predict how atomic systems will behave, and thus seemingly reasonable assumptions about the dynamics of a system are often dramatically wrong. When more computationally expensive methods are used to test these assumptions, researchers often find that atomic systems behave in ways that were not expected. In these cases, further investigation is warranted to determine how the system *actually* behaves.

Ideally, we would like to use more computationally expensive methods to learn about the true dynamics of a system. With this information, we can ensure that our assumptions are valid, and then we can use quicker

methods to accurately simulate the system at longer timescales. With this goal in mind, we will construct a simulation method that makes as few assumptions as possible about the target system.

1.1 Molecular Dynamics

Suppose that we would like to simulate the dynamics of an atomic system with N atoms, each with mass m. Every atom has three spatial degrees of freedom, and thus the state of the system can be specified as a point in a $3N$ dimensional configuration space. We denote this point with the $3N$ dimensional vector $\boldsymbol{x} \in \mathbb{R}^{3N}$. We call the subset of $\mathbb{R}^{3N}$ that is accessible to the system the *configuration space* of the system.

If we are given an initial system configuration $\boldsymbol{x}_i$, we would like to construct a function $\boldsymbol{x} : \mathbb{R} \to \mathbb{R}^{3N}$ such that $\boldsymbol{x}(0) = \boldsymbol{x}_i$, and $\boldsymbol{x}(t)$ specifies the configuration of the system at time t.

Before we can construct such a function, we must first make a few assumptions about the system. Let us assume that the nuclei of the atoms evolve classically over time, and further suppose that total force acting on a particular atom in the system is simply a function of the positions of all of the atoms in the system. Implicitly in the discussion in the previous paragraph, we have assumed that the system is deterministic. That is, we are assuming that, if we we precisely know the configuration of the system at a particular time t_0, then it is possible to compute the configuration of the system at any other time.

Our assumption that the force on an atom is only a function of the configuration of the system means that we can denote the total force on atom k by $\boldsymbol{F}_k : \mathbb{R}^{3N} \to \mathbb{R}^3$. We define the total force on the system to be the total force acting along each coordinate in the configuration space, and we denote this total force by $\boldsymbol{F}(\boldsymbol{x})$. Now, according to our first assumption, the system evolves in time according to the second order differential equation

$$\boldsymbol{F}(\boldsymbol{x}) = m\boldsymbol{a} = m\,\ddot{\boldsymbol{x}}, \tag{1.1}$$

where we have adopted the convention that time derivatives are denoted using a corresponding number of dots.

In principle, we are done. We have written an equation (really, a system of $3N$ equations) that exactly describes the time evolution of our system in terms of the forces between atoms. However, for most interesting systems, there is no analytical solution to this system of differential equations. Thus we must resort to numerical approximations.

In practice, it is more convenient to numerically integrate such a system using Hamiltonian mechanics. It can be shown that Hamiltonian and Newtonian mechanics are equivalent, and thus each formulation can be derived from the other [14]. However, for the molecular systems that we are considering, the Hamiltonian differential equations are only first order in time, and thus they are more numerically tractable.

We consider an adaptation of the derivation of Hamiltonian dynamics presented in [4]. The momentum of atom k is

$$\boldsymbol{p}_k = m\,\dot{\boldsymbol{x}}_k, \tag{1.2}$$

and we denote the momentum of the system using the $3N$ dimensional vector $\boldsymbol{p}$. Notice that we can now fully specify the state of the system at some point in time using the $6N$ dimensional point $(\boldsymbol{x}, \boldsymbol{p})$. We call the subset of $\mathbb{R}^{6N}$ that is accessible to the system the *phase space* of the system.

Let us further suppose that all of the forces acting on the system are conservative, and that they are only a function of the configuration of the system. Then we can find a potential energy function $\mathcal{U}(\boldsymbol{x})$ with the property that

$$\frac{\partial \mathcal{U}(\boldsymbol{x})}{d\boldsymbol{x}} = -\boldsymbol{F}(\boldsymbol{x}). \tag{1.3}$$

Finally, suppose that our system is isolated from the environment. Then the total energy of the system is some constant E. We denote the energy of a particular state of the system by $\mathcal{H}(\boldsymbol{x}, \boldsymbol{p})$ where $\mathcal{H}$ is called the *Hamiltonian* of the system. It clearly follows that

$$\mathcal{H}(\boldsymbol{x}, \boldsymbol{p}) = \frac{\boldsymbol{p}^2}{2m} + \mathcal{U}(\boldsymbol{x}) = E. \tag{1.4}$$

Now we can differentiate $\mathcal{H}$ with respect to $\boldsymbol{p}$ and $\boldsymbol{x}$ to obtain

$$\frac{\partial \mathcal{H}}{\partial \boldsymbol{p}} = \frac{\boldsymbol{p}}{m} = \dot{\boldsymbol{x}}, \qquad \text{and} \qquad \frac{\partial \mathcal{H}}{\partial \boldsymbol{x}} = \frac{\partial \mathcal{U}}{\partial \boldsymbol{x}} = -\boldsymbol{F}. \tag{1.5}$$

Notice that the left equation has specified a relationship between the momentum of the system and the time derivative of the position; this is one of Hamiltonian's two equations. To find the other equation, we consider the total time derivative of the Hamiltonian. We have that

$$\frac{d\mathcal{H}}{dt} = \frac{\partial \mathcal{H}}{\partial \boldsymbol{p}} \cdot \dot{\boldsymbol{p}} + \frac{\partial \mathcal{H}}{\partial \boldsymbol{x}} \cdot \dot{\boldsymbol{x}} + \cancel{\frac{\partial \mathcal{H}}{\partial t}} = 0. \tag{1.6}$$

Note that we have assumed that the energy of the system is constant in time, so the last term disappears. Then substituting the results of Equation 1.5 into Equation 1.6 we have that

$$\dot{\boldsymbol{x}} \cdot \dot{\boldsymbol{p}} - \boldsymbol{F} \cdot \dot{\boldsymbol{x}} = (\dot{\boldsymbol{p}} - \boldsymbol{F}) \cdot \dot{\boldsymbol{x}} = 0 \tag{1.7}$$

Each coordinate in the configuration space is independent, so if the equation is to be identically zero for each coordinate, the quantity in parenthesis must be zero. Thus Hamiltonian's equations of motion for our system are

$$\dot{\boldsymbol{x}} = \frac{\boldsymbol{p}}{m} \qquad \text{and} \qquad \dot{\boldsymbol{p}} = \boldsymbol{F}(\boldsymbol{x}). \tag{1.8}$$

We have replaced the $3N$ second order differential equations that we had in the Newtonian case with $6N$ first order differential equations. Each of these first order equations relates position to the rate of change of momentum, or vice versa. This substitution will make numerically integrating the system easier.

1.2 Integrators

Now that we have expressed the time evolution of our system in a system of differential equations, we need to devise a way to solve these equations. In all but the simplest systems, the equations will not have an analytical solution, so we are forced to use numerical integration.

At its most fundamental level, an integrator accepts the state of a system, and returns the state of the system after a short timestep, Δt. That is, an integrator is a function I with the property that

$$I(\boldsymbol{x}(t_0),\, \boldsymbol{p}(t_0)) = (\boldsymbol{x}(t_0 + \Delta t),\, \boldsymbol{p}(t_0 + \Delta t)). \tag{1.9}$$

The most computationally expensive part of any integrator is evaluating the force vector $\boldsymbol{F}$. Thus, any reasonable integrator must only perform the force computation once per timestep. In this section, we consider a few simple integrators.

Let's suppose that, for the duration of the timestep, the derivatives of position and momentum are essentially constant. Then the state of the system after the timestep is given by

$$\begin{aligned} \boldsymbol{p}(t_0 + \Delta t) &= \boldsymbol{p}(t_0) + \boldsymbol{F}(\boldsymbol{x}(t_0))\, \Delta t \\ \boldsymbol{x}(t_0 + \Delta t) &= \boldsymbol{x}(t_0) + \boldsymbol{p}(t_0)\, \Delta t / m. \end{aligned} \tag{1.10}$$

This scheme is called Forward Euler integration, and it is appealing because the integrator is clearly related to the equations of motion. However, it is rarely used in real applications because the timestep Δt must be prohibitively short for the trajectory to be accurate. We can improve the accuracy of the integrator by splitting the momentum integration across the position integration. The Leapfrog Integrator uses this approach, and is given by

$$\begin{aligned} \boldsymbol{p}\left(t_0 + \tfrac{\Delta t}{2}\right) &= \boldsymbol{p}(t_0) + \boldsymbol{F}(\boldsymbol{x}(t_0))\, \Delta t/2 \\ \boldsymbol{x}(t_0 + \Delta t) &= \boldsymbol{x}(t_0) + \boldsymbol{p}\left(t_0 + \tfrac{\Delta t}{2}\right) \Delta t / m \\ \boldsymbol{p}(t_0 + \Delta t) &= \boldsymbol{p}\left(t_0 + \tfrac{\Delta t}{2}\right) + \boldsymbol{F}(\boldsymbol{x}(t_0 + \Delta t))\, \Delta t/2. \end{aligned} \tag{1.11}$$

It might appear that the Leapfrog Integrator requires two force computations, and would thus it would be computationally unfavorable. But, the configuration of the system does not change between the second momentum calculation of one timestep and the first momentum calculation of the following timestep, and so the same force vector may be reused.

