Title
Fluctuations in Supercooled Fluids and Ionic Solutions

Author
Thorpe, Dayton Gray

Publication Date
2017

Fluctuations in Supercooled Fluids and Ionic Solutions

by

Dayton Gray Thorpe

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Phillip Geissler, Co-chair
Professor Ashvin Vishwanath, Co-chair
Professor Michael DeWeese
Professor Ronald Cohen

Spring 2017

Fluctuations in Supercooled Fluids and Ionic Solutions

Abstract

Fluctuations in Supercooled Fluids and Ionic Solutions

by

Dayton Gray Thorpe

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Phillip Geissler, Co-chair

Professor Ashvin Vishwanath, Co-chair

An overview of five studies is presented in two parts. The first part presents two studies of supercooled fluids. The second part presents three studies of water and aqueous solutions. Each study seeks a minimal model of a condensed matter system. In the first study, kinetically constrained models (KCM's) are compared to alternative theories of the glass transition in high dimensions. Dimensionality is used as a parameter to tune the connectivity of a lattice, where a higher dimensional model has more interactions between neighboring sites. This study finds that KCM's outperform alternative theories in high dimensions. The second study explores the possibility that bacteria have evolved to exploit the glass transition to enter a dormant state when environmental conditions are unfavorable. Although the available evidence shows that the bacterial cytoplasm does not meet the strict definition of a fragile glass former, much of its behavior is similar to and can be described using close analogies with the glass transition. In the second part, the third study describes the molecular mechanisms that gives rise to large electric field fluctuations, which in turn cause autoionization and ion dissociation. The fourth study analyzes several candidate order parameters as the basis for a Gaussian field theory of ion solvation. Finally, the fifth study discusses the most popular current explanation for observed charge asymmetry at liquid-vapor interfaces. This explanation, based on linear response of the surface polarization to the presence of an ion, is incorrect. Instead, the surface polarization responds non-linearly to the presence of an ion. Incorporating these non-linear fluctuations is essential to predict solvation free energies.

To my parents, Gail and Devin

Contents

Contents ii

List of Figures iv

List of Tables viii

1 Introduction 1

I Supercooled Liquids 4

2 The East Model Correctly Predicts the Behavior of Simulated Glassy Systems in High Dimensions 5

2.1 Asymptotic Result for High Dimensional Curvature 6
2.2 Hierarchical Dynamics Appears at Lower Temperatures in Higher Dimensions 9
2.3 Diffusion Decoupled from Stokes-Einstein In Any d 11
2.4 Conclusion 13

3 Glass-like Bacterial Cytoplasm Is Fluidized by Driving in Simulations 15

3.1 Bacterial Cytoplasm: Experiments 16
3.2 Bacterial Cytoplasm: Glassy Simulations 19
3.3 Bacterial Cytoplasm: Driven Simulations 22
3.4 Conclusion 28

II Water 29

4 Electric Field Fluctuations in Water 30

4.1 Autoionization 30
4.2 Ion Dissociation 33
4.3 Long-Range Contributions to Electric Field Fluctuations 34
4.4 Umbrella Sampling 35
4.5 Umbrella Sampling of the Electric Field on a Proton 36

4.6 Results 38
4.7 Conclusion 49

5 Hydrogen Bonds and Ions 50
5.1 Tetrahedral Order Parameter 52
5.2 Fluctuations of the Hydrogen Bond Density Field 53
5.3 Structure of Hydrogen Bond Network Around an Ion 59
5.4 Conclusion 61

6 The Surface Potential of Water Varies with the Presence of an Ion 63
6.1 Predictions of Dielectric Continuum Theory with an Interface 63
6.2 How to Measure Surface Interactions 66
6.3 Conclusion 70

Bibliography 71

List of Figures

2.1 The height function, h, measured numerically for the East model in dimensions 1, 2, and 3 and plotted against the asymptotic result from equation 5. 7
2.2 The curvature, J^2 in equation 3, observed in equilibrium simulations has an approximately constant offset from the asymptotic result $J^2 = 1/2d\log(2)$. . . . 8
2.3 The onset temperature as defined in equation 12. 10
2.4 We show diffusion constants for the East model across eight decades in relaxation time. Dashed lines show linear best fits calculated over $\log_{10}(\tau) \in (2, 8)$. 11
2.5 The deviation of the slope of $\log_{10}(D)$ versus $\log_{10}(\tau)$ from that predicted by the Stokes-Einstein relation in dimensions one through 10 and 13. 12

3.1 The three interaction potentials between the two particle types I simulated. . . . 21
3.2 Inverse diffusion constants for the three particle sizes in a confined cell with total dimension 15σ x 100σ. There may be a slightly larger slow down for the large probe particle, but the difference between the small and large particles is nowhere near the factor of five seen in experiments. 22
3.3 Snapshot of a 2D simulation of a bacterial cell. The smallest particles around the corners of the simulation never move; their purpose is to provide an approximately cell-shaped, bumpy confining boundary. 23
3.4 Three views of a snapshot of a 3D simulation of a bacterial cell: side view of the entire system (top); the bumpy, rounded cap made of small unmoving particles (bottom-left); a cross-section down the long axis of the cell (bottom-right). . . . 25
3.5 Mean square deviation of three particle sizes in the two-dimensional simulations. Adding an external force just to the smallest particles in the system speeds up the diffusion of other particles. 26
3.6 Ratio of diffusion constants in the presence or absence of an external force on the smallest particles for all other particles in the system. 26
3.7 Mean square deviation of two particle sizes in a three-dimensional simulation. The MSD is measured separately along each axis. The driving is along the x-axis, which is also the long axis. Diffusion along the other axes closely resembles that in a simulation with no driving. 27

4.1 The distribution of commitor values for configurations that are equilibrated from a transition state with all atoms through the first (a), second (b), and third (c) solvation shells frozen, reproduced from [31]. Configurations that are still in a transition state would have a commitor distribution sharply peaked around 0.5. Panel (c) suggests that freezing three solvation shells may hold the configuration in a transition state, but when the process is repeated with a larger simulation cell the distribution widens substantially (inset). 32
4.2 Reproduction of Figure 2 from [29]. The horizontal axis shows progression along the linear interpolation of protons between the charge-separated and neutral states. The vertical axis shows the potential energy difference from the neutral state. The top panel is calculated from a neutral configuration, the middle panel is a transition state, and the bottom panel is from a charge-separated configuration. 33
4.3 The distribution of electric fields experienced by the hydrogen atom of a water molecule reproduced from [32]. The x direction is parallel to the OH bond. The y direction is perpendicular to the OH bond and in the HOH plane. The z direction is perpendicular to the HOH plane. 35
4.4 Contributions to the electric field on a central atom as a function of distance reproduced from [32]. Different colored lines are for different sizes of electric fields on the central atom. 35
4.5 The very tight correlation between the total electric field on an atom and the electric field from nearby atoms only, excluding the long-ranged contribution from an Ewald sum, allows for importance sampling indirectly on the former by explicitly biasing the latter. 37
4.6 The free energy of the electric field experienced by a proton on a water molecule, projected along the OH bond axis. The black line is calculated from umbrella sampling windows that start at equilibrium and are then equilibrated to different biases. The red line is calculated from umbrella sampling windows that start from the last configuration in the trajectory with the largest, negative bias. . . . 39
4.7 The free energy of the electric field experienced by a proton on a water molecule, projected along the OH bond axis, under a bias. The bias is a harmonic penalty centered at the most negative bin observed, approximately -15 V/Å, with a spring constant chosen to force the free energy of that bin to 0. 40
4.8 Representative configurations taken from trajectories biased towards large negative (left) and positive (right) values of the electric field experienced by a tagged proton in water, projected along the OH-bond axis. 41
4.9 Radial distribution function (RDF) of hydrogen atoms around tagged hydrogen atom in trajectories with a harmonic bias around electric fields of -18 V/Å. This RDF quantitatively describes the configuration seen in the left panel of Figure **??**. 42
4.10 Radial distribution function (RDF) of oxygen atoms around the oxgen atom covalently bonded to the tagged hydrogen atom in trajectories with a harmonic bias around electric fields of -18 V/Å. This RDF quantitatively describes the configuration seen in the left panel of Figure **??**. 43

4.11 Radial distribution function (RDF) of oxygen atoms around the oxgen atom covalently bonded to the tagged hydrogen atom in trajectories with a harmonic bias around electric fields of 25 V/Å. This RDF quantitatively describes the configuration seen in the right panel of Figure **??**. 44

4.12 Representative configuration from a trajectory with a bias on the electric field on the blue hydrogen atom towards 25 V/Å. In this configuration, the water molecule with the hydrogen atom being pulled away strongly is accepting four hydrogen bonds, a large fluctuation from the equilibrium average of 1.8. 45

4.13 The potential of mean force for electric field fluctuations on two hydrogen atoms whose water molecules are within each other's first solvation shell. 45

4.14 The potential of mean force for electric field fluctuations on two hydrogen atoms whose water molecules are within each other's second solvation shell. When a central atom and an atom in the second solvation shell are biased towards large positive fields, a third atom forms a bridge between the two tagged atoms and remains there on the order of 1 ns, giving a path for a proton to diffuse 3 hydrogen bonds away from its initial position by the Grotthuss mechanism. 46

4.15 Representative configurations from trajectories with two molecules in each other's first shell biased to have large electric fields on one of their hydrogen atoms. Hydrogen atoms in blue are having their covalent OH bond stretched. Hydrogen atoms in orange are having their covalent OH bond compressed. 46

4.16 This plot shows what the 2D PMF of the electric field of two tagged hydrogen atoms in each other's first solvation shell would look like if their fluctuations were uncorrelated. 47

4.17 The difference between the actual 2D PMF of the electric field of two tagged hydrogens in each other's first solvation shell and the result that would have appeared for uncorrelated two hydrogen atoms individually following the 1D distribution. 47

4.18 The difference between the actual 2D PMF of the electric field of two tagged hydrogens in each other's second solvation shell and the result that would have appeared for uncorrelated two hydrogen atoms individually following the 1D distribution. 48

5.1 The tetrahedral order parameter does not fit a Gaussian. Although it looks bimodal, it also does not fit the sum of two Gaussians. These data come from a simulation of bulk SPC/E water at 298 K. 52

5.2 The distribution of the number of hydrogen bonds to a single water molecule. This distribution is found in a simulation of SPC/E water where a hydrogen bond is defined as in the main text. The average number of water molecules per molecule is 3.60 with a non-Gaussian distribution. 54

5.3 The distribution of the number of hydrogen bonds in a cell with sidelength $L =$ 3.104 Å, which on average holds 1 water molecule. 55

5.4 Fluctuations in the number of water molecules in a spherical probe volume (top) or a cubic probe volume (bottom). Lines are Gaussians with the same mean and variance as the data. Both cubic and spherical probe volume distributions are Gaussian to a high degree even in small volumes that hold an average of less than one molecule. 56
5.5 Donated and accepted hydrogen bond density fluctuations in three sizes of cubic probe volumes. Lines are Gaussians with the same mean and variance as the data. 57
5.6 Distribution of the normalized Fourier modes of the hydrogen bond density field. Each individual sample of $\hat{\rho}_{\mathrm{HB}}(k_n)$ has been divided by $\sqrt{\mathrm{Var}(\hat{\rho}_{\mathrm{HB}}(k_n))}$. The solid line is a Gaussian of mean 0 and variance 1. 57
5.7 Radial distribution functions of hydrogen bonds located at the midpoint between the Oxygen centers of the donating and accepting molecule.. 58
5.8 Radial distribution functions of hydrogen bonds located at the midpoint between the Oxygen centers of the donating and accepting molecule around a central ion. 58
5.9 Statistics of hydrogen bonds between water molecules. Each bond is tagged with a unit vector located at the O-O midpoint pointing from the donating hydrogen atom to the accepting oxygen atom. 59
5.10 Statistics of hydrogen bonds tagged with unit vectors. The vector is located at the ion-O midpoint. For water molecules, the vector is located at the O-O midpoint. The vector points from the donating molecule to the accepting molecule. 61

6.1 Comparison of the electric potential distributions measured at an air-water interface, where z=13.4 Å is the Gibbs Dividing Surface. 65
6.2 Average values of the electric potential, ϕ, for different cavity sizes and heights. The Gibbs Dividing Surface is at $z = 13.4$ Å, with bulk water above and vapor below. 65
6.3 Counts of oxygen atoms as a function of height over the entire course of a simulation. The vertical line denotes the Gibbs dividing surface, which is used as the interfacial height, z_i. 67
6.4 Electric potential distributions in the center of a cavity from simulations of bulk water with three sizes of cuboid simulation cells and one cubic cell all collapse onto each other with the corrections discussed in the text. 68
6.5 $\psi(q;\sigma,z) = \frac{F(q;\sigma,z)-F(-q;\sigma,z)}{2q}$ calculated for four sizes of cavities either at the interface - defined as the Gibbs dividing surface - or in the bulk. DCT with an interface predicts that each of these functions should be a constant. 69
6.6 $\Delta\psi(q;\sigma) = \psi(q;\sigma,z_i) - \psi(q;\sigma,z_b)$ for four cavity sizes. 70

List of Tables

3.1 Particle masses and counts for 2D and 3D in the more detailed cytoplasm simulations. 20

Acknowledgments

My growth as a scientist over the last six years is thanks both to the physics and chemistry communities I have had the privilege of being a part of for the last six years and my parents for their unending moral support. My advisor, Phill Geissler, has an invaluable but all too rare combination of scientific skills: he is able to conduct research with precision and mathematical rigor and to communicate it clearly to others. It's a tough balancing act and I hope some of that has rubbed off on me during my time in his group. The original ideas in chapters five and six are primarily due to him. I would also like to acknowledge my previous advisor, the late David Chandler. He was a great scientist and my two years with him surely had a large impact on me during those formative early years of graduate school. I hope that he is satisfied with the final form of my second chapter, which reports work I conducted in his group and which I edited with the lesson I learned from him in mind: don't back down.

PI's aren't there every week, but older graduate students and post-docs are. D.C. occasionally referred to David Limmer as my big brother. It was a fitting descriptor for a great mentor who was finishing graduate school just when I was getting started. I was thrilled to see him come back as an Assistant Professor. He and Aaron Keys made me feel especially welcome when I joined the Pitzer Center. Later on, Jonathan Landy was a pleasure to work with. His collegial style and humility in presenting impressive results always welcomed collaboration. Pat Shaffer had a similarly self-effacing way of describing his work. Science would be less stressful if more scientists were like them. I owe a special debt to Pat because my fifth and, even more so, my sixth chapters are extensions of his good work. Kranthi Mandadapu has been an encouraging and inspiring source of new ideas; he'll take an interest in anything you mention to him and has a breadth of knowledge across chemistry and engineering I haven't seen matched in anyone else. I wasn't surprised to see him get a job as a UC Berkeley Assistant Professor, too. Berkeley chose well in David Limmer and Kranthi. Nadine Schwierz was another helpful and productive collaborator and I regret that we didn't have the time to finish all our work together. Anna Schneider continues to be a friend; it's been exciting to watch everything she's been able to do in the worlds of tech and non-profits with a PhD from the Geissler group. In my last years here, Steve Cox has been both a productive collaborator and supportive friend. As a collaborator, he played a major role in my sixth chapter. I wish him and Julie well as they return to England for Steve's well-earned fellowship. Milo Lin was another source of new ideas during my PhD, reflected in my third chapter. Our slowly diffusing collaboration has turned into an exciting project that we'll be finishing off together in Texas. Milo has been a particularly helpful mentor as I enter the next stage of my career.

Graduate school can often be demoralizing and I wouldn't have made it through without the close friendship of other graduate students. Alex Hudson was one in those early years, as he continues to be. Katie Klymko's boundless optimism was often a boon during especially discouraging weeks. I'm glad the two of them found each other; they make quite the power couple. I gave serious thought to abandoning ship midway through and I likely would have if I didn't have friends like them. Plus, the marathons we've run together should add some

healthy years to my life. John Haberstroh and I have followed a path that at every step was closer than I would have guessed: hotel-mates at UIUC's visit weekend, classmates at Berkeley, super-group mates in the Pitzer Center, group mates in the Geissler group, roommates, and now office mates. Maybe next we'll be professor mates. In these last couple years, Layne Frechette has been a friendly addition to the Geissler group and a great friend to me. Judging not only by his *Science* paper in his second year of graduate school, but also our many conversations in the day-to-day grind of doing science, I expect him to go far. Finally, Kelsey Lyberger has been very patient with and supportive of me during this last, stressful year of graduate school. Luckily I'll get to return the favor when she starts writing her dissertation in a couple of years. They say you're stuck with the friends you make before you turn 30. I think I've found a good lot.