1.2.1 Langevin Integrators

In all of the preceding discussion, we have assumed that the system is isolated from its environment, and thus the total energy of the system is constant. However, this is rarely the case. More often, the system is in thermal contact with its environment, and so the energy of the system may fluctuate in time.

To model this reality, we typically assume that the environment forms a perfect heat bath. That is, we assume that the environment has sufficient thermal mass that it will stay at a fixed temperature T, regardless of what happens in the system. This heat bath fixes the *temperature* of the system, not its energy.

When the energy of the system is fixed, the phase space of the system is quite limited. We call the probability that the system occupies a particular point in this restricted space the *microcanonical ensemble*. However, now that the temperature of the system is fixed, the energy of the system fluctuates around some average energy, with larger deviations from the average energy becoming increasingly less likely. This means that a much larger phase space is (in principle) accessible. We call the probability that the system occupies a point in this enlarged space the *canonical ensemble*.

To simulate the canonical ensemble, we use *stochastic integration*, a type of numerical integration that is capable of simulating random processes. In particular, we will use Langevin dynamics which simulates the temperature regulating effects of the environment as random "kicks" in the momentum differential equation. The equations of motion for Langevin dynamics are

$$\dot{\boldsymbol{x}} = \frac{\boldsymbol{p}}{m} \qquad \text{and} \qquad d\boldsymbol{p} = \boldsymbol{F}(\boldsymbol{x})\,dt - \gamma \boldsymbol{p}\,dt + \boldsymbol{w}\sqrt{2k_B T\gamma m}. \tag{1.12}$$

Notice that the position differential equation is unchanged from the Hamiltonian case. However, the momentum equation has become a stochastic differential equation where γ is a constant that specifies how strongly the environment and the system are coupled, and $\boldsymbol{w}$ is a $3N$ dimensional vector of independent Brownian motions [1].

Numerical integrators for Langevin dynamics are similar to those for Hamil-

tonian dynamics, with additional terms to account for the stochastic nature of the momentum equation. For example, the forward Euler integrator for Langevin dynamics is

$$\begin{aligned} \boldsymbol{p}(t_0 + \Delta t) &= \boldsymbol{p}(t_0) + \boldsymbol{F}(\boldsymbol{x}(t_0))\,\Delta t - \gamma \boldsymbol{p}(t_0)\,\Delta t + \boldsymbol{R}\sqrt{2k_B T \gamma m\,\Delta t} \\ \boldsymbol{x}(t_0 + \Delta t) &= \boldsymbol{x}(t_0) + \boldsymbol{p}(t_0)\,\Delta t/m. \end{aligned} \tag{1.13}$$

where $\boldsymbol{R}$ is a $3N$ dimensional vector of independent random numbers with mean zero and standard deviation one. Similar to the Hamiltonian case, we do not use the forward Euler integrator because it has poor stability when used with a reasonable timestep length [1].

Instead, we will use an integrator called ABOBA [6]. This integrator splits both position and momentum integration into two parts, similar to the Leapfrog Integrator that we considered earlier. The five steps of the integrator are

$$\begin{aligned} \boldsymbol{x}\left(t_0 + \tfrac{\Delta t}{2}\right) &= \boldsymbol{x}(t_0) + \boldsymbol{p}(t_0)\,\Delta t/(2m) \\ \boldsymbol{p}\left(t_0 + \tfrac{\Delta t}{2}\right) &= \boldsymbol{p}(t_0) + \boldsymbol{F}(\boldsymbol{x}(t_0))\,\Delta t/2 \\ \hat{\boldsymbol{p}}\left(t_0 + \tfrac{\Delta t}{2}\right) &= c_1 \cdot \boldsymbol{p}\left(t_0 + \tfrac{\Delta t}{2}\right) + c_3 \cdot \boldsymbol{R} \\ \boldsymbol{p}(t_0 + \Delta t) &= \hat{\boldsymbol{p}}\left(t_0 + \tfrac{\Delta t}{2}\right) + \boldsymbol{F}(\boldsymbol{x}(t_0 + \Delta t))\,\Delta t/2 \\ \boldsymbol{x}(t_0 + \Delta t) &= \boldsymbol{x}\left(t_0 + \tfrac{\Delta t}{2}\right) + \boldsymbol{p}(t_0 + \Delta t)\,\Delta t/(2m). \end{aligned} \tag{1.14}$$

where $\boldsymbol{R}$ is a vector of random numbers as before,

$$c_1 = e^{\gamma \cdot \Delta t}, \qquad \text{and} \qquad c_3 = \sqrt{m k_B T \left(1 - c_1^2\right)}. \tag{1.15}$$

The notation $\hat{\boldsymbol{p}}$ indicates the stochastic step in the integrator. Notice that all of the momentum calculations are consecutive, and thus the force vector must be computed only once per step; namely between the first and second steps.

The development of integrators for MD simulations is an extremely active area of research. For a much more through discussion of the current state of MD integrators, see [11].

1.3 The Timescale Problem

The incredible power of MD comes with one very significant drawback: MD simulations require an enormous amount of computational power. To ensure that the numerical integration does not "blow up," the MD timestep must be a fraction of the fastest vibrational mode of the system. In a typical system, this means that the timestep length is on the order of one femtosecond.

In contrast, the chemical transition events that we are interested in observing often occur with frequencies measured in micro- or even milliseconds. This discrepancy in scales means that, in order to observe a series of chemical transition events, we would need to simulate the system for trillions of timesteps, which remains infeasible even with recent advances in computational power. We call this the *timescale problem*: the events we are most interested in studying using computational physics methods often occur on timescales that are far outside the reach of standard MD methods [18].

Typically, this problem is solved by resorting to more approximate computational methods, such as Kinetic Monte Carlo, that do not have the limitation that every vibration must be accurately simulated [18]. However, in this thesis, we explore another approximation that preserves the core assumptions of MD, while still allowing the simulation of longer timescales than standard MD.

1.4 Classical and Quantum MD

In our derivation of MD, we assumed that the atoms evolved classically on a single potential energy surface. Of course, we know that this is not really the case; systems at the molecular scale are inherently quantum mechanical. When the system is at a high temperature, and if the atoms in the system are relatively massive, the classical approximation is quite good. However, for systems with light atoms and at low temperatures, these quantum effects can be extremely significant [7].

We consider two quantum effects that can influence molecular systems: zero-point energy and tunneling. Zero-point energy refers to the intrinsic energy of a system that remains even when the system is in its ground state, and tunneling occurs when the system "passes through" an energy barrier, even though, according the classical predictions, it does not have the requisite amount of energy.

Later in this thesis, we will encounter a system for which the classical approximation is wrong by several orders of magnitude. In these cases, it is clearly important to reincorporate the quantum effects that we originally assumed were insignificant. Directly computing these effects is not possible in all but the simplest systems, however several excellent approximation schemes have recently been proposed.

1.5 Structure of the Thesis

For the remainder of this thesis, our goal will be to construct a MD simulation method that both incorporates quantum effects and reduces the timescale problem. In Chapter 2, we examine two important MD simulation methods: hyperdynamics and ring-polymer MD. Then in Chapter 3, we combine these two methods, and simulate a simple one-dimensional system. Finally, Chapter 4 reviews work we have done and suggests some future directions for research.

Chapter 2

Methods

In this chapter, we consider adaptations of the standard MD method that can reduce the timescale problem and simulate quantum effects. Hyperdynamics increases the state-to-state transition rate of systems, which allows for longer simulation times than standard MD methods. Ring-polymer MD (RPMD) approximates quantum effects using purely classical MD. We examine the details of each of these methods in turn.

2.1 Accelerated Molecular Dynamics

Accelerated Molecular Dynamics (AMD) is a class of methods that "speed up" MD simulations, which has the effect of diminishing the timescale problem. In the following sections, we will give a more precise definition of what it means to speed up an MD simulation, and also examine several AMD methods, and the compromises each must make to achieve acceleration.

2.1.1 Infrequent Event Systems

Before we explore AMD methods, we need to examine what exactly it is that they accelerate. The computational power of the computer running

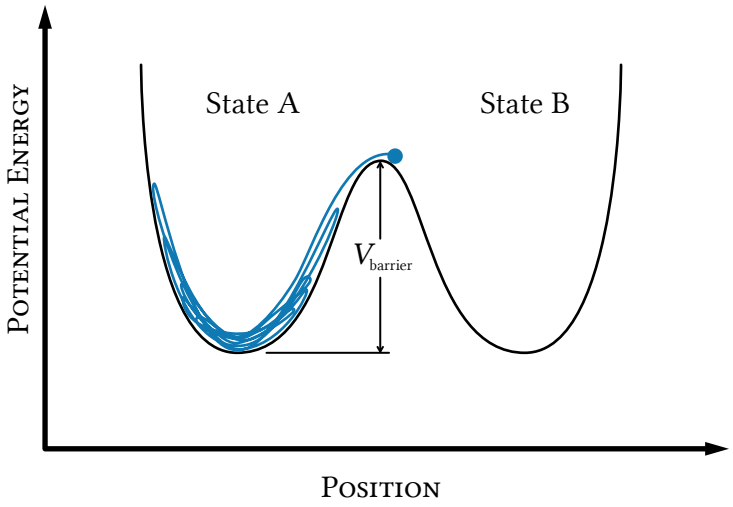

Figure 2.1. A schematic representation of a system undergoing a state transition. The line indicates the trajectory of the particle prior to the transition. Initially, the particle is confined in state A. The particle transitions only when enough energy has been concentrated along the reaction coordinate for the particle to overcome the barrier V_{barrier}.

the simulation is fixed, so clearly we must trade some sort of accuracy in return for increased simulation speed. For this reason, AMD methods cannot be used to simulate all systems; luckily, there are a large number of very interesting systems for which AMD methods are valid. For example, surface diffusion and radiation damage annealing are two important classes of systems that are amenable to AMD methods [16].

Consider a simple system in which a single particle is constrained to move along one axis, and suppose that the system is coupled to a heat bath that fixes the temperature of the system. For each position the particle may take along the axis, there is a corresponding potential energy $V(x)$, as shown in Figure 2.1. This double well potential has two energy basins, A and B, that are separated by a barrier with height V_{barrier}. Suppose that the particle initially begins in state A. Then, provided that the particle's energy is less than the barrier potential, the particle will be confined to vibrate within state A.

However, on occasion the particle's energy will increase sufficiently to make a transition between the energy basins. It might transition into state B, or it might rapidly bounce between states. This state transition on the potential energy surface corresponds to a chemical transition in the system. After a short period of time, the energy that has been concentrated into the particle will dissipate, and the particle will once again be confined to vibrate within a single well. We call this loss of energy *thermalization*.

We describe the time it takes for the particle to gather enough energy to transition from state to state with the rate constant, k_{rxn}. We call the average time that the system spends before transitioning $\mathcal{T}_{\text{rxn}}$ and we denote the average amount of time when the particle has sufficient energy to transition between energy wells by $\mathcal{T}_{\text{corr}}$.

Once the particle has settled into one of the wells, the random thermal energy from the environment will eventually make it impossible to determine how that particle entered the energy basin. We call the time it takes for the particle to "lose its memory" of how it came to be in the state the dephasing time, $\mathcal{T}_{\text{dephase}}$. If the time between transitions is much larger than the dephasing time, then each transition event appears to be independent of the transition before it.

We say that a system is an *infrequent event system* if the average time between transitions is much larger that the dephasing time following a transition. That is, an infrequent event system has $\mathcal{T}_{\text{rxn}} \gg \mathcal{T}_{\text{dephase}}$. This means that consecutive transition events will be not influence each other. If this condition does not hold, consecutive transition events will be *dynamically correlated*. For an infrequent event system, we can describe the probability per unit time that the system will undergo a transition using the rate constant k_{rxn} defined by

$$k_{\text{rxn}} = \frac{1}{\mathcal{T}_{\text{rxn}}}. \tag{2.1}$$

Our goal in all of the AMD methods is to spend less time simulating the vibrating system, and more time observing transition events. Infrequent event systems are particularly amenable to this goal because it is not critical to faithfully represent the vibrational activity between transitions to ensure that the state-to-state transition activity is accurate.

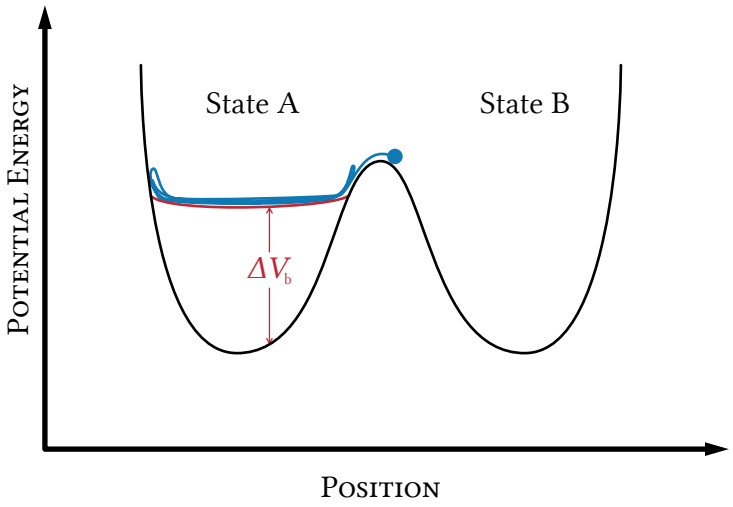

Figure 2.2. A one-dimensional example of a bias potential applied to a potential energy surface. The bias potential speeds the system's transition by reducing the depth of the energy basin. Notice that the bias potential is zero at the dividing surface.

2.1.2 Hyperdynamics

Hyperdynamics, the first of the AMD methods, was developed by Voter in the late 1990's. The method reduces the time spent computing vibrational activity of the system while still capturing the state-to-state dynamics. With the hyperdynamics method, "*at an accelerated pace, the system evolves from state-to-state in a sequence representative of the exact dynamics*" [15].

This method works by adding a bias potential ΔV_b to a system's potential energy. The system is then simulated on this biased potential energy surface, which accelerates the state-to-state dynamics of the system by reducing the activation energy required for the system to make transitions. The actual simulation time, and thus the reaction rate coefficients, are recovered as a statistical property of the simulation called *hypertime*, which is computed for each timestep as

$$\Delta t_{\text{hyper}} = \Delta t \cdot e^{\beta V_b(\boldsymbol{x})}, \tag{2.2}$$

where $\beta = 1/(k_B T)$ with k_B as the Boltzmann constant, and T as the

temperature.

In the limit as the simulation is carried out for an infinite number of timesteps, the hypertime becomes exact [17]. For the hyperdynamics method to be applicable, the system must meet three conditions:

1. The transitions between potential energy basins must be infrequent.
2. The bias potential ΔV_b must be zero at the dividing surfaces between potential energy basins.
3. There must be no recrossings from one basin to another in either the unbiased or the biased system. This means that when the system leaves the current energy basin, it must thermalize in the first energy basin into which it enters.

In practice, the third condition can often be relaxed to require that the biased and unbiased systems exhibit the same recrossing behavior, however the stronger form is assumed in the derivation of hyperdynamics.

Hyperdynamics is based on Transition State Theory (TST), so before we examine the derivation to hyperdynamics, we first review TST.