Katie Klymko once said my parents seem to have a higher opinion of me than most parents do of their children. At every bump in the road, they have encouraged me to press on. If my dissertation is half as good as they expect it to be, it will be very impressive indeed.

Chapter 1

Introduction

Thank you for asking. -Marshawn Lynch

In science, we seek the simplest possible models that can describe wide variations in behavior by a given system. Reducing complex behaviors to the simplest set of rules possible satisfies several goals. First, those rules are often able to reduce confounding, complex phenomena to more intuitive models. Some systems can be described by mean field theory, in which a single constituent of the system is said to interact only with the average of the rest of the system. Second, after all irrelevant details are stripped away, problems that once looked very different from each other may start to look very similar. At its most profound, statistical mechanics unveils universalities, relationships which are independent of the details of the material or system under investigation. The prototypical example is the renormalization group (RG) theory, developed by Nobel Laureate Kenneth Wilson in 1971 [1] [2] building on work by Leo Kadanoff [3]. From this abstract perspective, a phase transition between water and steam looks just like a magnet changing its poles.

I will discuss two classes of systems in this dissertation: supercooled liquids and ions in water. Supercooled liquids are those liquids which are below their freezing temperature but have not formed a crystalline solid. As their temperature is lowered, they turn into an amorphous solid, known commonly as glass. Glass is a very unusual state of matter. Microscopically, it is disordered just like a liquid. Macroscopically, it can be extremely strong. Its viscosity grows so large that it can be even more rigid than a crystalline solid. In a regular liquid, the viscosity grows exponentially in $\frac{1}{T}$, but fragile glass formers are defined by the fact that their viscosity grows even faster than exponentially in $\frac{1}{T}$. Explaining why they behave this way is one of the great remaining challenges in statistical mechanics. Two competing hypotheses are random first order transition (RFOT) theory and kinetically constrained models (KCM's). These two theories differ significantly in their underlying assumptions about the behavior of individual atoms in a glass. Unfortunately, it is very difficult to probe glasses experimentally at the resolution that would be required to refute one or both of these theories. One alternative route is to compare the theory to very detailed atomistic models simulated on a computer. In chapter 2, I will argue that recent computational experiments are well explained by KCM's.

In chapter 3, I will discuss a recent, surprising application of the theory of glassy dynamics. In response to environmental changes, bacteria can switch between active and dormant phases. The physical processes underlying this change are not yet understood. One possibility suggested by recent experiments is that bacteria actually exploit the glass transition. Switching between glassy and liquid phases may allow bacteria to preserve their inner structure for long periods of time without metabolizing. The structural similarity of supercooled fluids and equilibrium fluids allows the bacteria to return to the active, metabolizing phase quickly when environmental conditions improve. Using simulations of simple models of a bacterial cell, I will demonstrate that this hypothesis is physically plausible.

In part II, I change to the topic of ions and electric fields in water. An ion is approximately a hydrophobic cavity with a charge or dipole at its center. The solvation energy of an ion, therefore, depends on the spontaneous fluctuations of the electric field at the center of a cavity. In chapter 4, I will study the electric field fluctuations not on a foreign cavity inserted into water, but on the hydrogen atoms that are already a part of the water. These fluctuations lead to autoionization - the spontaneous breaking apart of two neutral water molecules into hydroxide and hydronium ions. Although a macroscopic volume of water has a large number of such ions in it, a given water molecule can go many hours without breaking apart. Such long timescales are completely intractable in a computer simulation. I will describe enhanced sampling techniques I have employed to collect many samples of this rare event in a computer simulation. I will then use those samples to describe the configurations of water molecules that conspire to cleave a hydrogen atom from its oxygen atom.

In chapter 5, I will move on to the problem of developing a Gaussian field theory of ion solvation. This challenge is another of the great open questions in statistical mechanics and it has been of particular interest at UC Berkeley for some years. Lum-Chandler-Weeks theory, developed at UC Berkeley in the 1990's, successfully described the solvation of hydrophobic cavities. Now, we seek to develop a similar theory for the more complicated solvation of ions. This project began before my dissertation - in particular, see the work of Patrick Shaffer - and will continue after mine. The first barrier to developing an effective field theory is choosing the right order parameter. Chapter 5 explores possible choices of order parameter. I find that the density field of hydrogen bonds is a promising option. It is highly Gaussian and sensitive to the presence of an ion.

Finally, in chapter 6, I will discuss the behavior of ions near interfaces. The simplest possible models of ions near interfaces predict that they will repel each other. In fact, the exact opposite happens: ions are strongly attracted to interfaces. Besides that attraction, there is significant charge asymmetry. Anions are much more strongly attracted to liquid-vapor interfaces than cations with identical radius and charge magnitude. The most popular theory for this interfacial attraction and charge asymmetry assumes that the polarization at the interface changes linearly in the presence of an ion and calculates the response of the ion to that polarization. I show, however, that the surface polarization changes non-linearly in the presence of an ion. The linearly responding surface potential hypothesis ultimately makes predictions that are quantitatively and qualitatively incorrect. In pursuit of a simple theory, the surface potential hypothesis goes one step too far. Its predictions are based only on the mean and variance of the electric potential at the center of a cavity. Following closely on the heels of work by Patrick Shaffer, I demonstrate that the higher moments of that distribution are required to calculate an ion's solvation free energy.

These five projects span a diverse set of systems in physical chemistry and biophysics. The primary unifying theme is the importance of fluctuations, or, alternatively stated, heterogeneity. In glassy systems, the dynamics are extremely heterogeneous. At any given time a small fraction of the system is free to move, but most is jammed. In bacterial cells, a small force injected into the system by metabolism is able to switch this jamming on and off, allowing the cell to go about the business of life or to wait for better days. In water, very rare fluctuations in the configuration of molecules result in one unlucky victim being torn apart. Although rare, the fraction of water molecules autoionized determines the pH of water, of central importance to biology. Developing a Gaussian field theory of ion solvation depends on finding an order parameter that captures the important details, and nothing more, of how water rearranges in the presence of an ion. Finally, correctly predicting charge asymmetry at liquid-vapor interfaces relies on understanding the non-linear fluctuations of the surface polarization when an ion draws near. These fluctuations are not merely higher order corrections that improve the quantitative accuracy of theories that are already approximately correct. Rather, theories that ignore these fluctuations completely miss qualitative features of the systems they seek to describe.

Part I

Supercooled Liquids

Chapter 2

The East Model Correctly Predicts the Behavior of Simulated Glassy Systems in High Dimensions

Supercooled or slowly compressed liquids maintain liquid-like structure but exhibit sluggish dynamics that, at low temperature, become slower than those of crystalline solids. The relaxation times of fragile glass-formers exhibit super-Arrhenius growth below an onset temperature, T_o, and above another crossover temperature, T_x. Below T_x the dynamic constraints of glassiness can be avoided at a constant energy cost E_x, giving a relaxation time $\tau_x \exp(E_x/T)$ [4]. Wolynes and other theorists explain the behavior of fragile glass-formers using mean-field theory for their random first-order transition theory (RFOT) [5]. Recent analytical work has demonstrated that RFOT has an upper critical dimension of 8 [6], followed by possible numerical verification in simulations of hard spheres [7]. We and other theorists, however, explain the behavior of fragile glass-formers using dynamical facilitation, a fundamentally non-mean-field concept [8, 9, 10]. Here, we use the East model [11] and comparisons to experimental data to argue that glassy dynamics are fluctuation dominated in all finite dimensions. We numerically observe that up to 13 dimensions the dynamics are fluctuation dominated and argue that this result should hold for any finite dimensionality. From there, we argue that relaxation of the East model in high dimensions qualitatively matches the behavior of glassy hard sphere simulations in high dimensions.

The East model is a d-dimensional lattice with the simple Hamiltonian

$$H = \sum_{i=1}^{N} n_i, \tag{2.1}$$

for a system with N spins, where we call $n_i = 0$ a jammed spin and $n_i = 1$ an excited spin or an excitation. The kinetics are constrained such that a spin at $\vec{x}_i$ can only change state if at least one of its nearest neighbor spins at $\vec{x}_i - \hat{u}_l$ is excited, where $\hat{u}_l$ is a unit vector in any dimension $l \leq d$. We call a spin that can change state facilitated, indicated by the function

$$f_i(\{n_{\vec{x}_i}\}) = 1 - \prod_{l=1}^{d}(1 - n_{\vec{x}_i - \hat{u}_l}),$$

where f_i is 1 if spin i is facilitated and 0 if not.

Below an onset temperature, T_o, the relaxation time grows as

$$\log\left(\frac{\tau}{\tau_o}\right) = \frac{E}{T} + J^2\left(\frac{1}{T} - \frac{1}{T_o}\right)^2, \tag{2.2}$$

where log with no subscript is taken to be the natural logarithm, τ_o is an elementary time for molecular relaxation, E is the energy barrier to relaxation in the simple liquid, and J is the energy scale for hierarchical dynamics. First, we generalize an asymptotic argument due to Sollich and Evans [12] to predict the curvature of the relaxation time in decreasing temperature in any number of dimensions and verify it numerically. This argument leads us to expect that the East model's upper critical dimension is ∞. Second, we measure the onset temperature of hierarchical dynamics in the East model in dimensions one through 13 and find that it decreases non-linearly in increasing dimensionality. We use this fact to interpret prior results from high dimensional hard sphere simulations. Third, we measure the diffusion of probe particles in the East model and find that their diffusion violates the Stokes-Einstein relation, even above RFOT's upper critical dimension of eight. Furthermore, our error analysis leads us to doubt claims that hard spheres have an upper critical dimension of eight.

2.1 Asymptotic Result for High Dimensional Curvature

The presence of hierarchical dynamics is indicated quantitatively by a non-zero value of the curvature, J. Before checking numerically if the curvature goes to zero in any dimension, we would like at least an asymptotic prediction. Sollich and Evans [12] made a heuristic argument for the curvature of the East model relaxation time in one dimension that can easily be generalized to any dimension. Following their reasoning, we consider the system's relaxation after a quench from high temperature, such that there is an excess of excited spins and flipping a spin down is approximately irreversible. Consider a d-dimensional domain spanning the cuboid from the origin to $(L-1, L-1, ..., L-1)$, consisting of an excited spin at the origin and no other excitations within the domain, but adjacent to another excitation at $(L, L, ..., L)$. Many trajectories exist to relax the spin just outside the domain. The energy barrier to a particular trajectory is the maximum number of spins that are simultaneously excited at any point along the trajectory. We will call the lowest such energy barrier the height, $h(L^d)$, where we consider only domains with all side lengths equal. Suppose $L = 2^n$ for some integer n. One possible path passes through a configuration with the spin at site $(2^{n-1}, 2^{n-1}, ..., 2^{n-1})$ excited, and all sites with any coordinate greater than 2^{n-1} remaining jammed. This configuration requires at least one more excitation than the original case, but

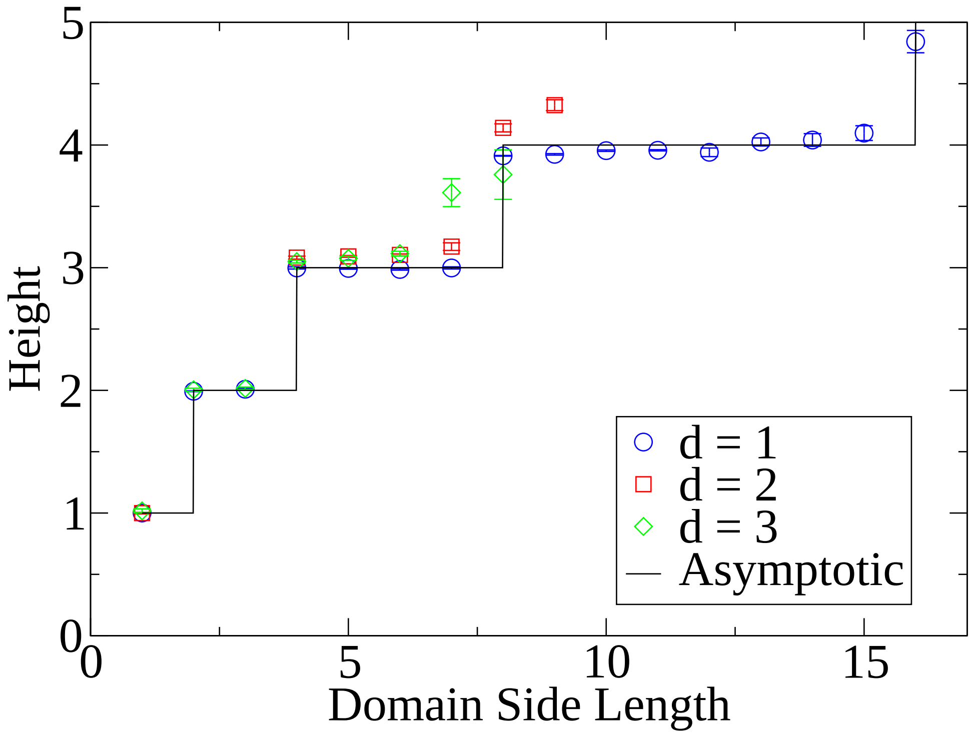

Figure 2.1: The height function, h, measured numerically for the East model in dimensions 1, 2, and 3 and plotted against the asymptotic result from equation 5.

leaves a domain with half the side length. This path, therefore, gives us the recursive relation $h(2^{dn}) = h(2^{d(n-1)}) + 1$. Sollich and Evans generalize this argument, in one dimension, to domains of lengths, M, that are not necessarily powers of two, giving, in any dimension,

$$h(M) = n + 1 \text{ for } 2^{d(n-1)} \leq M < 2^{dn}$$
$$h(M) = 1 + \lfloor \log_{2^d}(M) \rfloor$$

In turn, this implies $\log(\tau_M) = 1 + \lfloor \log_{2^d}(M) \rfloor \beta$, where τ_M is the relaxation time for a domain of size M, in the regime where M is much smaller than the equilibrium domain size. In our simulations, we measure the relaxation time, τ, as the $1/e$ decay time of the persistence function $\langle P(t) \rangle = \left\langle \frac{1}{N} \sum_{i=1}^{N} p_i(t) \right\rangle$, where $p_i(t)$ is 1 if site i has never flipped up to time t and 0 otherwise and the angled brackets indicate an average over independent trajectories. The relaxation barrier, which we call the height, of a domain is then the slope of $\log(\tau)$ versus β for that domain. See the resulting heights in Figure 2.1. Our simulations demonstrate this asymptotic argument holds in dimensions 1 through 3.

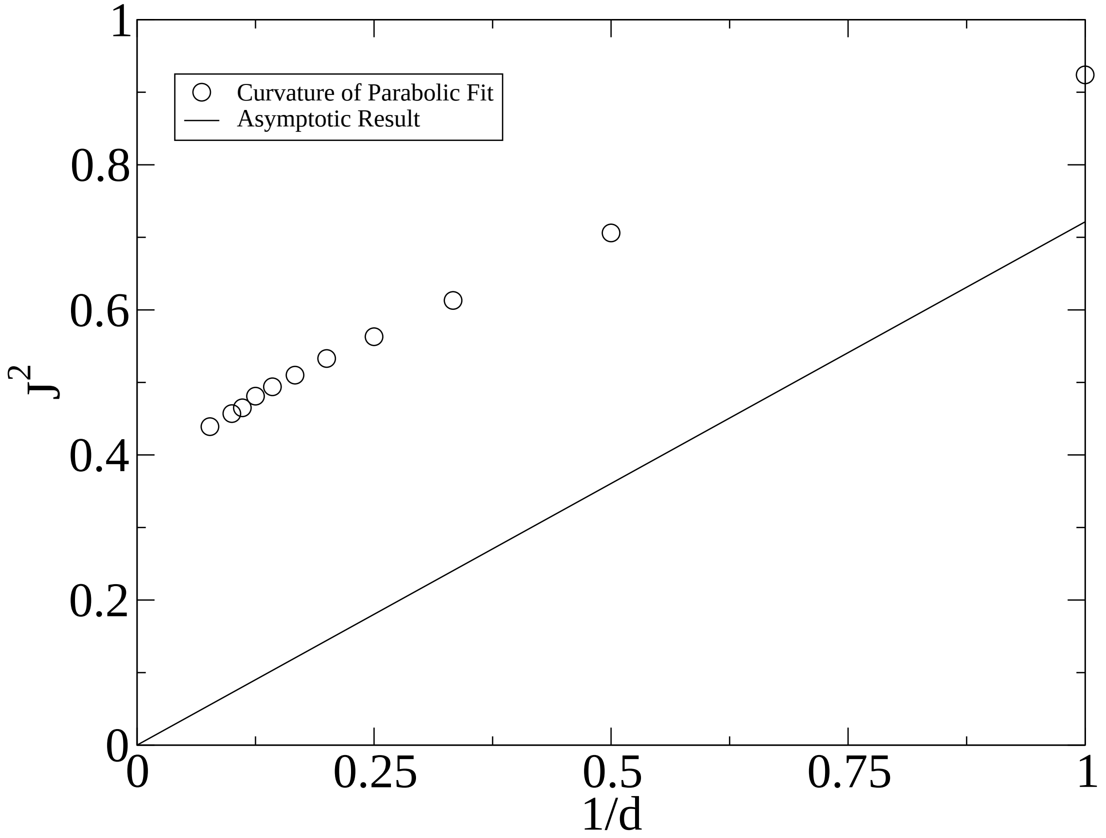

Figure 2.2: The curvature, J^2 in equation 3, observed in equilibrium simulations has an approximately constant offset from the asymptotic result $J^2 = 1/2d\log(2)$.

Given τ_M, we can substitute the equilibrium domain size, M_{eq}, and calculate the equilibrium relaxation time, keeping in mind that the equilibrium domain size is outside the rigorous result's range of applicability. A recent and more detailed calculation by Martinelli and coworkers [13] that accounts for the entropy of all possible relaxation pathways, though still in one dimension, reduces $\log(\tau)$ by a factor of 2. Assuming this extends to higher dimensions, we get

$$\begin{aligned}\log(\tau) &= \frac{\log(M_{\text{eq}})}{2T\log(2^d)}\\ &= \frac{1}{2T^2 d\log(2)},\end{aligned}$$

implying $J^2 = 1/2d\log(2)$.

To test this prediction in simulations, we chose our system sizes to have at least 100 excitations at the lowest temperature simulated in a given dimension. In one dimension, our lattice has 5,041 sites. In 13 dimensions, each side of our lattice has only three spins, for a total of 3^{13} spins. In larger dimensions, it becomes intractable to simulate enough spins to have 100 excitations. For experimental systems and atomistic simulations, accurately fitting the relaxation time of hierarchical dynamics requires first measuring the contribution from simple liquid dynamics by fitting a line to the data above the onset temperature. However, the dynamics of the East model lack any simple liquid regime and are entirely determined by dynamical facilitation, so we measure J by setting the parameter E in equation (2.2) to 0 and fitting the remaining parameters simultaneously over the data available at all temperatures. See the resulting curvatures in Figure 2.2. The analytical result qualitatively matches the trend seen in the curvature across dimensionality calculated from simulations, but with a roughly constant offset. Extrapolating the observed linear trend to infinite dimensions implies that even in that limit the system still has positive curvature. In infinite dimensions every lattice site is connected to every other lattice site, so if even one lattice site has its spin up, every lattice site in the system is facilitated. If every lattice site is facilitated, the curvature of the log relaxation time is zero. This apparent contradiction requires further investigation, but simulating an East model above 13 dimensions becomes extremely difficult.

2.2 Hierarchical Dynamics Appears at Lower Temperatures in Higher Dimensions

Reorganizations in a glass forming liquid are rare events we can treat approximately using transition state theory with a single, temperature dependent barrier to relaxation: $\log\tau \propto \beta E(\beta)$. From above we know that after a temperature quench the energy barrier is $E(M_{\text{eq}}) = \frac{1}{2}[1 + \log_{2^d}(M)]$ where $M_{\text{eq}}(\beta) = c^{-1}(\beta) = 1 + \exp(\beta)$ is the equilibrium domain size. Then

$$\begin{aligned}\partial_\beta \log\tau =& E(\beta) + \beta\partial_\beta E(\beta)\\ =& \frac{\log(1+\exp(\beta))}{2d\log 2} + \\ & \frac{\beta}{2d\log 2}\frac{1}{1+\exp(\beta)}\exp(\beta)\\ & \to 2E(\beta)\end{aligned}$$

Let T* be the temperature at which the slope of the relaxation time versus beta crosses 2:

$$\left.\frac{d\log\tau}{d\beta}\right|_{T^*} = 2. \tag{2.3}$$

See the values of T^* we measured in our simulations of the East model in Figure 2.3. In an experimental system or atomistic simulation, the onset temperature is the temperature where dynamical facilitation becomes significant. The East model's dynamics, however, are

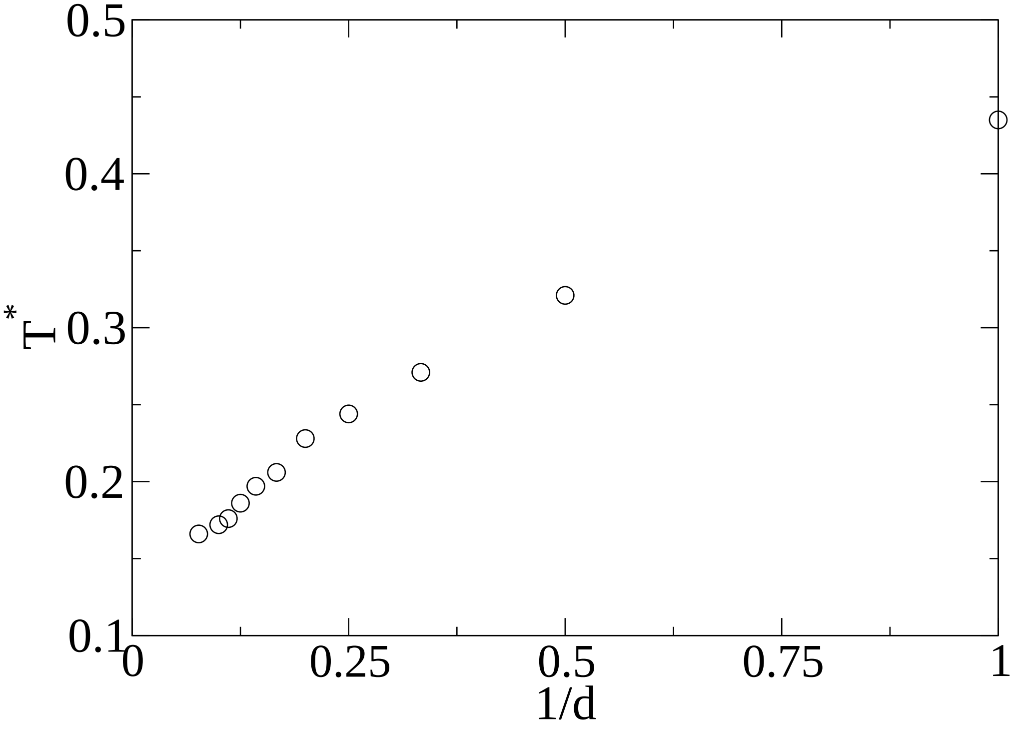

Figure 2.3: The onset temperature as defined in equation 12.

controlled by facilitation at all temperatures. Hierarchical dynamics with an infinitesimal energy scale would be physically irrelevant, so measuring the temperature at which the energy scale for hierarchical dynamics crosses some finite threshold is a preferable way of comparing the East model to real glasses. Although the choice of threshold is arbitrary, different choices for which T^* is still well defined have the same qualitative result that T^* decreases monotonically as d increases. In higher dimensions there are more pathways for one spin to relax another, so a given relaxation time is reached at a lower temperature. Simulations of hard spheres or other atomistic systems across dimensions cannot be compared at fixed packing fraction or temperature. Doing so results in comparing systems with different energy scales and dimensions, confounding any changes due only to the differing dimensionality. It would even be possible to compare a high dimensional system above its onset temperature to a low dimensional system below the onset, giving the appearance of a fragile-to-strong crossover in dimensionality where none exists. If T^* for the East model or T_o for experimental systems is nonzero for all finite dimensions, these systems are fundamentally non-mean field.

2.3 Diffusion Decoupled from Stokes-Einstein In Any d

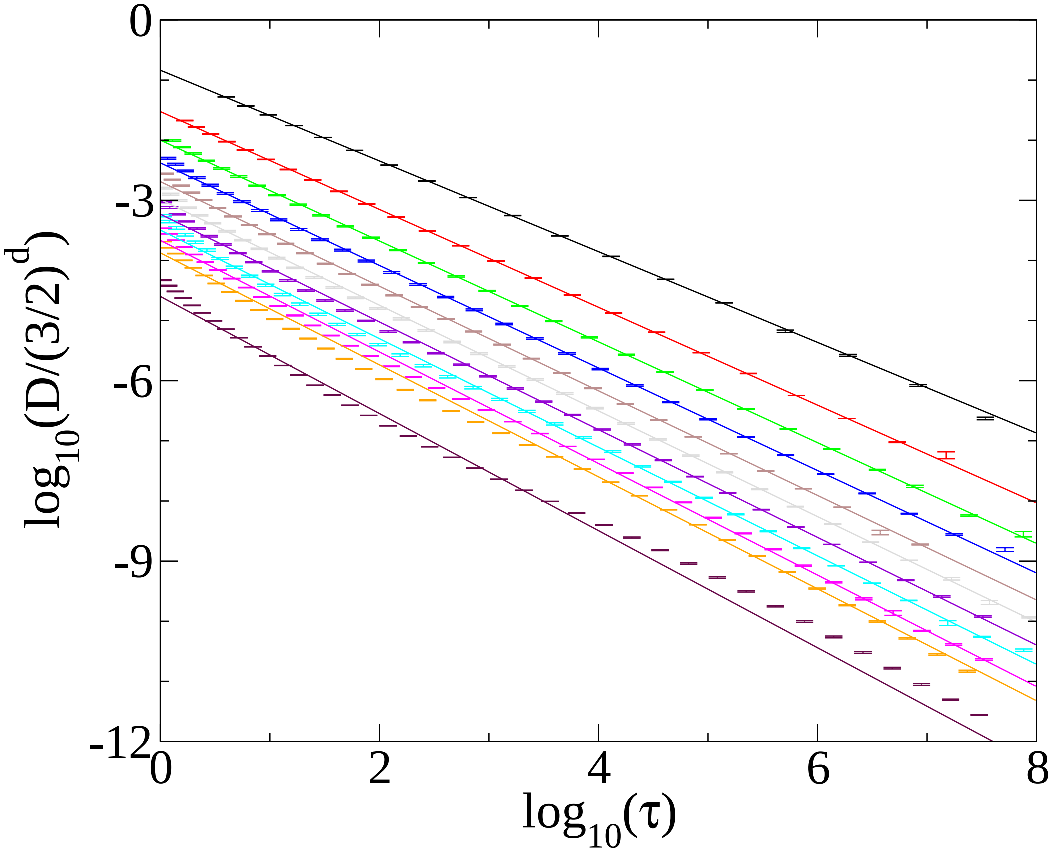

Figure 2.4: We show diffusion constants for the East model across eight decades in relaxation time. Dashed lines show linear best fits calculated over $\log_{10}(\tau) \in (2, 8)$.

Another test of mean-field theory in the East model is the diffusion of probe particles coupled to the lattice spins. After each Monte Carlo sweep over the spins, the probe particle attempts a move. If a particle is located at a site with a jammed spin, it does not move. If it is located at a site with an excited spin, it randomly chooses a nearest neighbor site. If that site also holds an excitation, the particle moves to that site. If that site holds a jammed spin, the particle does not move. If mean-field theory correctly described the diffusion of these particles, they would obey the Stokes-Einstein relation, $D \propto \tau^{-1}$. Instead, we observe that diffusion of the probe particles fails to obey the Stokes-Einstein relation; see Figure 2.4. To test that we have run our simulations at sufficiently low temperatures to capture the asymptotic behavior, we vary the minimum temperature we use for fitting the slope of

$\log D$ versus $\log \tau$ and obtain qualitatively similar results (not shown) to those in Figure 2.4. The diffusion constant instead goes as $D \sim \tau^{-\xi} = \tau^{-1+\omega}$ where $\xi < 1$. This was previously observed in one dimension across 12 decades of relaxation times [14, 15] and is seen here up through 13 dimensions over 8 decades of relaxation times. Figure 2.5 shows our fits to ω. The slope of the diffusion data violates the Stokes-Einstein relation in dimensions one through ten. In dimension 13, ω is still positive but the error bar includes ω=0, which prevents a strong conclusion but is consistent with the possibility of an infinite upper critical dimension.

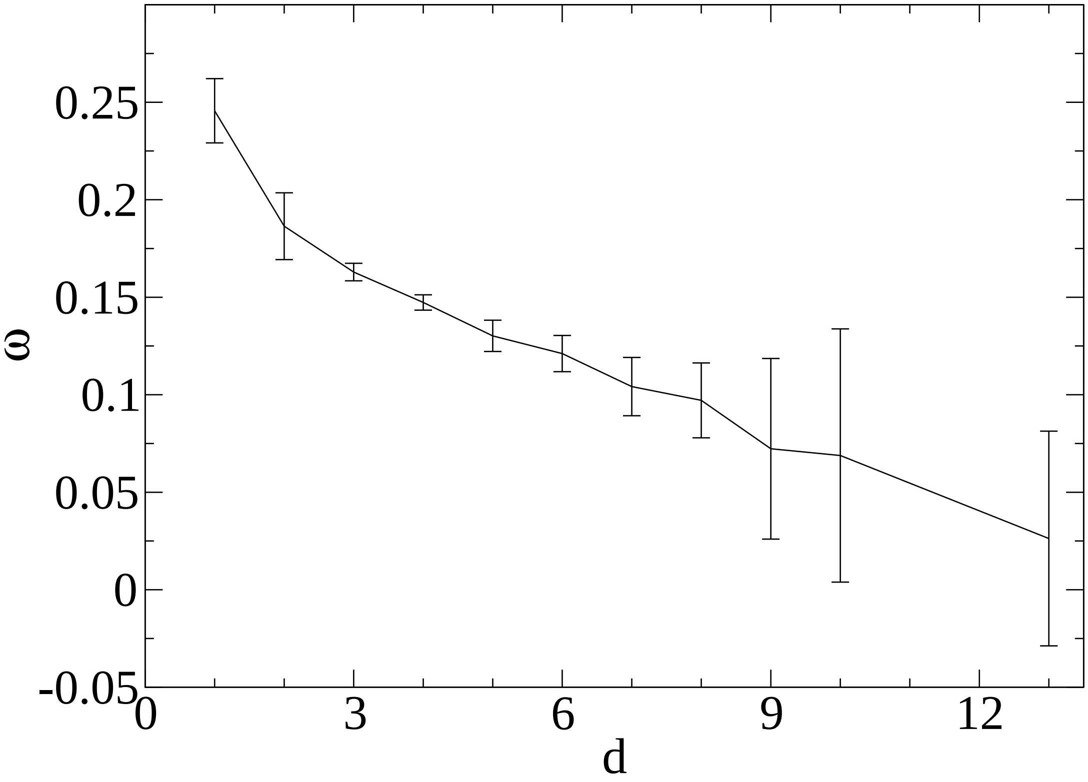

Figure 2.5: The deviation of the slope of $\log_{10}(D)$ versus $\log_{10}(\tau)$ from that predicted by the Stokes-Einstein relation in dimensions one through 10 and 13.