Transition State Theory

TST is a method of computing approximate reaction rate coefficients for systems in thermal equilibrium with their environment. Consider a potential energy surface that has an energy well A that is separated from the surrounding configuration space by an energy barrier. We can construct a dividing surface that encloses the state A, and then the TST rate for a transition from state A is

$$k^{\mathrm{TST}} = \left[\begin{array}{c} \text{outward flux through} \\ \text{the dividing surface} \end{array} \right] \times \left[\begin{array}{c} \text{probability system is} \\ \text{near dividing surface} \end{array} \right]. \quad (2.3)$$

More formally, consider a prototypical system of N atoms with position $\boldsymbol{x}(t)$ and momentum $\boldsymbol{p}(t)$. Suppose that the system has a potential energy $V(\boldsymbol{x})$, and it has kinetic energy $K(\boldsymbol{p})$ given by

$$K(\boldsymbol{p}) = \frac{\boldsymbol{p}^2}{2m} \quad (2.4)$$

where m is the mass of each atom. The total energy of the system at a particular point in phase-space $E(\boldsymbol{x}, \boldsymbol{p}) = K(\boldsymbol{p}) + V(\boldsymbol{x})$. We define a function $F_A(\boldsymbol{x})$ with the property that

$$F_A(\boldsymbol{x}) = \begin{cases} > 0, & \boldsymbol{x} \text{ is in state } A; \\ 0, & \boldsymbol{x} \text{ is on the boundary of state } A; \\ < 0, & \boldsymbol{x} \text{ is outside state } A. \end{cases} \tag{2.5}$$

The occupation function $\Theta_A(\boldsymbol{x})$ for state A is given by

$$\Theta_A(\boldsymbol{x}) = \theta[\, F_A(\boldsymbol{x}) \,], \tag{2.6}$$

where θ is the standard unit step function. Notice that the dividing surface between A and the rest of the space is the subset of the configuration space that satisfies $F_A(\boldsymbol{x}) = 0$, and thus we define the dividing surface to be

$$\delta_A(\boldsymbol{x}) \equiv \delta[\, \Theta_A(\boldsymbol{x}) \,]. \tag{2.7}$$

We must also determine the velocity v_A which is perpendicular to the boundary surface of state A. This is simply given by

$$v_A = -\frac{\nabla F_A(\boldsymbol{x})}{|\nabla F_A(\boldsymbol{x})|} \cdot \frac{\boldsymbol{p}}{m}. \tag{2.8}$$

Now we have all of the pieces required to define the TST rate. The outward flux through some small patch of the dividing surface is

$$|v_A|\, \delta_A(\boldsymbol{x})\, \Theta_A(\boldsymbol{x})\, d\boldsymbol{x}\, d\boldsymbol{p}. \tag{2.9}$$

The normal velocity accounts for the flux through the patch of dividing surface, the delta function ensures that the point is actually on the dividing surface, and the step function enforces that the system is *leaving* state A. To compute the total rate constant, we simply weight this factor by the probability of the system being in a particular state (the Boltzmann factor), and then integrate over the entire phase space. Thus the TST approximation for the total escape rate is

$$k_{A\rightarrow}^{\mathrm{TST}} = \frac{1}{Z} \iint |v_A|\ \delta_A(\boldsymbol{x})\, \Theta_A(\boldsymbol{x})\, e^{-\beta E(\boldsymbol{x},\boldsymbol{p})}\, d\boldsymbol{x}\, d\boldsymbol{p}, \tag{2.10}$$

The factor Z is the partition function, which is defined as

$$Z \equiv \iint e^{-\beta E(\boldsymbol{x},\boldsymbol{p})}\, d\boldsymbol{x}\, d\boldsymbol{p}. \tag{2.11}$$

Many thermodynamic properties are computed by integrating over the Boltzmann weighted phase space, so we have a more compact notation to indicate this. For example, the average value of a property P over the phase space is

$$\langle P \rangle \equiv \frac{1}{Z} \iint P(\boldsymbol{r},\boldsymbol{p})\, e^{-\beta E(\boldsymbol{x},\boldsymbol{p})}\, d\boldsymbol{x}\, d\boldsymbol{p}. \tag{2.12}$$

With this new notation, we can write the TST rate as

$$k_{A\rightarrow}^{TST} = \langle |v_A|\, \delta_A(\boldsymbol{x})\, \Theta_A(\boldsymbol{x}) \rangle = \langle |v_A|\, \delta_A(\boldsymbol{x}) \rangle_A\,, \tag{2.13}$$

where, in the last equality, we indicate that the ensemble average should only occur over the phase space of state A by using the subscript.

The TST rate approximates the true transition rate of the system, however, if the system satisfies a few conditions, the TST rate is actually exact. First, we assume that the atoms are evolving classically, and that they are only subject to conservative forces. This means that we can think of the motion of the system as a movement on a $3N$ dimensional potential energy surface. Second, we assume that the reactant state is in thermal equilibrium with its surroundings, and thus we may assume a Boltzmann distribution within the reactant state.

Finally, we assume that there are no recrossing events. This means that we are assuming that every trajectory that crosses the dividing surface leaving state A will actually thermalize outside of the state, and not return to state A. This assumption is often the most troublesome because recrossings are present in almost all systems to some degree. Several methods have been developed to relax this assumption such as variational TST [9] and dynamical correction factors [19]. However, these methods are beyond the scope of this thesis.

If a system satisfies these three conditions, then the TST rate will be the exact reaction rate.

Derivation of Hyperdynamics

Now that we have defined the TST rate, we are in a position to examine the derivation of hyperdynamics. In this section we expand on the original derivation given in [15], and fill in some of the omitted details.

Consider the example system from the last section, and suppose we add a non-negative bias potential $\Delta V_b(\boldsymbol{x})$ to the system's potential energy $V(\boldsymbol{x})$. We design the bias potential in such a way that $\Delta V_b(\boldsymbol{x})$ is zero along the dividing surface, does not block any escape paths from state A, and does not introduce any new energy wells into the system.

We define $E_b(\boldsymbol{x}, \boldsymbol{p}) = K(\boldsymbol{p}) + V(\boldsymbol{x}) + \Delta V_b(\boldsymbol{x})$ to be the total biased energy at a point in phase-space. Now modifying our TST rate equation, we have that

$$
\begin{aligned}
k_{A\rightarrow}^{TST} &= \left\langle |v_A|\, \delta_A(\boldsymbol{x}) \right\rangle_A && \text{(2.14a)}\\
&= \left\langle |v_A|\, \delta_A(\boldsymbol{x}) \Theta_A(\boldsymbol{x}) \right\rangle && \text{(2.14b)}\\
&= \frac{\iint |v_A|\, \delta_A(\boldsymbol{x}) \Theta_A(\boldsymbol{x}) \cdot e^{-\beta E(\boldsymbol{x},\boldsymbol{p})}\, d\boldsymbol{x} d\boldsymbol{p}}{\iint e^{-\beta E(\boldsymbol{x},\boldsymbol{p})}\, d\boldsymbol{x} d\boldsymbol{p}} && \text{(2.14c)}\\
&= \frac{\iint |v_A|\, \delta_A(\boldsymbol{x}) \Theta_A(\boldsymbol{x}) \cdot e^{-\beta E(\boldsymbol{x},\boldsymbol{p})} \cdot e^{-\beta(\Delta V_b(\boldsymbol{x}) - \Delta V_b(\boldsymbol{x}))}\, d\boldsymbol{x} d\boldsymbol{p}}{\iint e^{-\beta E(\boldsymbol{x},\boldsymbol{p})} \cdot e^{-\beta(\Delta V_b(\boldsymbol{x}) - \Delta V_b(\boldsymbol{x}))}\, d\boldsymbol{x} d\boldsymbol{p}} && \text{(2.14d)}\\
&= \frac{\iint |v_A|\, \delta_A(\boldsymbol{x}) \Theta_A(\boldsymbol{x}) \cdot e^{-\beta E_b(\boldsymbol{x},\boldsymbol{p})} \cdot e^{+\beta \Delta V_b(\boldsymbol{x})}\, d\boldsymbol{x} d\boldsymbol{p}}{\iint e^{-\beta E_b(\boldsymbol{x},\boldsymbol{p})} \cdot e^{+\beta \Delta V_b(\boldsymbol{x})}\, d\boldsymbol{x} d\boldsymbol{p}} && \text{(2.14e)}\\
&= \frac{\left\langle |v_A|\, \delta_A(\boldsymbol{x}) \cdot e^{\beta \Delta V_b(\boldsymbol{x})} \right\rangle_{A_b}}{\left\langle e^{\beta \Delta V_b(\boldsymbol{x})} \right\rangle_{A_b}}, && \text{(2.14f)}
\end{aligned}
$$

where the subscript A_b indicates that the ensemble average is being taken on the biased potential energy surface within state A. That is, we define the state A_b to be the state A with the bias potential "turned on."

Recall that $\Delta V_b(\boldsymbol{x}) = 0$ whenever $\delta_A(\boldsymbol{x})$ is nonzero, and thus we may simplify the previous equation:

$$
k_{A\rightarrow}^{\mathrm{TST}} = \frac{\left\langle |v_A|\, \delta_A(\boldsymbol{x}) \right\rangle_{A_b}}{\left\langle e^{\beta \Delta V_b(\boldsymbol{x})} \right\rangle_{A_b}}. \tag{2.15}
$$

The numerator in this equation is just the TST escape rate from the state A_b, and the denominator is the ratio of the partition functions for states A and A_b. If we specified a TST dividing surface, we could evaluate this equation to compute the TST rate using any number of sampling methods.

However, this would require a precise knowledge of the dividing surface, and a complete sampling of the state A in order to compute a transition rate. Instead, we would like to create a dynamic method that can advance the system from state-to-state more quickly than a typical MD simulation without requiring prior knowledge of the possible transitions out of the state.

Suppose we have a system for which the TST approximation is actually exact. That is, any crossing of the TST dividing surface corresponds to a true transition into a new state, and the crossing is not dynamically correlated with any past or future crossing events. Further suppose that modifying the potential $V(\boldsymbol{x})$ by adding ΔV_b does not affect either of these conditions.

The bias potential increases the TST escape rate because the biased energy well is not as deep as the normal well. Also, because $\Delta V_b(\boldsymbol{x})$ is zero along the TST dividing surface, the ratio of TST escape rates to each of the states adjacent to A are preserved because the partition functions cancel in a ratio of rates regardless of whether the potential is biased or unbiased. That is

$$\frac{k^{\mathrm{TST}}_{A\to B}}{k^{\mathrm{TST}}_{A\to C}} = \frac{k^{\mathrm{TST}}_{A_b\to B}}{k^{\mathrm{TST}}_{A_b\to C}}. \tag{2.16}$$

Now we consider the meaning of the simulation time on the biased potential energy surface. We would like to find a way to relate the simulation time to the true time between transitions in the simulation. For a given timestep, the simulation time does not have a physical meaning, but, as we will see, it does have a well defined meaning in the limit as the simulation is carried out for an infinite number of timesteps.

Consider an experiment in which we evaluate the averages in Equation 2.15 in state A_b using a single, long MD simulation. Let Δt_{MD} be the length of a timestep in the simulation, n_{tot} be the total number of MD steps taken in

the simulation, and t_i be the time at the ith MD step. We couple the system to a heat bath to guarantee that, over time, the entire canonical phase space of the system is correctly sampled. If the simulation is sufficiently long, we can compute the averages to arbitrarily high accuracy.

Imagine that we place a reflecting barrier along the TST dividing surface, and count the number of times n_{esc} in which the system attempts to escape state A_b. Then, recalling that the average escape time $\mathcal{T}_{\text{esc},A}$ from state A is the inverse of the rate constant, we have that

$$\mathcal{T}_{\text{esc},A} = \frac{1}{k_{A\rightarrow}^{\text{TST}}} = \frac{\left\langle e^{\beta \Delta V_b(\boldsymbol{x})} \right\rangle_{A_b}}{\left\langle |v_A| \, \delta_A(\boldsymbol{x}) \right\rangle_{A_b}} \tag{2.17}$$

$$= \frac{\frac{1}{n_{\text{tot}}} \sum_{i=1}^{n_{\text{tot}}} e^{\beta \Delta V_b(\boldsymbol{x}(t_i))}}{n_{\text{esc}} / (n_{\text{tot}} \cdot \Delta t_{\text{MD}})} \tag{2.18}$$

$$= \frac{1}{n_{\text{esc}}} \sum_{i=1}^{n_{\text{tot}}} \Delta t_{\text{MD}} \, e^{\beta \Delta V_b(\boldsymbol{x}(t_i))}. \tag{2.19}$$

We have used the equivalence between an ensemble average and a time average by evaluating the numerator using the n_{tot} samples from the canonical ensemble, and evaluating the denominator as a time average. We do not actually wish to compute $\mathcal{T}_{\text{esc},A}$; rather, we would like to find a relationship between the timescale of the biased trajectory and the "true" dynamics of the system. By examining Equation 2.19, it seems that a natural definition for the time evolved per MD step (which we will call the *hypertime* of the system) is

$$\Delta t_{\text{hyper},i} = \Delta t_{\text{MD}} \, e^{\beta \Delta V_b(\boldsymbol{x}(t_i))}. \tag{2.20}$$

Thus we estimate the total hypertime that this simulation has accumulated as

$$t_{\text{hyper}} = \sum_{i=1}^{n_{\text{tot}}} \Delta t_{\text{MD}} \, e^{\beta \Delta V_b(\boldsymbol{x}(t_i))}. \tag{2.21}$$

This definition is completely meaningless on short time scales, however, note that Equation 2.19 can be manipulated to read $t_{\text{hyper}} = \mathcal{T}_{\text{esc},A} \cdot n_{\text{esc}}$. Obviously, on short time scales, the average escape time from state A will not be exactly equal to the true average. However, in the long time limit,

the average escape time must converge to the correct result. Consequently, in the long time limit, the hypertime becomes exact.

2.1.3 Other AMD Methods

Though they will not be our main focus, we will briefly describe two other AMD methods: Temperature Accelerated Dynamics (TAD) and Parallel Replica Dynamics (ParRep). Both of these methods share the requirement that the system's activations be infrequent, however they make different assumptions in order to accelerate the state dynamics.

Temperature Accelerated Dynamics

In the TAD method, a physical system confined in energy basin A is simulated at a high temperature T_{high}. The basin A is surrounded by a reflecting barrier, and times at which the system bounces off the barrier are recorded. Using this information, the correct escape path, and corresponding escape time, is computed for a lower temperature T_{low}, and then the system is advanced to this new state where the procedure is repeated [12].

TAD assumes a more approximate form of TST called harmonic TST, and, like hyperdynamics, it has the advantage that the possible transition pathways do not need to be known prior to the simulation.

Parallel Replica Dynamics

As noted previously, the increase in the availability of computing power has enabled larger MD simulations that incorporate many more atoms. Many of the innovations that have driven the sharp increase in computing power are parallel processing techniques that allow moderately powerful computers to collaborate, resulting in extremely powerful clusters.

These computer clusters have allowed MD simulations of larger simulations by assigning each processor a different parts of the system. This means that the forces acting on the system can be computed efficiently

in parallel. However, the increases in the spatial capability of MD have traditionally not come with a corresponding increase in temporal capability because MD simulations are linear in time. The results of one timestep are dependent on the results of the previous timestep, and as a result, multiple processors can only work on the current step. While this does provide some speedups as more computers are added to the cluster, there is a point of diminishing marginal returns when the communication between new processors exceeds the computation benefits they provide.

Parallel Replica Dynamics (ParRep) solves this problem and allows parallel computational resources to be used to accelerate MD simulations. To do this, ParRep duplicates the system n times, and then simulates each copy of the system independently. When one copy of the system escapes from a potential energy basin, the total simulation time from all of the systems is added to the simulation clock, and then the non-escaping copies are destroyed. New copies are made of the escaping copy, and the process repeats.

For more information about the ParRep method, see [10].

2.2 Ring-Polymer Molecular Dynamics

Ring-polymer MD (RPMD) is an MD method that incorporates quantum tunneling and zero point energy effects into classical MD simulations. It is based on the Feynman path integral formulation of quantum mechanics. Unfortunately, there is not a completely rigorous derivation to show that the state-to-state dynamics produced by RPMD are actually exact. Despite this lack of rigor, RPMD has been extremely successful in accurately simulating quantum effects, and attempts to rigorously derive RPMD are ongoing [5].

To perform an RPMD simulation, a system is replicated n times, and the n replicas are arranged in a "ring." The same atom in adjacent replicas are connected by weak harmonic springs, and different atoms in the same replica experience the same forces as in a classical MD simulation. This duplication and arrangement process is depicted in Figure 2.3.

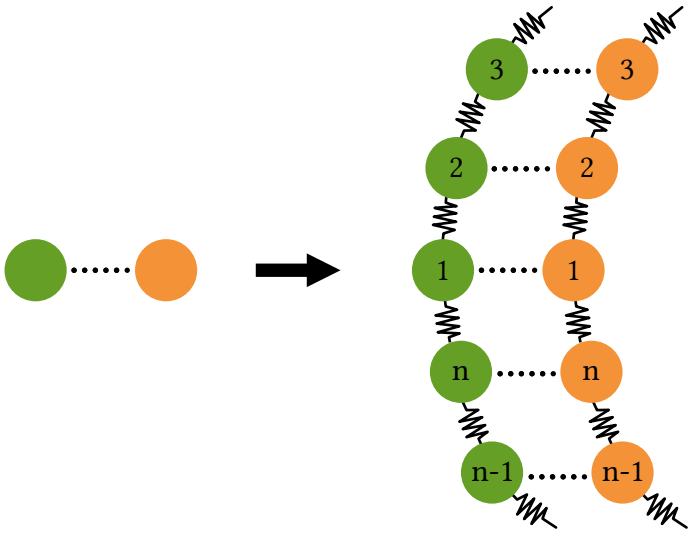

Figure 2.3. A conceptual view of the RPMD process. The system is replicated n times, and the replicas are arranged in a ring. The same atom in adjacent replicas are connected by weak springs (indicated by the jagged lines), and different atoms in the same replica experience the same forces as a classical MD simulation (the dotted lines).