We calculated the error bars for ω with the following procedure. First, we fit to the three $\log_{10}(D)$ versus $\log_{10}(\tau)$ points with the lowest value of τ and calculated the mean square error between the fit and those three points. Second, we sequentially added points at lower values of τ until the mean square error rose above 0.01. Finally, we recorded ω for each fit and reported the error as half the difference between the maximum and minimum values. The middle of the reported range is the value obtained from the widest range in

τ. Charbonneau et al. [7] use ξ as the key test of a hard sphere system's upper critical dimension, observing that ξ rises approximately to one in eight dimensions and remains there in nine and 10 dimensions. However, their fits are over a much narrower range of data in the relaxation time, covering as few as two decades with only seven individual points. We find that propagating error from the diffusion constant data into ω using the Central Limit Theorem yields significantly smaller error bars than does comparing fits to different subsets of our data. Similar analysis may also yield much larger error bars for the slope of D versus τ in hard sphere systems, casting doubt on the calculation of their upper critical dimension.

2.4 Conclusion

The fundamental question in the theory of glassy dynamics is whether the behavior of a glass results from a mean field, which ultimately implies that there will be a first-order phase transition at finite temperature, or that behavior results from kinetic constraints, which then implies that there is no first-order phase transition and that the relaxation time diverges only at zero temperature. One way to test the mean field approximation is to find its upper critical dimension, where the approximation becomes exact. Simulations of an atomistic model of glass above that dimension should then reveal whether the mean field-based theory is correct. Zamponi and co-workers found that the upper critical dimension of RFOT is eight [6]. They then went on to measure the behavior of a glassy hard sphere system up through 10 dimensions [7]. When measuring the degree of violation of Stokes-Einstein diffusion, ω, they find that it decreases quickly with dimension. They go on to claim that the decrease is linear in dimension until reaching 8 dimensions, when the violation goes away.

We agree that ω decreases quickly with dimension, but the difficulty of the simulations in high dimensions and the wide error bars resist a final conclusion that ω decreases linearly until hitting zero in eight dimensions. Here, we show that ω also decreases quickly, though not linearly, in the multidimensional East model. We highlight sources of error that grow with dimension and affect simulations of lattice models such as the East model as well as atomistic models such as those simulated by Zamponi and co-workers. In work with Soree Kim, Juan Garrahan, David Chandler, and YounJoon Jung, we have demonstrated that finite size effects become severe for the very small systems that are tractable in high dimensions but that the system sizes used here are sufficiently large to have reached a plateau in τ and D versus system size [16]. Additionally, the onset temperature decreases with growing dimension; equivalently for hard sphere systems, the onset packing fraction increases with growing dimension. This drift in the onset could result in a glassy system looking liquid simply because the onset temperature has dropped below the lowest temperature for which the simulation is tractable.

We argue that even in the face of the high-dimensional hard sphere data and the upper critical dimension of RFOT, kinetically constrained models remain well suited to explain glassy dynamics. In the East model ω drops very quickly with dimension, a result qualitatively consistent with that for hard spheres. However, ω remains above zero at least up through 13 dimensions and likely up through any finite dimension of the East model, which remains consistent with our claim that real glasses will not show a first-order phase transition at any finite temperature.

Chapter 3

Glass-like Bacterial Cytoplasm Is Fluidized by Driving in Simulations

The possibility that living organisms have learned to exploit thermodynamic phase transitions in the course of evolution is an exciting line of interdisciplinary research in physics and biology [17]. Near a phase transition, correlation lengths diverge, which might help a living organism or group of organisms coordinate a very rapid response to a small change in the environment. On a macroscopic scale, flocks of birds and swarms of insects show behavior that resembles a system near a critical point. Individuals who are too far apart to interact directly still move in a correlated direction. As an example of biological phase transitions gone awry, neurons can transition into a hyperactive state and cause a seizure.

One widespread biological phenomenon that remains poorly understood is bacteria's ability to alternate between active and dormant states [18]. Somehow, when environmental conditions are unfavorable, bacterial cells are able to stop metabolizing and enter a dormant state. Later, the bacteria are able to reverse this transition and resume metabolism [19]. Somehow, the bacteria are able to preserve the proteins and ribosomes that are essential to their survival, even in inhospitable conditions. It is difficult to understand how bacteria can survive in a dormant state given their lack of internal cytoskeleton. They do not, for example, enter a crystalline solid phase.

A possible explanation for the mechanism of this behavior comes from experiments that found the bacterial cytoplasm - everything in bacterial cells except for the nucleus - turns glassy when the cell's metabolism stops [19]; that is, without becoming substantially colder or crystallizing, diffusion of particles inside the cell drops precipitously. In this chapter, I will consider ways in which that bacterial dormancy could be enabled by the onset of glass-like dynamics. This hypothesis is a slight variation on the theme of biological phase transitions described above because the onset of glassy dynamics is not actually a phase transition, but a super-exponential slowing down of the dynamics. It would also make for an unusual result in the study of glassy dynamics because the early results suggest that the largest particles turn glassy while smaller particles continue to behave like a normal liquid.

The central result from the experiments is that the diffusion of large particles slows down much more than that of small particles when the cell stops metabolizing. After reviewing the experimental results, I will describe a minimal model of a cell that includes glassy dynamics. I build the model on the assumption the cell is near the onset of glassy dynamics. By construction, as I lower the temperature, T, in my simulations the inverse of the diffusion constant, $\frac{1}{D}$, will grow super-exponentially in $\frac{1}{T}$. My first hypothesis is that particles of different sizes will either crossover for simple liquid behavior to glassy behavior at different temperatures or that once they crossover, their super-exponential slowing down will happen at different rates. Either result would be surprising and new.

This hypothesis is not borne out by the simulation results, but fortunately they do ultimately suggest a path towards a better model. The behavior seen in experiments on the diffusion of particles in the bacterial cytoplasm is not quite true glassiness. The experiments find a material in which particles diffuse much more slowly than a reference state but have not crystallized. This behavior does match the properties of a glass that make glass unusual and useful. However, the experiments have not shown that the bacterial cytoplasm is a liquid that slows down super-exponentially in $\frac{1}{T}$, a necessary condition for a material to be deemed a glass. To be clear, I am not claiming that the experimental results are wrong. Rather, my original hypothesis over-interpreted the statement that, “the bacterial cytoplasm exhibits physical properties typically associated with glass-forming liquids approaching the glass transition” [19]. The experimental results are better interpreted as suggesting a useful analogy between glasses and the bacterial cytoplasm, not a direct application of the theory of glasses to the bacterial cytoplasm.

Following these results, I will change my minimal model of bacterial cytoplasm to include non-equilibrium driving forces. My second hypothesis is that the bacterial cytoplasm of a non-metabolizing cell exhibits physical properties of glass-forming liquids that disappear when metabolism begins. As I will soon show, modeling a cell as a densely packed collection of many sizes of spherical particles far away from the glassy regimes of temperature or density gives an improved reproduction of an actual cell compared to my previous, glassy model. When I add an external force to the smallest particles in the system, which are roughly proportional to the size of motor proteins found in actual bacterial cells, the larger particles around them speed up as well. Even better, the speed-up of the non-driven particles is roughly proportional to their size. The largest particles are the ones most strongly affected by the presence or absence of a driving force on the smallest particles. This qualitatively captures the central experimental result: the largest particles tracked in the experiments had the largest speed up when the cells started metabolizing.

3.1 Bacterial Cytoplasm: Experiments

Although other labs have worked on related experiments before and since, the impetus for this project came from the groups of Corey O’Hern, Eric Dufresne, and Christine Jacobs-Wagner at Yale, and in particular their 2014 *Cell* paper, “The Bacterial Cytoplasm Has Glass-like Properties and Is Fluidized by Metabolic Activity” [19]. Diffusion in simple liquids is a thoroughly studied and well understood problem [20]. At least two properties of the bacterial cytoplasm make it a little less simple, though. First, it is composed of particles that vary by orders of magnitude in size. At one extreme, a water molecule is approximately 3 Å across. At the other extreme, the ribosome is about 200 Å in diameter [21]. Second, as the cell metabolizes it pumps energy into the cytoplasm, driving the system far from equilibrium.

While studying the dynamics of a protein that had been made non-functional by the attachment of a bulky GFP tag, Jacobs-Wagner and associates found that in metabolizing bacterial cells the large probe diffused through the entire cell in minutes via a random walk, but that in a starved, non-metabolizing cell the probe rarely made movements of any measurable size over a period of nine hours. They obtained qualitatively similar results with multiple probe particles in *Caulobacter crescentus* and *Escherichia coli*, so the behavior is not unique to a specific bacterium or probe particle. As a further robustness check, they forced the cells into a non-metabolizing state three different ways. First, they starved the cells until they stopped metabolizing. Second, they selected late-stationary-phase cells that had stopped metabolizing of their own accord. Third, they treated the cells with 2,4-dinitrophenol (DNP), which depletes the cells of ATP and GTP, again arresting their metabolism.

To quantify these differences, Jacobs-Wagner and associates measure the probe particles' mean square displacement (MSD), where $\mathrm{MSD}(t) = \langle (\vec{x}_t - \vec{x}_0)^2 \rangle$. In an infinite or periodically replicated system, the diffusion constant is $D = \lim_{t->\infty} \frac{\mathrm{MSD}(t)}{2dt}$. However, in a finite, bounded system $\mathrm{MSD}(t)$ will reach an asymptote, so at long times the MSD will grow sub-linearly. At very short times, the particles move ballistically - without collisions - and their MSD grows quadratically in time. That is, for a particle moving at constant velocity, $\langle (\vec{x}(t) - \vec{x}(0))^2 \rangle = (vt)^2$, where v is the particle's velocity. So, to measure diffusion constants in the finite environment of a cell, we will look for some intermediate time regime where the MSD is growing approximately linearly. Experimentally, at these intermediate times the MSD's of probe particles in the metabolizing cells grows at rates that are orders-of-magnitude faster than in the non-metabolizing cells.

If this were the entire result, the experiment would provide little new information. The metabolizing cell is a highly driven environment, so it comes as no surprise that particles diffuse more quickly inside metabolizing cells than in non-metabolizing cells. The surprise is that the magnitude of the speed-up in diffusion between the two kinds of cells depends on the size of the probe particle. Jacobs-Wagner and co-workers are able to tune the number of GFP particles that stick to the central molecule of their probe, letting them construct probes of different sizes. Then, they measure the diffusion of the various probe particles in both metabolizing and non-metabolizing cells. As usual for diffusion, larger particles diffuse more slowly than smaller ones in the same conditions. However, unusually the ratio between the diffusion constants in the metabolizing and non-metabolizing cells also depends on the probe particle size, with larger particles experiencing a greater increase in diffusion when metabolism restarts.

Initially, some might argue that these results are a trivial consequence of motor proteins essentially turning off when metabolism stops. However, the probe particles are inert objects that diffuse passively through the cell even when it metabolizes. They are not actively driven by any of the molecular motors inside the cell. Alternatively, the differences in diffusion may be attributed to a rise in temperature during metabolism. First, even if the temperature of the cell drops slightly when it stops metabolizing, that does not explain the orders-of-magnitude reduction in diffusion. Second, even if metabolism effectively raised the cell's temperature a great deal, that would increase the diffusion of all particles proportionally, contrary to the size dependent effect that is actually observed. Finally, Jacobs-Wagner and co-workers report control experiments where they chemically inactivated molecular processes that might somehow specifically drive the probe particles. In their words: "Through drug inactivation of transcription, translation, or peptidoglycan wall synthesis, we also showed that metabolism-dependent motion does not originate from the sole (motor-like) action of RNA polymerases, ribosomes, or penicillin-binding proteins on their DNA, mRNA, or peptidoglycan substrates" [19].

More to the point, the experiments also show certain properties that might be interpreted as dynamical heterogeneity, a hallmark of glassy physics. Rather than measuring MSD's, Jacobs-Wagner and associates computed a related quantity, the radiation of gyration, R_g, the root-mean-square distance from the center of a given particle's trajectory. It is, in effect, the MSD averaged over the time of observation. They report that the distribution of R_g's for the probe particles is bimodal. The MSD's of the particles with $R_g < 0.3$ μm plateau to a low, constant value, indicating they are caged. The MSD's of the particles with $R_g > 0.3$ μm grow linearly for 10 minutes and grow to values commensurate with the size of the cell, indicating free diffusion.

Although these measurements do prove the existence of two populations of probe particles with very different dynamics, they are very different from the "heterogeneous dynamics" discussed in the previous chapter. The meaning of glassy dynamics is explained thoroughly by Aaron Keys, Lester Hedges, and their advisors in their 2011 paper "Excitations are Localized and Relaxation is Hierarchical in Glass-Forming Liquids" [22]. In brief, at any given time in a glassy system, most particles are vibrating over a distance that is much smaller than their own diameter. Excitations, localized in time and space, are groups of atoms that are able to make persistent displacements in their positions, on the order of their own diameter. When particles in one location rearrange, they make room for adjacent particles to rearrange as well. This process allows the excitation to move slowly through the entire system. Whether or not the excitations themselves can diffuse freely through the whole system, individual atoms do not. As the excitation moves through the system, the identity of the atoms composing the excitation changes. If any single atom could move freely through the whole system, the system would be a liquid and not a supercooled fluid. Also, a histogram of R_g for each atom in the system over a long trajectory in an atomistic simulation of a glass would not be bimodal. As the excitations move through the system they do not favor or disfavor the inclusion of certain atoms. The results in non-metabolizing bacterial cells are not the same as glassy dynamics. Taking this analogy literally is the mistake of my first hypothesis. My second hypothesis changes from a precise consideration of glassy dynamics to a focus on how driving preferentially speeds up larger particles.

3.2 Bacterial Cytoplasm: Glassy Simulations

How, exactly, could the characteristics of glassy dynamics explain the observed phenomena? Recall that the salient feature of glassy dynamics is super-Arrhenius growth of relaxation times. That is,

$$\log\left(\frac{1}{D}\right) = \frac{E}{T} + J^2\left(\frac{1}{T} - \frac{1}{T_o}\right)^2, \tag{3.1}$$

where D is the diffusion constant, E is the energy scale of relaxation in the simple liquid, J is the energy scale of hierarchical dynamics, T is the temperature, and T_o is the onset temperature for hierarchical dynamics. The consequence is that small changes in temperature lead to abnormally large changes in the diffusion constant compared to the response of a typical liquid. However, size dependent effects are not a component of the onset of glassy dynamics in systems studied to date.

There are two possibilities for how glassy dynamics could relate to the extreme and size dependent changes in the diffusion of probe particles in bacterial cells. First, the energy scale of glassy dynamics could itself scale with particle size. If J is larger for larger particles and if bacterial cells live in a glassy regime, then small changes in the temperature could cause a large, size dependent slow down. Second, the onset temperature, T_o, could be higher for larger particles. If a cell that stops metabolizing drops below the onset temperature for a large particle but is still above the onset temperature for a small particle, that could also cause a large, size dependent slow down.

Either one of these results would suggest that bacteria utilize the physics of glassy systems in order to preserve themselves during unfavorable environmental conditions. Jacobs-Wagner and co-workers highlight that the bacterial cell is polydisperse, crowded, and driven. Polydispersity is a necessary component of molecular simulations of glassy dynamics without size dependent effects. The standard molecular model of a three dimensional glass is the Kob-Andersen mixture of 80% particles with diameter σ and 20% particles with a radius 0.88σ [23]. Two dimensional glass models also require polydispersity [22]. At high packing fraction or low temperature, a monodisperse system will just crystallize. As for the cell's crowdedness, that is another necessary ingredient of a typical glass. Below some critical packing fraction, which will vary based on the details of the intermolecular interactions, the system will fluidize [24]. On the other hand, the metabolism that drives the cells far from equilibrium should keep the cell fluidized. Supercooled fluids are in a metastable equilibrium [23], very different from a far-from-equilibrium system. Perhaps the combination of these two ingredients of a typical glass and one ingredient preventing the formation of a glass will give rise to the peculiar, size dependent arrested dynamics seen in experiment.

Particle No.	Mass (kDa)	σ (Å)	N_{2D}	N_{3D}
1	11.4	15.15	107	480
2	29.1	20.70	210	940
3	46.0	24.12	148	662
4	67.2	27.37	106	476
5	91.4	30.32	91	410
6	107.1	31.97	31	138
7	132.8	34.35	33	147
8	156.1	36.26	39	175
9	161.7	36.69	6	29
10	186.5	38.47	12	54
11	346.0	47.26	165	738
12	2700.0	93.74	22	102

Table 3.1: Particle masses and counts for 2D and 3D in the more detailed cytoplasm simulations.