The spring constant k_n between the same atom in adjacent replicas is given by

$$k_n = m \cdot \left(\frac{n k_B T}{\hbar} \right)^2 . \tag{2.22}$$

The springs create closed loops, each with n atoms (the same atom in each of the n replicas). We call these closed loops *ring-polymers*, and sometimes refer the the atoms in the ring-polymer as *beads*. The spatial extent of the ring-polymer is the maximum distance between two beads in the polymer.

The interactions between adjacent systems, or equivalently, the coupling of the ring-polymers, approximate the equilibrium properties (and to a lesser extent, the quantum dynamics) of the system. The approximation of the equilibrium properties is exact in the infinite bead limit. However, the results converge as n increases, so we may simulate the system with an appropriately large n to observe approximately correct dynamics.

2.2.1 Formal Definition

Consider our prototypical $3N$ dimensional system. We would like to use RPMD to compute the quantum canonical partition $\tilde{Z} = \mathrm{Tr}[e^{-\beta\hat{H}}]$, which can, in turn, be used to compute many interesting quantities including reaction and diffusion rates. We can also express the quantum partition function as

$$\tilde{Z} = \lim_{n\to\infty} Z_n = \lim_{n\to\infty} \frac{1}{h^{3nN}} \iint d\boldsymbol{X}\, d\boldsymbol{P} e^{-\beta_n H_n(\boldsymbol{X},\boldsymbol{P})}, \tag{2.23}$$

where $\boldsymbol{X} = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n\}$ and $\boldsymbol{P} = \{\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n\}$ are the position and momentum vectors of n replicas of the system. Together, $\boldsymbol{X}$ and $\boldsymbol{P}$ form the $6N \times n$ dimensional space called the *ring-polymer phase space* [7]. The extended phase space Hamiltonian is given by

$$H_n(\boldsymbol{X},\boldsymbol{P}) = \sum_{i=1}^{n} \left(\frac{\hat{p}_i^2}{2m} + \frac{1}{2} k_n (x_i - x_{i+1})^2 + V(\hat{x}_i) \right), \tag{2.24}$$

where $\beta_n \equiv \beta/n$ and k_n is the spring constant defined in Equation 2.22. We also enforce a cyclic condition on the subscripts such that $x_1 = x_{n+1}$.

The key idea in RPMD is that Equation 2.23 can be accurately approximated for a finite number of replicas, and thus we may compute canonical ensemble averages by simulating a system with the Hamiltonian $H_n(\boldsymbol{X},\boldsymbol{P})$ for an appropriately large n.

For example, suppose that our Hamiltonian has an energy basin A. Then in analogy to section 2.1.2, we define a function $F_A(\boldsymbol{X})$ with the property that

$$F_A(\boldsymbol{X}) = \begin{cases} > 0, & \boldsymbol{X} \text{ is in state } A; \\ 0, & \boldsymbol{X} \text{ is on the boundary of state } A; \\ < 0, & \boldsymbol{X} \text{ is outside state } A. \end{cases} \tag{2.25}$$

Now we let $\Theta_A(\boldsymbol{X}) = \theta[F_A(\boldsymbol{X})]$ and $\delta_A = \delta[\Theta_A(\boldsymbol{X})]$. Then the RPMD-TST rate from state A, which approximates the true quantum TST rate form state A, is

$$k^{\text{RPMD-TST}} \approx \frac{1}{Z_n} \cdot \frac{1}{h^{3Nn}} \iint d\boldsymbol{X}\, d\boldsymbol{P}\, \bar{v}_A\, \delta_A(\boldsymbol{X})\, \Theta_A(\boldsymbol{X})\, e^{-\beta_n H_n(\boldsymbol{X},\boldsymbol{P})}, \tag{2.26}$$

where $\bar{v}_A$ is the normal velocity of the center of mass of the ring-polymer relative to the dividing surface,

$$\bar{v}_A = -\frac{1}{n} \frac{\nabla F_A(\boldsymbol{X})}{|\nabla F_A(\boldsymbol{X})|} \cdot \frac{\boldsymbol{P}}{m}. \tag{2.27}$$

In the limit as n goes to infinity, we find that the quantum TST (qTST) rate is exactly the RPMD-TST rate [2]. That is

$$\lim_{n \to \infty} k^{\text{RPMD-TST}} = k^{\text{qTST}}. \tag{2.28}$$

High Temperature Limit

At very high temperatures, we expect that the quantum zero-point and tunneling effects will be insignificant, and thus the RPMD-TST rate should collapse to the TST rate in the high temperature limit. To see that this is the case, note that the springs between adjacent beads on the ring-polymer become very tight, and thus the spatial extent of the ring-polymer collapses to zero at very high temperatures. In this limit, the RPMD-TST rate can be evaluated with just a single bead, so $n = 1$. With this simplification, Equation 2.26 simply collapses to the TST rate [3], as we expected.

2.2.2 Extended Phase Space

The RPMD method is appealing because, fundamentally, it is just classical molecular dynamics carried out in an extended phase space, namely the ring-polymer phase space. This means that a variety of existing atomistic simulation methods, and the codes that have been written to implement these methods, can be easily adapted to carry out RPMD simulations.

However, when a system is duplicated n times, the number of dimensions in the simulation also increases n fold. Even more troubling, the computational cost of the simulation increases superlinearly with increasing dimension, and thus the cost of an n replica RPMD simulation is even more than n times the cost of a corresponding classical MD simulation.

In this case, the need for AMD methods becomes especially apparent. In the following chapter, we demonstrate a new method, hyperRPMD, that uses hyperdynamics to accelerate the state-to-state dynamics of RPMD simulations.

Chapter 3

Implementation and Results

The goal of this project is to use hyperdynamics to accelerate the state-to-state transitions of RPMD. As was previously noted, the use of AMD methods to accelerate RPMD is particularly nice because simulations in the RPMD extended state space are extremely computationally expensive, even for relatively small systems. In this chapter, we describe a new method, hyperRPMD, that combines hyperdynamics and RPMD. We also test the effectiveness of this new method on a simple, one-dimensional model system.

3.1 Implementation

The combined implementation of hyperdynamics and RPMD is largely straightforward because hyperdynamics is unaffected by the extended phase space the RPMD creates. However, there are a few places where the best choice is not obvious. In this section, we highlight the non-obvious design choices that we made while implementing hyperRPMD.

Recall that the RPMD approximation to the quantum canonical partition function is

$$Z_n = \frac{1}{h^{3nN}} \iint d\boldsymbol{X}\, d\boldsymbol{P} e^{-\beta_n H_n(\boldsymbol{X},\boldsymbol{P})}. \tag{3.1}$$

Notice, in particular, the factor of $\beta_n H_n$. There are two ways to interpret this term: either $(\beta/n)H_n$ or $\beta_n(H_n/n)$. In the first interpretation, the temperature of the extended phase space system is multiplied by a factor of n, while in the second interpretation, the Hamiltonian (and thus the effective mass m and potential energy $V(\boldsymbol{x}_i)$) are divided by a factor of n. The second interpretation is perhaps more intuitive because, for example, the sum of the mass of each bead on a ring-polymer is equal to the classical mass of the atom that the ring-polymer represents. However, this choice quickly becomes confusing when the bias potential is added to the Hamiltonian. For this reason, we instead chose to multiply the simulation temperature by a factor of n.

Another potentially confusing point is the construction of the dividing surface between states. We choose to define the dividing surface between states as the location where the *center of mass* of the ring-polymer crosses the dividing surface in the classical phase space.

Finally, the construction of bias potentials that both satisfy the hyperdynamics assumptions and also provide a reasonable boost factor is always challenging, even without the added complexity of RPMD. Suppose that a bias potential $V_b(\boldsymbol{x})$ has already been designed for the classical system. We then compute the bias force on the extended phase space system as $V_b(\bar{\boldsymbol{x}})$, where $\bar{\boldsymbol{x}}$ is the center of mass of the ring-polymer. We apply the full bias potential to each replica of the system.

This definition of the bias potential, along with our definition of the dividing surface, ensures that the bias potential will be zero at the dividing surface, as required. It is also advantagous because this method is not dependent on the particular form of the classical bias potential; any bias that works in the classical case will also work for hyperRPMD.

3.2 Testing

Recall that the key advantage of the hyperRPMD method is that it can replicate state-to-state dynamics of a molecular system, while also incorporating quantum effects. To test these claims, we simulate a simple one

dimensional system, and use the simulation to compute the reaction rate coefficients between states.

Computing these rates is *not* the primary goal of this method; rather, this method provides a way to explore the quantum dynamics of a system without a complete knowledge of the transition pathways. If the transitions are already known, then there are much more efficient methods to compute reaction rate coefficients. However, computing transition rates provides a convienent way to compare the results of hyperRPMD against known theoretical results.

3.2.1 Eckart Barrier

To test hyperRPMD, we will use the one dimensional symmetric Eckart barrier that was designed to simulate the collinear $H + H_2 \rightarrow H_2 + H$ reaction. The Eckart barrier is defined as

$$V(x) = \frac{V_0}{\cosh^2(x/a)} \tag{3.2}$$

where $V_0 = 0.425\,\text{eV}$, $a = 0.3884\,\text{Å}$, and $m = 9.7 \times 10^{-28}\,\text{kg}$. The rate of this reaction is highly dependent of quantum effects, particularly at low temperatures, and thus it has become somewhat of a standard test case for quantum molecular methods [7, 2, 8]. The classical transition rate for the Eckart barrier is given by

$$k_{\text{cl}} = e^{-\beta V_0}\sqrt{\frac{1}{2\pi\beta m}}, \tag{3.3}$$

and the exact quantum transition rate can be computed by numerically evaluating the integral

$$k_{\text{q}} = \frac{2\pi}{\hbar \cdot Z_r}\int_0^E \mathcal{T}(E)e^{-\beta E}\,dE, \tag{3.4}$$

where $Z_r = \sqrt{m/(2\pi\beta\hbar)}$, is the reactant partition function, $\hbar$ is the reduced Planck constant, and $\mathcal{T}(E)$ is the transmission coefficient at energy E, given

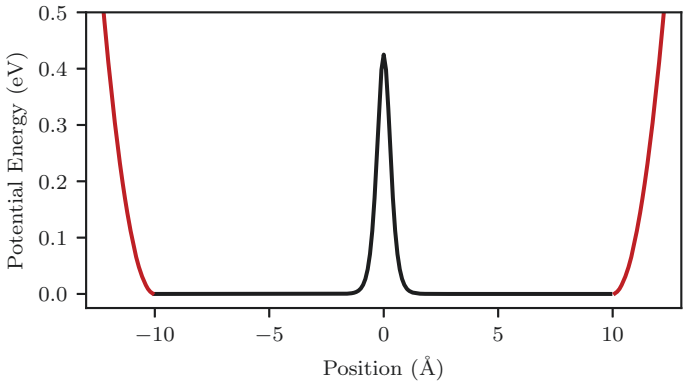

Figure 3.1. An Eckart barrier with reflecting walls at $L = 10$ to create energy wells.

by

$$\mathcal{T}(E) = \frac{\sinh^2(\pi a \cdot A)}{\cosh^2(\pi a \cdot B) + \sinh^2(\pi a \cdot A)} \tag{3.5}$$

with $A = \hbar^{-1}\sqrt{2mE}$ and $B = \hbar^{-1}\sqrt{2mV_0 - \hbar^2/(4a^2)}$ [7].

To make the Eckart barrier suitable for simulation using hyperRPMD, we bounded the barrier by adding two harmonic reflecting walls, one on each side of the barrier at a distance L from the origin. With this modification, we can compute the transition rate as $k = L/\langle\mathcal{T}\rangle$ where $\langle\mathcal{T}\rangle$ is the average transition time. In the limit as L goes to infinity, this rate is exact [7], though experience has shown that the rate is extremely accurate for values as small as $L = 5$ Å. The Eckart barrier with reflecting walls at $L = 10$ Å is shown in Figure 3.1.

3.2.2 Results

We simulated the bounded Eckart system using MD, hyperdynamics, and hyperRPMD and then computed the transition rates at a variety of temperatures. In all cases, we used the ABOBA integrator with a Langevin thermostat to ensure the the system sampled the entire canonical ensemble while it was confined in a state. For the hyperdynamics and hyperRPMD simulations, we used a so called "pond-ice" bias potential which creates

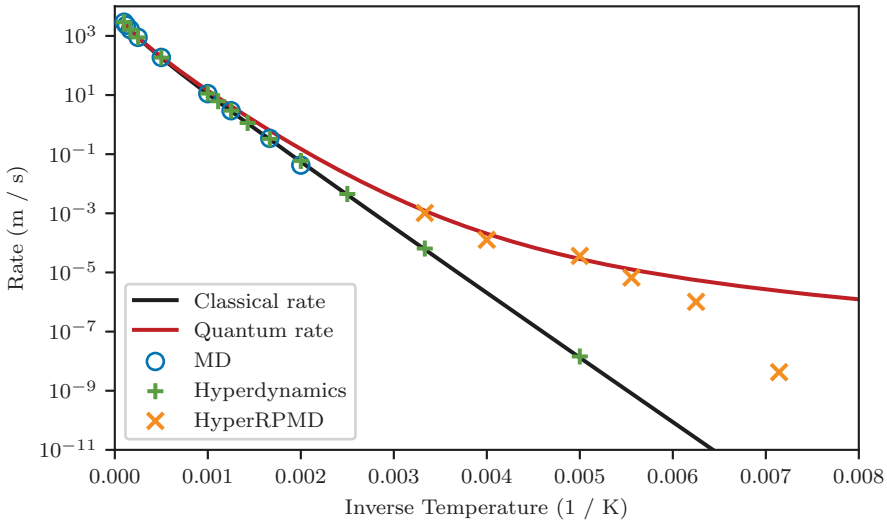

Figure 3.2. Arrhenius plot for the Eckart barrier. The solid lines indicate exact classical and quantum transition rates for the Eckart barrier. Notice that we were able to simulate much lower temperatures using the AMD methods (hyperdynamics and hyperRPMD).

a flat surface on the biased potential surface for any location where the original potential is below a cutoff energy [13]. In real atomic systems, it is difficult to use a pond-ice potential because it requires a detailed knowledge of the lowest barrier for escape from the state. Additionally, the pond-ice bias does not greately accelerate the dynamics in larger systems.

The results of the simulations are shown in Figure 3.2. At the lower temperatures, the transition rates computed using RPMD clearly deviate from the exact quantum rate. This deviation is consistent with other Eckart transition rates computed using RPMD [7, 8]. However, some preliminary testing has indicated that this error may not be inherent in the RPMD method, but rather from an over-aggressive bias potential that prevents the system from fully thermalizing prior to the next transition. In the future, we will preform low temperature simulations using less aggressive bias potentials to determine if this is the case. We do not have any speculation as to what would have caused the errors in other RPMD based transition rate computations, none of which used bias potential based methods.

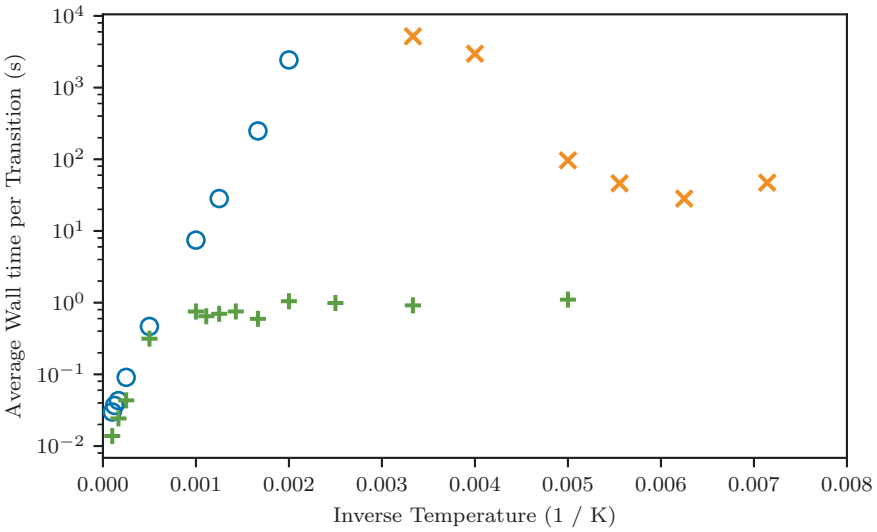

Figure 3.3. Wall time per transition for the Eckart barrier. They symbols used in this graph are the same as in Figure 3.2. Notice that the hyper based methods both achieve a constant time per transition at lower temperatures.