To start checking this hypothesis, I need to build a minimal model of a bacterial cell with glassy dynamics. I will start with a model of a fluid, then add in the ingredients of glassy dynamics, and finally modify the system to resemble a cell. The standard model of a fluid used in Monte Carlo or Molecular Dynamics simulations is the Lennard-Jones potential [25]:

$$V_{\text{LJ}}(r) = 4\epsilon \left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6} \right],$$

where ϵ is an energy scale and σ is a length scale. The r^6 attraction term comes from Van der Waals attraction [26]. Consider two neutral atoms with some separation, R. If the relative separation of their nuclei and electrons were fixed, the center of their electron clouds would coincide with their nuclei, and there would be no reason to expect them to feel each other's presence. However, two adjacent neutral atoms induce (oscillating) dipole moments in each other because each atom's nucleus is attracted to the electrons in the other atom. The induced dipole magnitude scales as R^{-3} and the potential between two dipoles scales again as R^{-3}, so the total attraction scales as R^{-6} [26]. As the atoms closely approach each other, their electron clouds overlap. The Pauli exclusion principle then forces some of the electrons to jump to a higher energy level. Rather than deriving the exact scaling from first principles, experimental fits show the repulsion scales approximately as R^{-12} [26]. Combining Van der Waals attraction with the Pauli exclusion principle gives the Lennard-Jones potential. The Lennard-Jones potential has liquid, vapor, and crystal phases and lets us study first-order phase transitions and phase coexistence. All these properties make it a staple of statistical mechanics. However, it is difficult to form an amorphous solid with a Lennard-Jones potential.

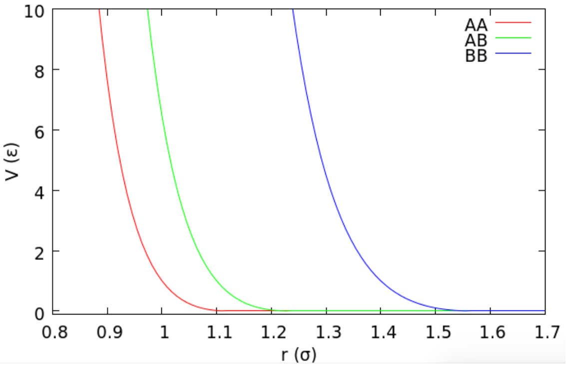

Figure 3.1: The three interaction potentials between the two particle types I simulated.

With a good model of a fluid, some small modifications give a glass. I will start with a two dimensional model, approximating the view experimentalists get when looking through a bacterium, and work up to a three dimensional model if necessary. Instead of a single particle type, I will have two, with the following parameters: $\epsilon_{\mathrm{AA}} = \epsilon_{\mathrm{BB}} = \epsilon_{\mathrm{AB}} = \sigma_{\mathrm{AA}} = 1.0, \sigma_{\mathrm{AB}} = 1.1, \sigma_{\mathrm{AB}} = 1.4$. Finally, the interaction potentials are truncated at their minimum, $2^{1/6}\sigma$, and shifted so they go smoothly to 0 where they are truncated. See Figure 3.1. This truncation removes the attractive portion of the Lennard-Jones potential and turns the particles into purely repulsive Weeks-Chandler-Anderson (WCA) particles [27]. WCA particles have the same phases as Lennard-Jones particles, but because we can truncate their interactions at $2^{1/6}\sigma$ instead of the standard Lennard-Jones truncation of 2.5σ they are faster to simulate. Previous simulations at Berkeley showed these parameter choices give rise to heterogeneous dynamics and super-Arrhenius relaxation times [22].

The final step is to take this model of glassy physics and modify it to resemble a bacterial cell. A physicist's approximation of a cell as an environment for a diffusing particle is a bounded, roughly rectangular space. A two-dimensional simulation is a reasonable approximation to the view experimentalists have. To start, I take a box with shape 15σ by 100σ. (Simulation cells with shapes 30σ x 100σ or 100σ x 100σ give qualitatively similar results, which I will not show.) Then I randomly place my mixture of A and B particles as well as a single particle of a third type with $\sigma = 10.0$. (Simulations where the large particle has $\sigma = 2.0$, 4.0, 5.0, or 7.5 also give qualitatively similar results, which I also will not show.) I make sure to place the large particle, which will stand in for the experimental probe, at least 2σ from any edge of the simulation. After energy minimization and 1,000 steps of equilibration, I freeze the particles within 2σ of the edge in order to confine the probe and to mimic the bumpy cell wall the probe particle experiences in the experiments. Then I run the system for 10^6 time steps of $\Delta t = 0.005\tau$ each, where $\tau = \sigma\sqrt{\frac{m}{\epsilon}}$, and measure the MSD of the non-frozen small particles of type A and B and the single large particle. As I repeat this procedure at lower and lower temperature, the diffusion constants for the three particle types all drop proportionally, contrary to what would be expected if my first hypothesis were

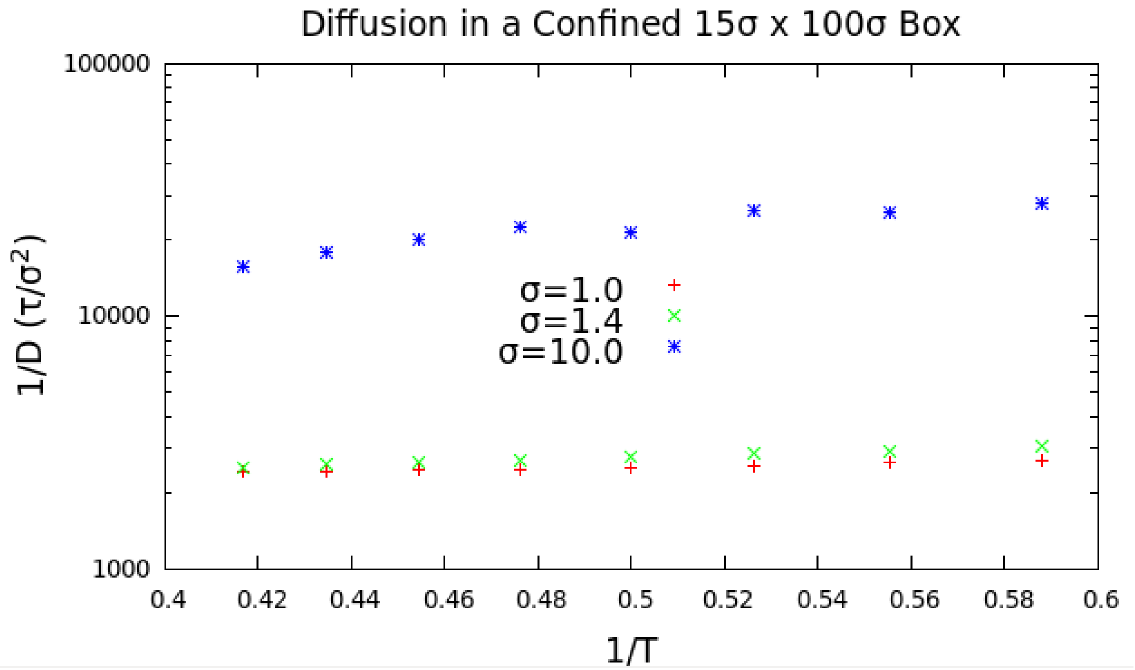

Figure 3.2: Inverse diffusion constants for the three particle sizes in a confined cell with total dimension 15σ x 100σ. There may be a slightly larger slow down for the large probe particle, but the difference between the small and large particles is nowhere near the factor of five seen in experiments.

correct. See Figure 3.2. The various particle sizes all have approximately the same onset temperature and the same curvature, J. The onset of glassy dynamics does not differentially affect particles of different sizes.

3.3 Bacterial Cytoplasm: Driven Simulations

My original model may have failed to give the desired results simply because it was too simplified. I will (partially) follow the example of Fabio Trovato and Valentina Tozzini in their simulations of the bacterial cytoplasm [28]. They define 12 particle types with the concentration and size of each particle type corresponding to experimentally determined protein abundances. Although their simulations are all in 3D, I want to run in 2D as well as 3D. In Table 3.1 I show particle types with masses and counts in 3D pulled from [28] as well as the particle counts I used in my 2D simulations. In 2D, I use a simulation cell size of 1152.5 Å x 2305 Å, for a total area of $2.656 * 10^6$ Å^2. In 3D, I use a simulation cell size of 1152.5 Å x 1152.5 Å x 2305 Å, for a total volume of $3.061 * 10^9$ Å^3. For the Lennard-Jones

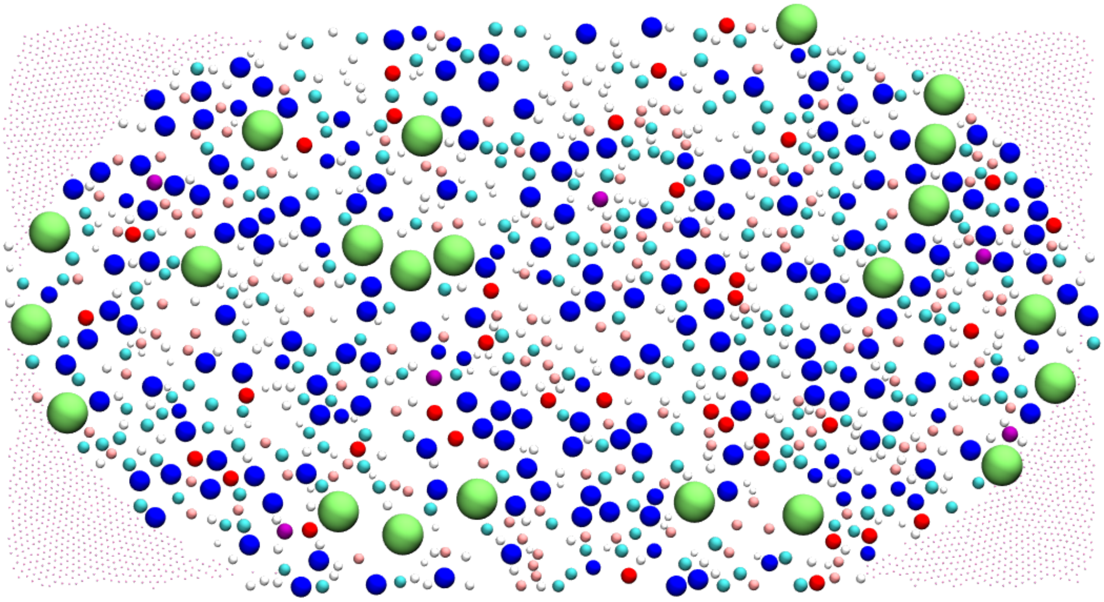

Figure 3.3: Snapshot of a 2D simulation of a bacterial cell. The smallest particles around the corners of the simulation never move; their purpose is to provide an approximately cell-shaped, bumpy confining boundary.

potentials between two particles of the same type, I set $\epsilon_i = 0.146(\sigma - 3) + 0.25$. Between different kinds of particles I use the mixing rules $\epsilon_{ij} = \sqrt{\epsilon_i \epsilon_j}$ and $\sigma_{ij} = \frac{\sigma_i + \sigma_j}{2}$. See Figure 3.3 for a snapshot from one of the two dimensional trajectories. Next, see Figure 3.4 for several views of a snapshot from one of the three dimensional trajectories.

I truncate and shift the interaction potentials at $\sqrt{2}\sigma$ to make them purely repulsive. I use 1 kcal/mol as my unit of energy, 1 kDa as my unit of mass, and 1 Å as my unit of length, so $\tau = \sqrt{\frac{m\sigma^2}{\epsilon}} = 4.889 * 10^{-14} s$. I run my simulations for $12.5 * 10^6$ time steps of 0.02τ each, for a total simulation time of approximately one minute.

In addition to adding more structural detail into my model, I need to add a key component of the dynamics that was lacking in the previous model: driving. Instead of turning the temperature up and down, turning an external force on and off is likely a better model of metabolism. A driven system is far from equilibrium, completely unlike the metastable equilibrium a glassy system is in. A very crude model of metabolism is that a small fraction of particles, representing motor proteins, is driven. I have mimicked such a system in my simulations with 15 lanes running parallel to the long axis of the cell, which I will call the x-axis. Each lane has a width of 76.8 Å, roughly as wide as five of the smallest particles. Across the 15 lanes, I have a pattern of an external force that acts to the right or to the left, or no force at all. The external force acts only on the smallest particles, parallel to the x-axis. Measuring each particle's area as $\pi\sigma^2$, the smallest particles occupy only 2.3% of the area of all the particles in the system.

In the two-dimensional simulations, the results begin to capture the qualitative features of the experiments. See Figure 3.5. A driving force that only acts directly on the smallest particles in the system causes all other particles in the system to speed up. In simulations it is simple to control for the effect of temperature. With a thermostat we hold temperature constant, so we can attribute all the increase in diffusion to the driven motion of the small particles and not just an increase in temperature. Furthermore, the ratio between the diffusion constants in the simulations with and without driving depends on the size of the particles. See Figure 3.6. Larger particles experience a greater increase in diffusion than smaller particles. The strength of the dependence on particle size is sensitive to the system's density. Making a direct comparison to a real bacterial system requires a three dimensional model with a well calibrated density.

The three-dimensional simulations, which take much longer to run, are in progress. Early results are quite promising. See Figure 3.7. After some delay, larger particles that are not directly acted upon by the driving force diffuse much more quickly along the forced direction than the other directions. The time of the delay and the magnitude of the speed up for different particle sizes and system densities will be presented in forthcoming work with Milo Lin. It is already possible to conclude that in an appropriately chosen range of driving forces and densities, an external driving force on the smallest particles in the system substantially increases the diffusion of larger particles in the system in a size-dependent manner.

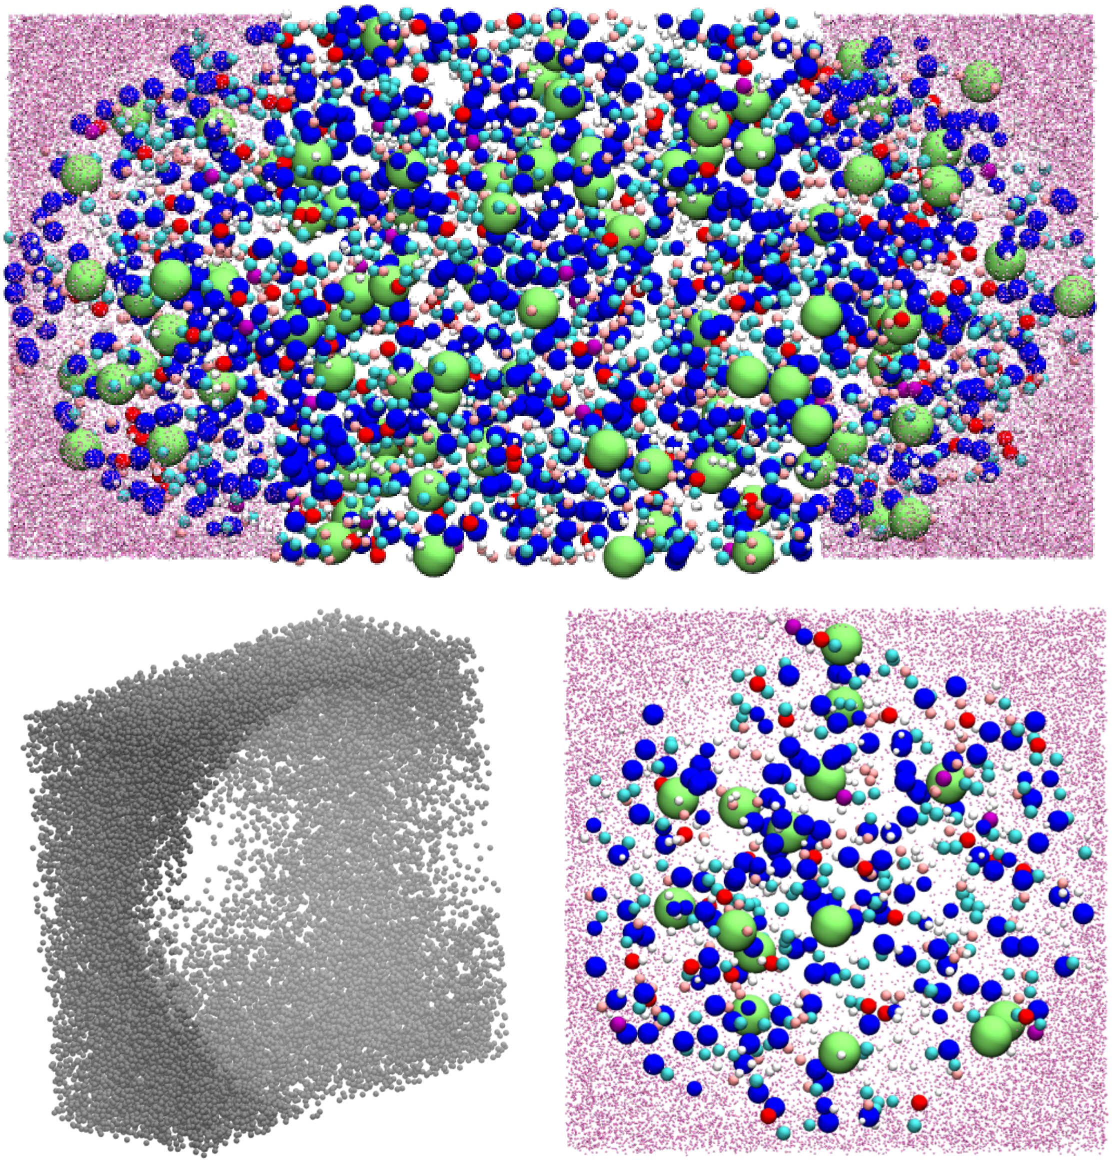

Figure 3.4: Three views of a snapshot of a 3D simulation of a bacterial cell: side view of the entire system (top); the bumpy, rounded cap made of small unmoving particles (bottom-left); a cross-section down the long axis of the cell (bottom-right).

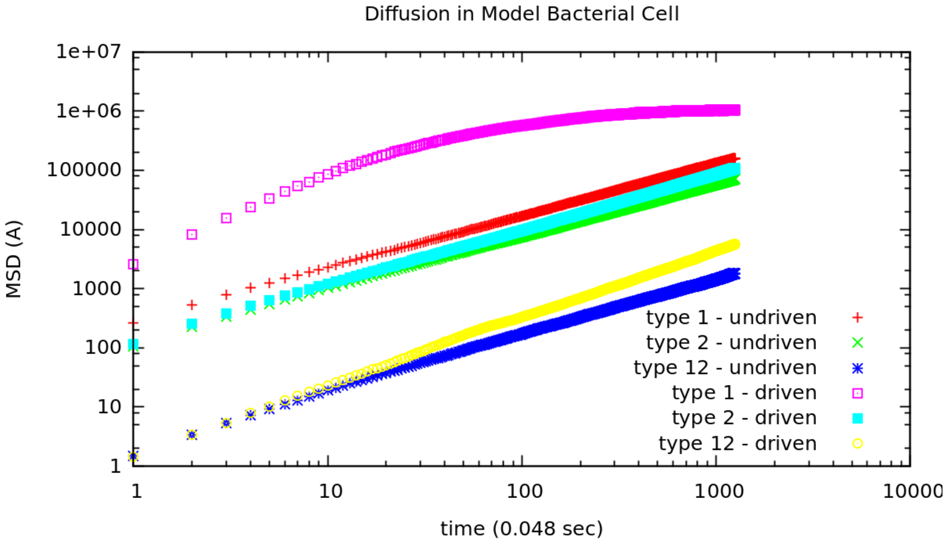

Figure 3.5: Mean square deviation of three particle sizes in the two-dimensional simulations. Adding an external force just to the smallest particles in the system speeds up the diffusion of other particles.

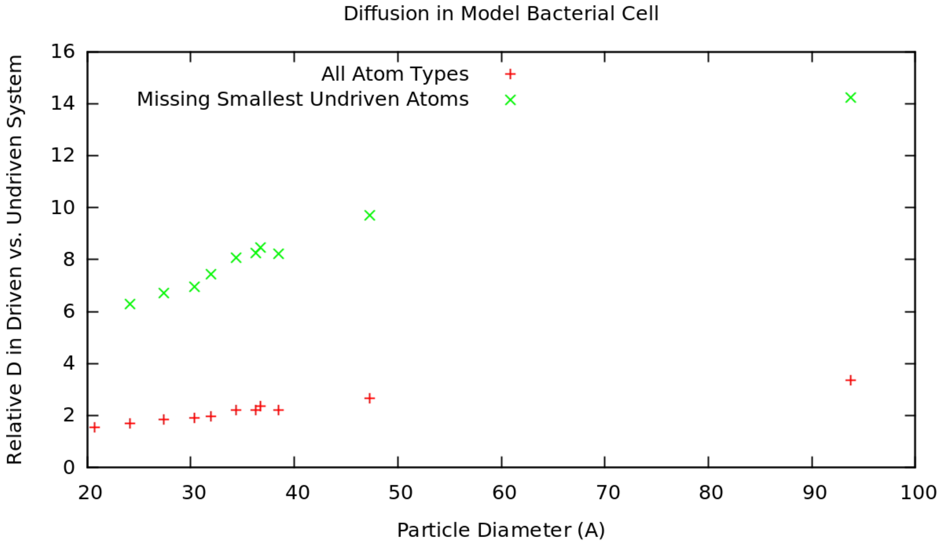

Figure 3.6: Ratio of diffusion constants in the presence or absence of an external force on the smallest particles for all other particles in the system.

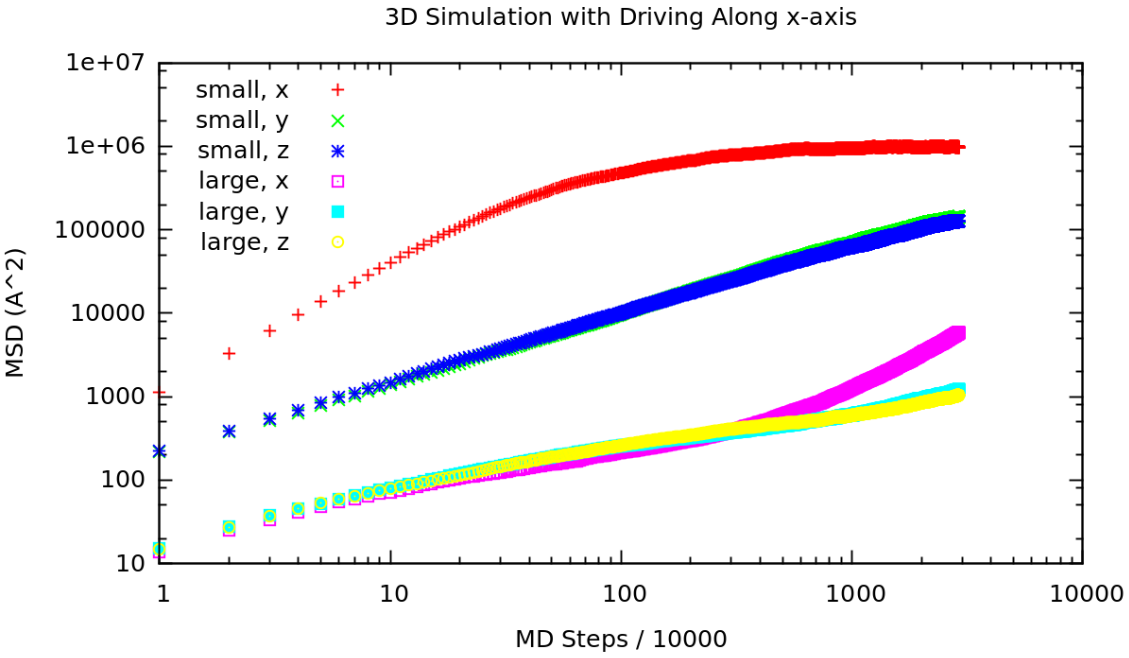

Figure 3.7: Mean square deviation of two particle sizes in a three-dimensional simulation. The MSD is measured separately along each axis. The driving is along the x-axis, which is also the long axis. Diffusion along the other axes closely resembles that in a simulation with no driving.

3.4 Conclusion

The mechanism for bacterial dormancy is not yet understood. Experiments on non-metabolizing cells show that diffusion of probe particles in these cells drops dramatically and in a size-dependent manner. Those experiments suggest that bacterial cells may exploit properties of glassy physics to preserve themselves in unfavorable environments. By selectively slowing down the motion of ribosomes and large proteins, bacterial cells may preserve their structure, which is necessary for their function in the cell when metabolism resumes.

I tested two hypotheses for how bacterial cells could utilize glassy physics. First, I tested the hypothesis that particles of different sizes in bacterial cells have different onset temperatures for glassy dynamics or different energy scales for hierarchical dynamics. If that were the case, a small decrease in temperature when metabolism stops could cause large, size dependent decreases in diffusion. However, neither of these scenarios appeared in the simulations.

Second, I tested the hypothesis that the bacterial cytoplasm without driving is jammed on the time scales over which the experimental measurements are taken, but that the addition of a small driving force on a small fraction of the particles in the system causes a large, size dependent increase in diffusion. In two-dimensional simulations, this hypothesis was correct. For better comparison to experiment, I am reproducing these simulations in three dimensions. Early results confirm the hypothesis. Although these simulations cannot prove that bacterial cells exploit properties of glassy physics in order to enter a reversible dormant state, they successfully demonstrate that such an explanation is plausible. Forthcoming work with Milo Lin and in collaboration with experimentalists should lend further credence to this result.

Part II

Water

Chapter 4

Electric Field Fluctuations in Water

In this chapter, I will use enhanced sampling techniques to study how fluctuations in the arrangement of solvent molecules can exert large enough forces on a molecule to break it apart. The central molecule can be an ion like NaCl, in which case we call this ion dissociation, or another water molecule, in which case we call this autoionization. Studying these fluctuations is a challenge because they happen very rarely - requiring a very long simulation - and may involve many water molecules - requiring a large system size.

4.1 Autoionization

Autoionization in water - the spontaneous breaking apart of a charge neutral water molecule into positively and negatively charged components - has both fundamental and practical importance in physical chemistry. Fundamentally, charge transfer is one of the main topics of study in physical chemistry and autoionization of water is one of its most prevalent examples. Practically speaking, autoionization determines the pH of pure liquid water. No matter how good a job you do of distilling your water, you will always find it to be full of ions because water molecules are constantly splitting apart and recombining. Through extensive effort and the development of entirely new tool sets at UC Berkeley and elsewhere - namely transition path sampling (TPS) - research culminating with that of Phillip Geissler finally revealed the molecular mechanism for autoionization in liquid water [29].

Somehow, a collection of water molecules must fluctuate into a configuration that makes it energetically favorable for a water molecule to split into a positively charged hydronium and a negatively charged hydroxide. Geissler and co-workers found that the appropriate reaction coordinate was the electric field experienced by the hydrogen atom on the verge of changing its bond from one oxygen atom to another. There are many, fleeting autoionization events in which the disloyal proton is quickly pushed back to its parent water molecule. The rare event is not the bond breaking, but the large electric field that quickly drives the excess proton away and breaks the chain of hydrogen bonds connecting the ions. If the hydrogen bond chain does not break, rapid recombination follows. To be clear, the proton we label "excess" at the end of this process is not the same proton that originally broke its chemical bond. Rather, when an excess proton attaches itself to a water molecule, another proton on the same molecule breaks off and starts the process again. Instead of hopping one molecule at a time, the excess charge can move across an $H_5O_2^+$ cluster (called a Zundel cation) or an $H_9O_4^+$ cluster (an Eigen cation). This process, known as the Grotthus mechanism, allows for faster diffusion of the excess charge than if a single excess proton had to diffuse through the system [30].

Directly observing autoionization in a simulation was a great challenge overcome by Geissler and co-workers. Chemical bond breaking is a quantum mechanical process. Simulating that with density functional theory may have given them the correct behavior - as long as they were lucky enough to choose correctly from a zoo of possible functionals - but scales so poorly with system size that they would have been unable to see an excess proton get even a few molecular lengths away from the hydroxide it abandoned. Instead, they used Car-Parinello Molecular Dynamics (CPMD), simulating nuclear coordinates in a classical potential. CPMD allowed them to obtain trajectories of 32 water molecules on the order of 1 ps long. Though modest, this system size is large enough to see stably separated ions. However, a given water molecule only ionizes roughly once in 10 hours. Even supposing all 32 water molecules fluctuate independently - which they do not - a single 1-ps simulation would have roughly a one in 10^{12} chance of seeing an autoionization event. They overcome this problem by developing TPS. Starting with a single trajectory that goes from a neutral state to an ionized state, they then make small perturbations to the trajectory. This serves as a form of importance sampling in trajectory space rather than configuration space.

Careful analysis of the resulting trajectories shows that the separation between the ions does not accurately describe whether they will progress to the stable, charge separated state. First, to test whether a configuration is in a transition state, Geissler and co-workers run many trajectories from the same configuration but with velocities chosen at random from a Boltzmann distribution. Call the number of bonds in the shortest hydrogen bond path between the ions l. If roughly half the trajectories end with $l \geq 3$ and the other half end in configurations with no ions, then the initial configuration is said to be a transition state. With a set of transition state configurations identified, the second task is to identify which features of a given configuration cause it to be a transition state. With an appropriate reaction coordinate, a configuration from the transition state can be allowed to fluctuate and as long as the reaction coordinate is held constant the system will remain in a transition state. If the fixed order parameter is not actually the reaction coordinate, configurations that were originally in the transition state will migrate to the initial or final state. For an example of how the probabilities change when the system is allowed to fluctuate without the reaction coordinate held constant, see Figure 4.1. For autoionization, because l is used to define the final state reached by the trajectory, it is also a natural candidate for reaction coordinate. However, for intermediate values of l some configurations have a very high probability of going to the charge separated state while others have a very high probability of going to the charge neutral state.

Instead, the proper reaction coordinate is the potential energy difference between the

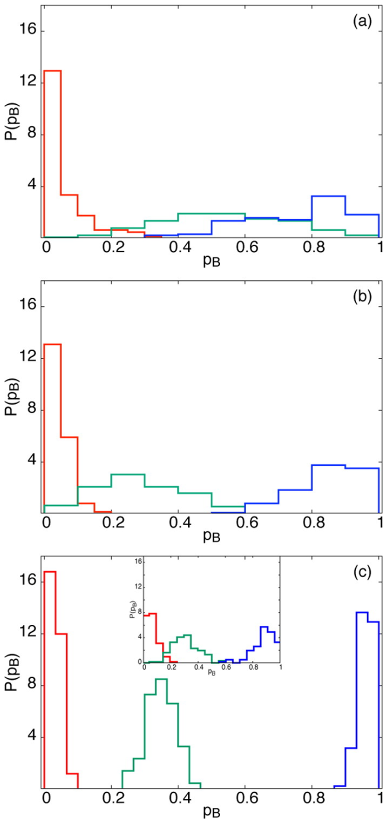

Figure 4.1: The distribution of commitor values for configurations that are equilibrated from a transition state with all atoms through the first (a), second (b), and third (c) solvation shells frozen, reproduced from [31]. Configurations that are still in a transition state would have a commitor distribution sharply peaked around 0.5. Panel (c) suggests that freezing three solvation shells may hold the configuration in a transition state, but when the process is repeated with a larger simulation cell the distribution widens substantially (inset).

two charge states when only protons along the hydrogen bond wire connecting the ions are allowed to move. To calculate this difference, Geissler and co-workers linearly interpolate the coordinates of those protons between the charge-separated configuration and a charge-neutral configuration. They hold all other coordinates fixed. As expected, a stable, neutral configuration strongly disfavors separation of the charges and a stable, charge-separated configuration strongly disfavors restoring the ions to charge neutrality. However, a neutral configuration in a transition state shows metastability of the charge-separated state. See their results in Figure 4.2.