For the lowest temperature (140 K), we used 273 replicas in our RPMD simulation. This means that, compared to a classical simulation, the RPMD simulation takes more than 273 times the computational power to simulation a transition. However, the addition of hyperdynamics allows us to simulate the system at these extremely low temperatures. To see how remarkably effective hyperdynamics is at accelerating the rate of transitions, we plot the inverse temperature against the average wall time per transition in Figure 3.3. This indicates how long (in the real world) we would expect to simulate the system before observing a transition.

Notice that, as the temperature decreases, the expected wall time per transition for a standard MD simulation rapidly increases because the classical transition rate falls rapidly as temperature decreases. In comparison, the hyperdynamics and hyperRPMD methods fair much better, with both eventually approaching an apparently constant wall time per transition.

Another possible way to interpret the speedup is to examine the ratio of the hypertime to the simulation time. We call this ratio the *boost factor*, and it indicates the relative speeds of the hyperdynamics simulation and a

standard MD simulation. For the 140 K hyperRPMD simulation, the boost factor was 6.78×10^{12}, which means that a RPMD simulation would need to run on the order of 10^{12} times longer than the hyperRPMD simulation to reach the same total time.

One particularly enticing aspect of the hyperRPMD method is that the wall time per transition may actually decrease again at extremely low temperatures because the exact quantum transition rate does not approach zero as the temperature goes to zero. This tendency can be seen in the orange points in Figure 3.3. Once the bias potential for low temperature simulations has been corrected, we intend to further investigate this possibility.

Chapter 4

Conclusion and Future Work

The hyperRPMD method uses hyperdynamics to accelerate the state-to-state dynamics of RPMD simulations. This new method allows for rapid simulation of the state-to-state dynamics of infrequent event systems while also incorporating quantum zero-point and tunneling effects.

Traditionally, the key challenge in performing hyperdynamics simulations has been the construction of appropriate bias potentials that both satisfy the hyperdynamics assumptions and also produce large boost factors. As a result, the design of appropriate bias potentials has been an active area of research for many years, and several generally applicable bias potentials have been developed. Fortunately, all of this research is also applicable to hyperRPMD because, by construction, every bias potential that is appropriate for hyperdynamics may also be used for hyperRPMD simulations.

To test hyperRPMD, we simulated the one-dimensional Eckart barrier using MD, hyperdynamics and hyperRPMD at a variety of temperatures. For each temperature and simulation method, we computed transition rates, and compared these rates to classical and quantum exact transition rates. At low temperatures, hyperRPMD consistently computed transition rates that were less than exact quantum transition rate. These results are consistent with other simulation methods based on RPMD. In our particular tests, we believe that this incorrect rate is due to an overly aggressive bias potential that confines the system to vibrate near the dividing surface. This

is particularly interesting because the other RPMD simulation methods that similarly deviate from the exact quantum rate do not utilize bias potentials.

In the future, we will reexamine the low temperature behavior of the hyperRPMD simulations on the Eckart barrier. In particular, we would like to determine if, at low temperatures, the bias potential is constraining the system from sampling the full reactant state prior to transitioning. Once this behavior has been studied, we would also like to examine the intriguing possibility that frequency (measured in wall time) of transition events will actually increase at low temperatures.

We will also implement hyperRPMD for a larger and more complex system to validate the method on higher dimension systems. Two key features of the hyperRPMD method will ease this implementation. First, we can use a bias potential designed for hyperdynamics to carry out the hyperRPMD simulation. Second, hyperRPMD is just classical MD in an extended phase space, so any computer package designed for MD simulations can be easily adapted to carry out hyperRPMD simulations.

Bibliography

[1] Nawaf Bou-Rabee. "Time Integrators for Molecular Dynamics". In: *Entropy* 16 (2013), pp. 138–162. DOI: 10.3390/e16010138.

[2] Ian R. Craig and David E. Manolopoulos. "Chemical reaction rates from ring polymer molecular dynamics". In: *The Journal of Chemical Physics* 122.8 (2005), p. 084106. DOI: 10.1063/1.1850093.

[3] Scott Habershon et al. "Ring-Polymer Molecular Dynamics: Quantum Effects in Chemical Dynamics from Classical Trajectories in an Extended Phase Space". In: *Annual Review of Physical Chemistry* 64 (2013), pp. 387–413. DOI: 10.1146/annurev-physchem-040412-110122.

[4] J.M. Haile. *Molecular Dynamics Simulation. Elementary Methods*. John Wiley & Sons, 1997. ISBN: 0-471-81966-2.

[5] Timothy J. H. Hele and Stuart C. Althorpe. "An alternative derivation of ring-polymer molecular dynamics transition-state theory". In: *The Journal of Chemical Physics* 144.17 (2016), p. 174107. DOI: 10.1063/1.4947589.

[6] Benedict Leimkuhler and Charles Matthews. "Rational construction of stocastic numerical methods for molecular sampling". In: *Applied Mathematics Research eXpress* 144 (2013). DOI: 10.1093/amrx/abs010.

[7] Chun-Yaung Lu, Danny Perez, and Arthur F. Voter. "Accelerating ring-polymer molecular dynamics with parallel replica dynamics". In: *The Journal of Chemical Physics* 144.244109 (2016). DOI: 10.1063/1.4954311.

[8] William H. Miller et al. "Quantum instanton approximation for thermal rate constants of chemical reactions". In: *The Journal of Chemical Physics* 119.3 (2003), pp. 1329–1342. DOI: 10.1063/1.1580110.

[9] Roar A. Olsen. *An introduction to Transition State Theory*. Leiden University. URL: http://www.acmm.nl/molsim/han/2006/TSTAndQTSTAndTSLectures.pdf.

[10] Danny Perez, Blas P. Uberuaga, and Arthur F. Voter. "The parallel replica dynamics method a Coming of age". In: *Computational Materials Science* 100 (2015). Special Issue on Advanced Simulation Methods, pp. 90–103. ISSN: 0927-0256. DOI: https://doi.org/10.1016/j.commatsci.2014.12.011.

[11] Ulf D. Schiller. *An overview of integration schemes for molecular dynamics simulations*. 2008. URL: http://www2.mpip-mainz.mpg.de/~andrienk/journal_club/integrators.pdf.

[12] Mads R. Sørensen and Arthur F. Voter. "Temperature-accelerated dynamics for simulation of infrequent events". In: *The Journal of Chemical Physics* 112 (2000), pp. 9599–9606. DOI: 10.1063/1.481576.

[13] M. M. Steiner, P.-A. Genilloud, and J. W. Wilkins. "Simple bias potential for boosting molecular dynamics with the hyperdynamics scheme". In: *Phys. Rev. B* 57 (1998). DOI: 10.1103/PhysRevB.57.10236. URL: https://link.aps.org/doi/10.1103/PhysRevB.57.10236.

[14] John R. Taylor. *Classical Mechanics*. Mill Valley, California: University Science Books, 2005. ISBN: 1-891389-22-X.

[15] Arthur F. Voter. "A method for accelerating the molecular dynamics simulation of infequent events". In: *The Journal of Chemical Physics* 106.11 (1997), pp. 4665–4677. DOI: 10.1063/1.473503.

[16] Arthur F. Voter. *Accelerated Molecular Dynamics Methods. A very brief introduction*. Los Alamos National Lab. URL: http://public.lanl.gov/afv/AMD-intro.pdf.

[17] Arthur F. Voter. "Hyperdynamics: Accelerated Molecular Dynamics of Infrequent Events". In: *Physical Review Letters* 78.20 (1997), pp. 3908–3911.

[18] Arthur F. Voter. "Introduction to the Kinetic Monte Carlo Method". In: *Radiation Effects in Solids*. Ed. by Kurt E. Sickafus, Eugene A. Kotomin, and Blas P. Uberuaga. Springer Netherlands, 2007, pp. 1–23. ISBN: 978-1-4020-5295-8.

[19] Arthur F. Voter and Jimmie D. Doll. "Dynamical corrections to transition state theory for multistate systems: Surface self-diffusion in

the rare-event regime". In: *The Journal of Chemical Physics* 82.1 (1985), pp. 80–92.