With this reaction coordinate in hand, they would like to know which solvent fluctuations determine its value. As an initial check, to test whether the contributions to this potential energy are primarily short- or long-ranged, Geissler and co-workers remove the outer sol-

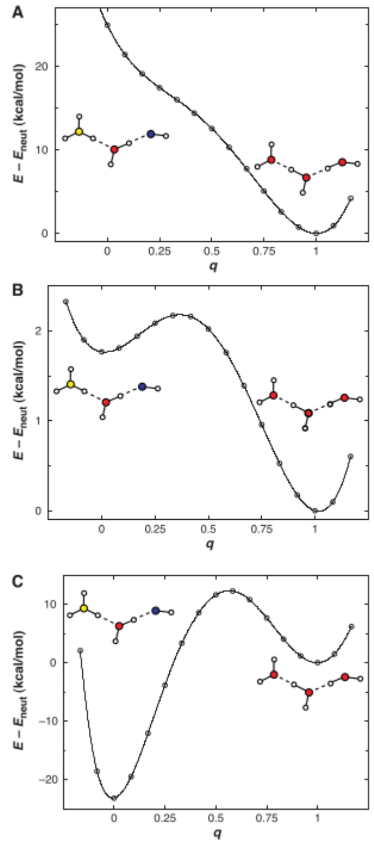

Figure 4.2: Reproduction of Figure 2 from [29]. The horizontal axis shows progression along the linear interpolation of protons between the charge-separated and neutral states. The vertical axis shows the potential energy difference from the neutral state. The top panel is calculated from a neutral configuration, the middle panel is a transition state, and the bottom panel is from a charge-separated configuration.

vation shells and recalculate the same reaction coordinate. With the outer shells removed, the stability or metastability of the charge-separated state is substantially reduced. This experiment tells us that there is some important long-ranged contribution to the electric field, but it does not tell us what that contribution is.

4.2 Ion Dissociation

Another simple case of ionization, with just a couple more ingredients than autoionization, is the dissociation of sodium chloride in water. Simulations done by Andrew Ballard and Christoph Dellago [31] show that, as in autoionization of water, solvent fluctuations are crucial to the ionization process. Similar to Geissler and co-workers, Ballard and Dellago identify transition states between the associated and dissociated states. Then, to test which coordinates are essential elements of the transition state, they take a transition state, freeze all molecules out to the first, second, or third solvation shell, and let the other molecules evolve at equilibrium. They then check whether the newly obtained configurations are still transition states. Even when they freeze three solvation shells they find that a significant fraction of the equilibrated configurations are no longer in or near the transition state (see Figure 4.1), indicating that information from the fourth solvation shell and farther contributes meaningfully to the transition state for dissociation of sodium chloride. As in the case of autoionization, we still don't know what these long-range contributions are, just that they are important to the transition state.

4.3 Long-Range Contributions to Electric Field Fluctuations

To understand how long-range solvent fluctuations contribute to the rare, large electric field that break chemical bonds, we need a method that lets us run a simulation with more than 32 water molecules for longer than a few picoseconds. I turn now to work by Bernhard Reischl, Jurgen Kofinger and Christoph Dellago [32]. Instead of using CPMD to propagate the nuclear dynamics, as Geissler and co-workers did, Dellago and co-workers used a rigid water model, TIP4P. Although this prevents them from actually observing bond breaking, it lets them run a much larger and longer calculation so they can find what solvent configuration gives rise to a large electric field.

These rigid water simulations reveal several properties of the electric field distribution on a water molecule. First, the distribution is Gaussian over the measured range. See Figure 4.3. Second, they decompose the electric field at a given location into sources as a function of distance from the point of measurement. This decomposition reveals that even for the largest fields observed, the contributions come completely from atoms within 7 Å and mostly from atoms within approximately 4 Å. See Figure 4.4.

Brute force simulations are able to tell us quite a bit about the solvent contributions to electric fields of above average magnitude. However, a given water molecule will only autoionize approximately once every 10 hours [29]. Dellago and co-workers' data come from 5 ns of simulation of 1000 water molecules. Even if every molecule were independent, this would give us only 5 ms of electric field time series, giving us roughly a 1 in 10^{10} chance of observing the conditions for autoionization emerging spontaneously. To probe this regime, we must use some sort of importance sampling to drive our simulation towards autoionization.

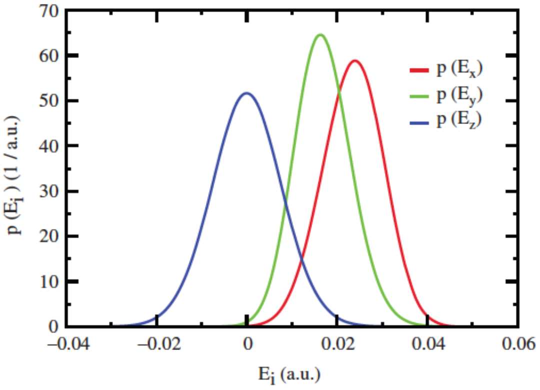

Figure 4.3: The distribution of electric fields experienced by the hydrogen atom of a water molecule reproduced from [32]. The x direction is parallel to the OH bond. The y direction is perpendicular to the OH bond and in the HOH plane. The z direction is perpendicular to the HOH plane.

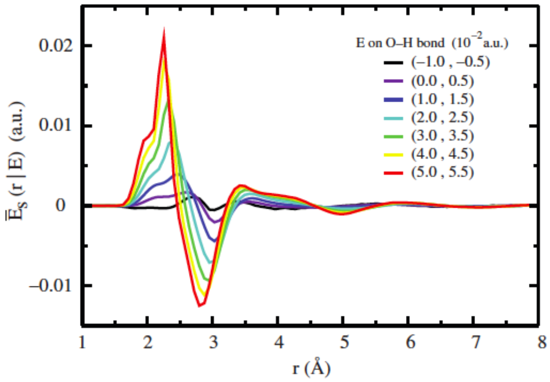

Figure 4.4: Contributions to the electric field on a central atom as a function of distance reproduced from [32]. Different colored lines are for different sizes of electric fields on the central atom.

4.4 Umbrella Sampling

The importance sampling technique I will use is umbrella sampling, which I will briefly review now. My review will closely resemble that of Patrick Varilly [33]. An unbiased equilibrium simulation will generate samples proportional to a Boltzmann weight, that is, $e^{-\beta U(\{\vec{r}\})}$, where $\beta = \frac{1}{k_B T}$, k_B is Boltzmann's constant, T is the temperature, and $U(\{\vec{r}\})$ is the energy of a configuration defined by the set of coordinates $\vec{r}$. Since the number of samples decreases exponentially in the energy, sampling high energy configurations is extraordinarily difficult. However, we can run a simulation with some energy penalty, $\Delta U(\{\vec{r}\})$, that is chosen carefully to enhance the frequency of whichever set of configurations we are interested in. In particular, suppose we want to know the mean of some variable, x, at equilibrium, $\langle x\rangle_0$, where the subscript 0 indicates an unbiased simulation. We can calculate this expected value as follows:

$$\langle x\rangle_0 = \frac{\int d\{\vec{r}\}e^{-\beta U(\{\vec{r}\})}x(\{\vec{r}\})}{\int d\{\vec{r}\}e^{-\beta U(\{\vec{r}\})}} \tag{4.1}$$

$$= \frac{\int d\{\vec{r}\}e^{-\beta U(\{\vec{r}\})-\beta\Delta U(\{\vec{r}\})}e^{\beta\Delta U(\{\vec{r}\})}x(\{\vec{r}\})}{\int d\{\vec{r}\}e^{-\beta U(\{\vec{r}\})-\beta\Delta U(\{\vec{r}\})}e^{\beta\Delta U(\{\vec{r}\})}} \tag{4.2}$$

$$= \frac{\int d\{\vec{r}\}e^{-\beta U(\{\vec{r}\})-\beta\Delta U(\{\vec{r}\})}e^{\beta\Delta U(\{\vec{r}\})}x(\{\vec{r}\})/\int d\{\vec{r}\}e^{-\beta U(\{\vec{r}\})-\beta\Delta U(\{\vec{r}\})}}{\int d\{\vec{r}\}e^{-\beta U(\{\vec{r}\})-\beta\Delta U(\{\vec{r}\})}e^{\beta\Delta U(\{\vec{r}\})}/\int d\{\vec{r}\}e^{-\beta U(\{\vec{r}\})-\beta\Delta U(\{\vec{r}\})}} \tag{4.3}$$

$$= \frac{\langle xe^{\beta\Delta U(\{\vec{r}\})}\rangle_{U+\Delta U}}{\langle e^{\beta\Delta U(\{\vec{r}\})}\rangle_{U+\Delta U}} \tag{4.4}$$

In some cases, the choice of ΔU is simple and obvious. If you want to bias a local order parameter, x, say the distance between two atoms, towards a value, x', you can add a harmonic energy penalty for deviating from that value: $\Delta U = \frac{k}{2}(x-x')^2$. In other cases, the choice is more difficult. In particular, if the parameter of interest is a global order parameter, the path from equilibrium to the desired value may pass through kinetic traps or states that aren't accessed by that simple biasing force. Additionally, to apply the bias in a molecular dynamics trajectory you need to differentiate the biasing energy with respect to every atom's position at every time step. If that is a long calculation it will substantially slow down your simulation.

4.5 Umbrella Sampling of the Electric Field on a Proton

Using umbrella sampling on the total electric field on an atom in a simulation with periodic boundary conditions is computationally intractable because it has a contribution from an Ewald sum, the derivative of which is difficult to calculate at every time step. However, the biasing energy ΔU does not actually need to be a function of the order parameter of interest, even though it often is in practice. Once you calculate the equilibrium weight of a configuration generated under any (validly defined) biasing energy, you can calculate whatever order parameter you like. As long as the biased trajectory sufficiently samples the distribution of the order parameter you are interested in, your sampling is done. In an MD simulation, the electric field is calculated as the sum of a short-ranged contribution from Coulomb's law and a long-ranged Ewald sum [25] [34]. The short-ranged contribution

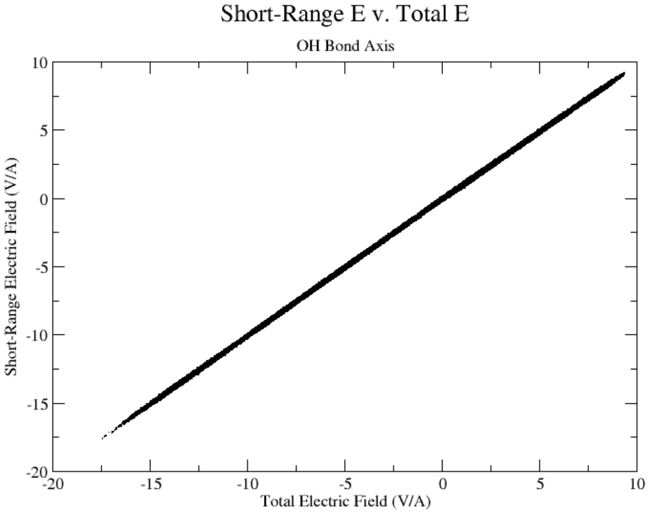

Figure 4.5: The very tight correlation between the total electric field on an atom and the electric field from nearby atoms only, excluding the long-ranged contribution from an Ewald sum, allows for importance sampling indirectly on the former by explicitly biasing the latter.

is highly correlated with the total electric field, so by adding a much simpler bias on the short-ranged electric field only, I am able to sample the distribution of the total electric field. See Figure 4.5. The electric field in a molecular dynamics simulation is calculated as

$$U_{\text{coul}} = \frac{1}{2V} \sum_{\vec{k} \neq 0} \frac{4\pi}{k^2} |\rho(\vec{k})|^2 \exp(-k^2/4\alpha) - (\alpha/\pi)^{1/2} \sum_{i=1}^{N} q_i^2 + \frac{1}{2} \sum_{i \neq j}^{N} \frac{q_i q_j \text{erfc}(\sqrt{\alpha} r_{ij})}{r_{ij}} \tag{4.5}$$

where V is the volume of the simulation cell, $\rho(\vec{k})$ is the Fourier transform of the charge density, and α is a parameter that can be tuned to ensure quick convergence of the reciprocal space sum and acceptable error. The third term is the short-ranged contribution in real space.

Adding a biasing force in LAMMPS happens through a “fix” command. LAMMPS has no fix implemented to bias the electric field experienced by an atom. I introduced a fix that adds the following bias to the potential energy of the system:

$$\Delta H = \tfrac{\kappa}{2}(\vec{v}_{\text{OH}} \cdot E_H - E_0)^2$$

where $\vec{v}_{\text{OH}}$ is the vector pointing from the oxygen atom to the tagged, covalently bonded hydrogen atom and E_{H} is the short range electric field experienced by a tagged proton. κ and E_0 are parameters I can choose in each simulation. To propagate dynamics under this additional energy term, I have to differentiate the third term in Equation 4.5 above.

Although I am only directly biasing the short range contribution to the electric field and not the total, the tight correlation between the two means I successfully, if indirectly, bias the total electric field. Now that I have the ability to bias my simulations along the reaction coordinate for autoionization discovered by Geissler and co-workers, I can get exponentially better sampling of the solvent configurations that naturally give rise to a large electric field.

Each simulation was in a cubic box with sidelength 21.73 Å with periodic boundary conditions. Each box contained 343 SPC/E water molecules. The initial configuration of the water molecules came from the last step of an equilibration run without any umbrella bias that had, in turn, started from the water molecules occupying a cubic lattice. I ran each simulation for 10^7 timesteps of 1.0 fs each, for a total of 10 ns in each window and 110 ns total, with the velocity Verlet integrator. I recorded the electric field along the OH bond on the atom subject to the umbrella bias every 1000 time steps. I held temperature at 298 K with a Langevin thermostat. For the one dimensional potential of mean force of the electric field along an OH bond, I used 11 umbrella windows with the electric field bias centered at 0, ±1, ±4, ±9, ±16 and ±25 $V/\text{Å}$. As a check that I had a properly equilibrated and stationary system, I redid all the simulations with the initial conditions chosen from the last timestep of the umbrella window biased to $-25V/\text{Å}$. For the two dimensional potentials of mean force of the electric fields along two OH bonds of different pairs of molecules, I used 121 umbrella windows with every combination of the 11 biases referenced above on the tagged pair of molecules.

4.6 Results

The energy displayed in Figure 4.6 is the same calculation as in Figure 4.3, where the previous figure shows the probability of a given electric field instead of the energy. We can now evaluate fluctuations below a probability of e^{-100}. This lets us see which solvent configurations account for these extreme electric field fluctuations. Note that, as defined here, a positive electric field would pull the tagged proton away from its parent water molecule, making that the relevant direction for autoionization. We see that those fluctuations are highly Gaussian, confirming the result of Dellago and co-workers [32], but much further into the tails of the distribution.

In the other direction, towards negative electric fields that squeeze the OH bond, what we see is a little more surprising. Above roughly 20 k_BT the energy grows sub-quadratically. Although there is no metastable state, the sub-quadratic growth of the free energy means that a sufficiently large external electric field, which creates a harmonic bias, can create a second metastable state; see Figure 4.7. 20 k_BT is a large energy for a single molecule, though, so this effect may be an artifact of pushing the SPC/E water model too far outside the conditions it was parameterized for.

What we want to know is what solvent configurations give rise to these large electric field fluctuations. In Figure 4.8 I have chosen typical configurations from the simulations run under biases towards large negative or positive values of the electric field. I only show the

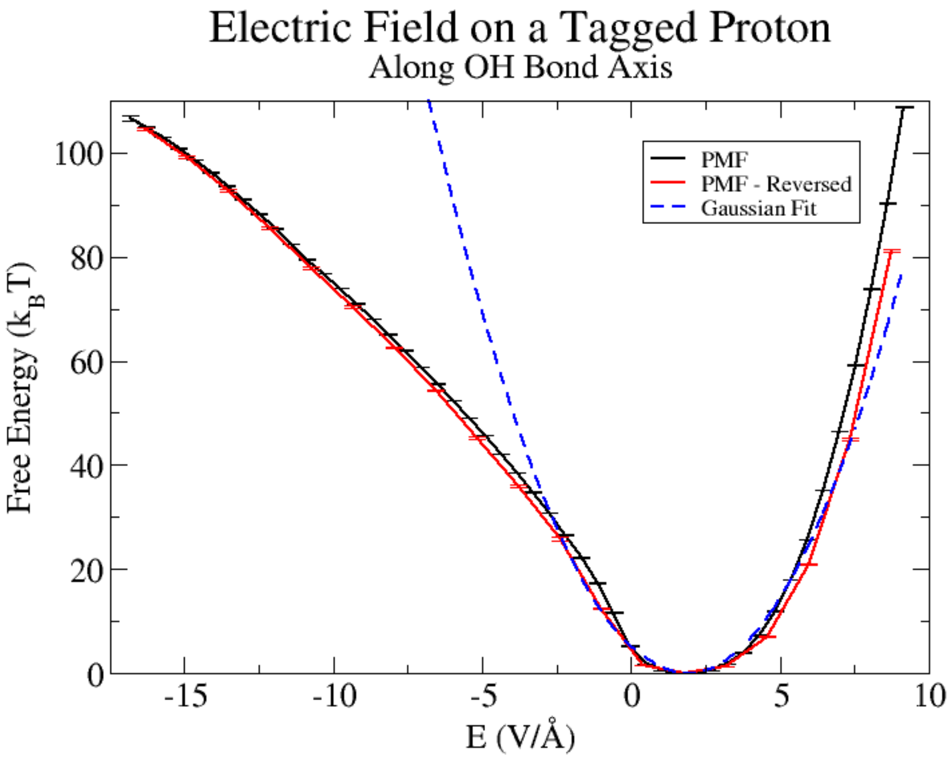

Figure 4.6: The free energy of the electric field experienced by a proton on a water molecule, projected along the OH bond axis. The black line is calculated from umbrella sampling windows that start at equilibrium and are then equilibrated to different biases. The red line is calculated from umbrella sampling windows that start from the last configuration in the trajectory with the largest, negative bias.

tagged molecule and its nearest neighbor because, in both cases, that one neighbor accounts for nearly the entire electric field on the tagged hydrogen atom. For the large, negative field, two water molecules form a dimer analagous to a Bjerrum defect in ice. A Bjerrum defect is a pair of water molecules with both donating and accepting a hydrogen bond to and from the other [35]. To my knowledge, such a defect has never been reported in liquid water and could very well be an artifact of the SPC/E model under a large external force. I will give more details of these configurations below. The configurations giving rise to a large, positive electric field seem more sensible. With a larger and larger bias, the neighboring oxygen atom pulls closer and closer to the tagged hydrogen atom.

To quantify the distribution of solvent configurations under these extreme conditions, I will calculate the radial distribution functions (RDF's) of hydrogen atoms around the tagged hydrogen and oxygen atoms around the oxygen atom bonded to the tagged hydrogen. The

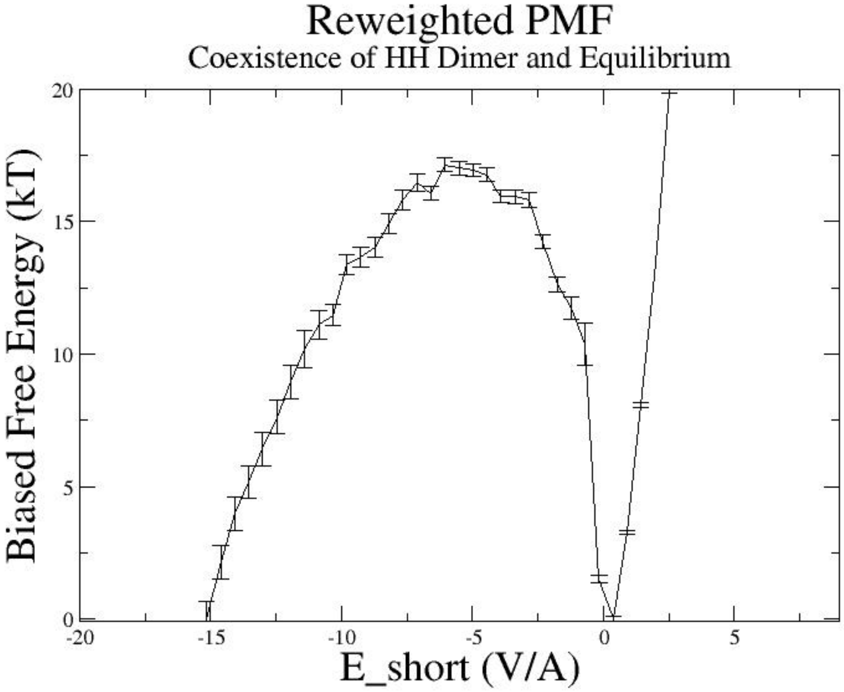

Figure 4.7: The free energy of the electric field experienced by a proton on a water molecule, projected along the OH bond axis, under a bias. The bias is a harmonic penalty centered at the most negative bin observed, approximately -15 V/Å, with a spring constant chosen to force the free energy of that bin to 0.

radial distribution function, $g_{\mathrm{AB}}(r)$, is proportional to the probability of finding a particle of type B at a distance r from a particle of type A. We can write this as $g_{\mathrm{AB}}(r) = \frac{\langle \rho_B(r) \rangle}{\rho_B}$, where $\langle \rho_B(r) \rangle$ is the density of particles of type B at a distance r from a particle of type A and ρ_B is the average density of particle type B over the whole system. For the density of hydrogen atoms around the tagged hydrogen, see Figure 4.9 and for the density of oxygen atoms around the oxygen covalently bonded to the tagged hydrogen, see Figure 4.10. Taken together, these RDF's tell us that, compared to equilibrium, in configurations with a large negative electric field the tagged hydrogen comes more than an Angstrom closer to the nearest adjacent hydrogen atom, while the nearest oxygen atom is pulled roughly a quarter of an Angstrom closer. The second shell - that is, water molecules farther away than roughly 3.5 Å- is largely unaffected. These facts suggest that the large electric field results primarily from nearby fluctuations in the solvent.

The oxygen-oxygen RDF for trajectories where the tagged hydrogen is biased towards very large and positive electric fields shows both a splitting of the first peak into two peaks

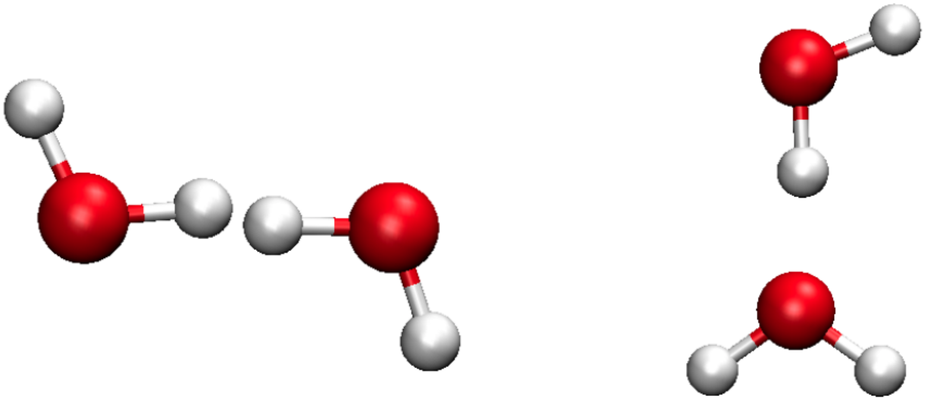

Figure 4.8: Representative configurations taken from trajectories biased towards large negative (left) and positive (right) values of the electric field experienced by a tagged proton in water, projected along the OH-bond axis.

and an enhancement of order farther out from the first shell. See Figure 4.11. The heights of the two peaks in the first shell are both higher than the single peak in the equilibrium RDF and have a similar width at their base. This new RDF suggests that the configurations with large electric fields are more ordered than the typical liquid. The molecules experiencing the large electric field have a larger coordination number and accept more hydrogen bonds than at equilibrium. See Figure 4.12. The region around the boundary between the first and second shells is significantly depleted as those molecules are pulled in towards the central atom and the first shell is pulled roughly 0.5 Å closer. The second shell peak, around 4.5 Å, and the second minimum, around 5.5 Å, are also noticeably enhanced compared to equilibrium. Farther out, the available statistics don't allow us to see any change in the liquid structure. If there is a change, it is quite small. These results suggest that the large electric field at the center results primarily from fluctuations of molecules within 4 Å, with some influence of molecules out to around 6 Å.

As a further test of whether the large electric fields on a hydrogen are due primarily to the single nearest neighbor, I bias the electric field on pairs of hydrogen atoms. In one test, I pick a hydrogen atom to bias, then find a molecule in the first solvation shell of the first molecule and bias one of its hydrogen atoms as well. In another test, I repeat the procedure for a molecule in the second solvation shell of the first molecule. In both cases, I bias the distance between the chosen water molecules to keep them in each other's first or second solvation shell. For the first solvation shell, the bias is a harmonic potential centered

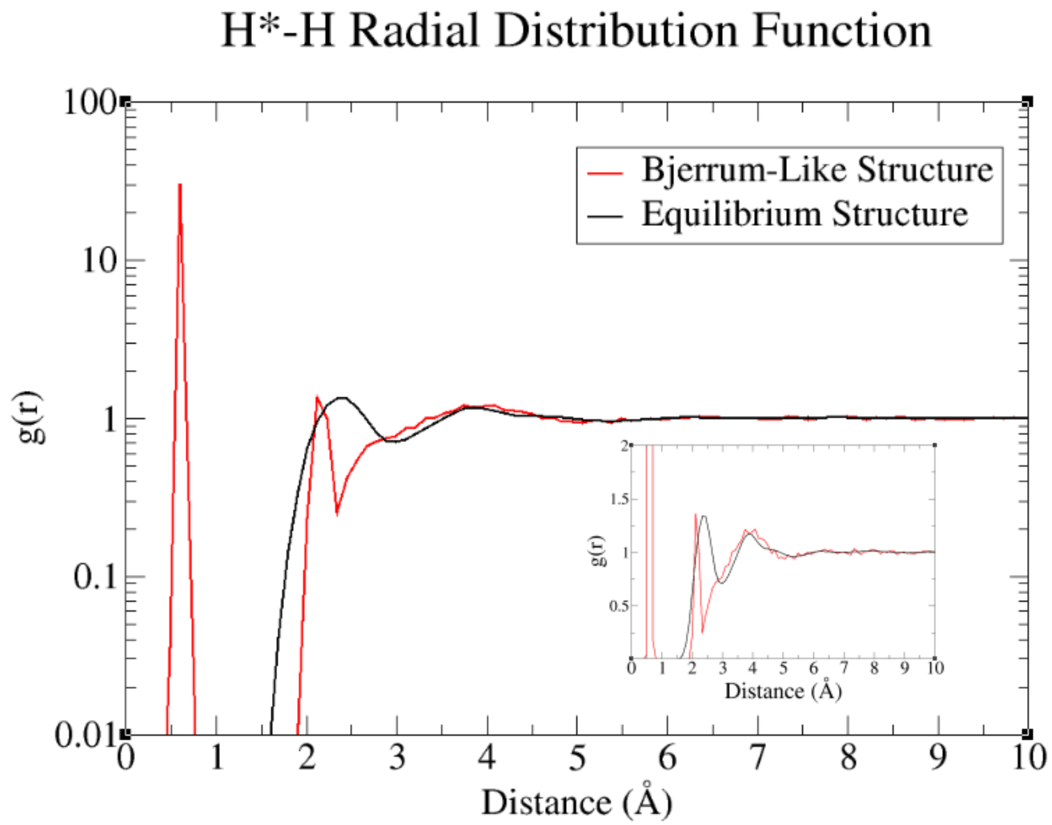

Figure 4.9: Radial distribution function (RDF) of hydrogen atoms around tagged hydrogen atom in trajectories with a harmonic bias around electric fields of -18 V/Å. This RDF quantitatively describes the configuration seen in the left panel of Figure 4.8.

on a distance 2.75 Å between the molecules. For the second solvation shell, the harmonic potential is centered at a distance of 4.5 Å. See the resulting potentials of mean force in Figure 4.13 and Figure 4.14. Representative configurations from different regions of the first shell PMF are in Figure 4.15.

To see how the biases on the two molecules interact with each other, compare to the case of what the two-dimensional PMF would look like if the two electric fields were uncorrelated; see Figure 4.16. To help make sense of these results, I also compute the difference between the uncorrelated PMF and the first shell PMF in Figure 4.17 and the difference between the uncorrelated PMF and the second shell PMF in Figure 4.18. If long range fluctuations are required to induce a large electric field at a point, then the differences between the PMF for two uncorrelated electric fields and the actual PMF's will be quite large. In fact, the differences are not very large. In the first shell, the free energy difference ranges from -27 k_BT to 36 k_BT. Compared to the uncorrelated case, the first shell pair more easily develops

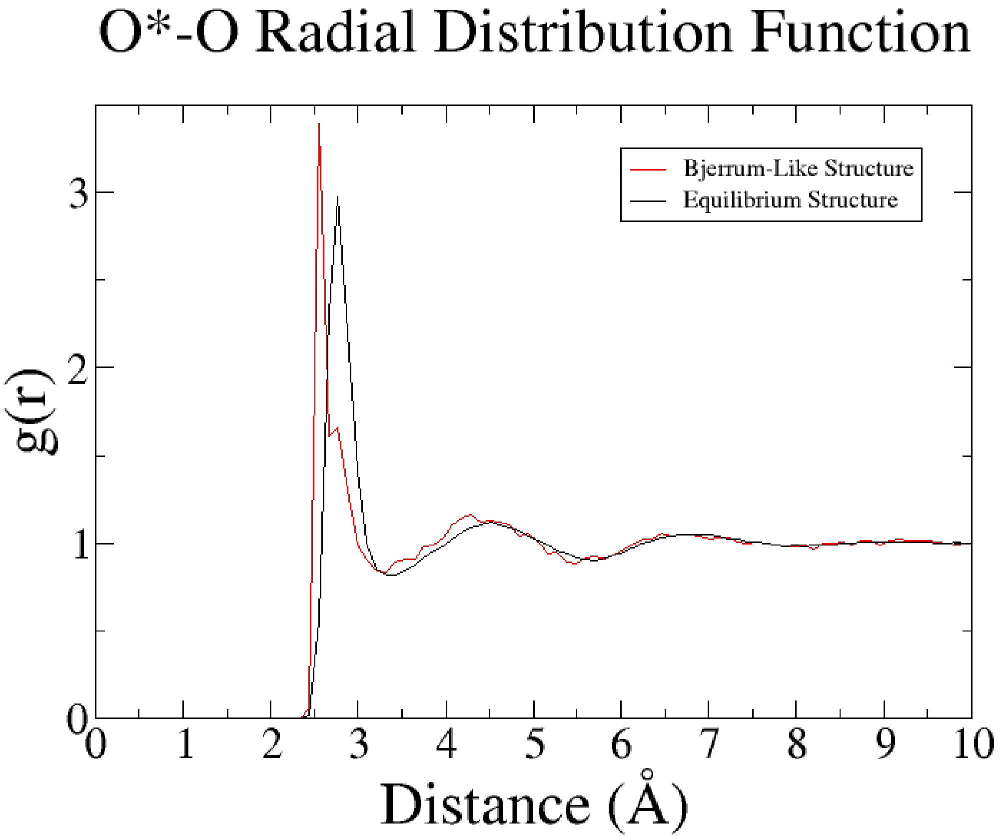

Figure 4.10: Radial distribution function (RDF) of oxygen atoms around the oxgen atom covalently bonded to the tagged hydrogen atom in trajectories with a harmonic bias around electric fields of -18 V/Å. This RDF quantitatively describes the configuration seen in the left panel of Figure 4.8.

large negative electric fields on both molecules. In the case of the second solvation shell pair, this enhancement goes away. There is no fluctuation that is more likely for the second shell pair than for the uncorrelated case. The two molecules slightly inhibit each other from experiencing large fields. Although the maximum free energy in the second shell PMF is lower than in the first shell PMF, on average it is closer to the uncorrelated PMF, as expected. These two-dimensional PMF's further support the hypothesis that no significant long range cooperation is required to create a large electric field on a given hydrogen atom.

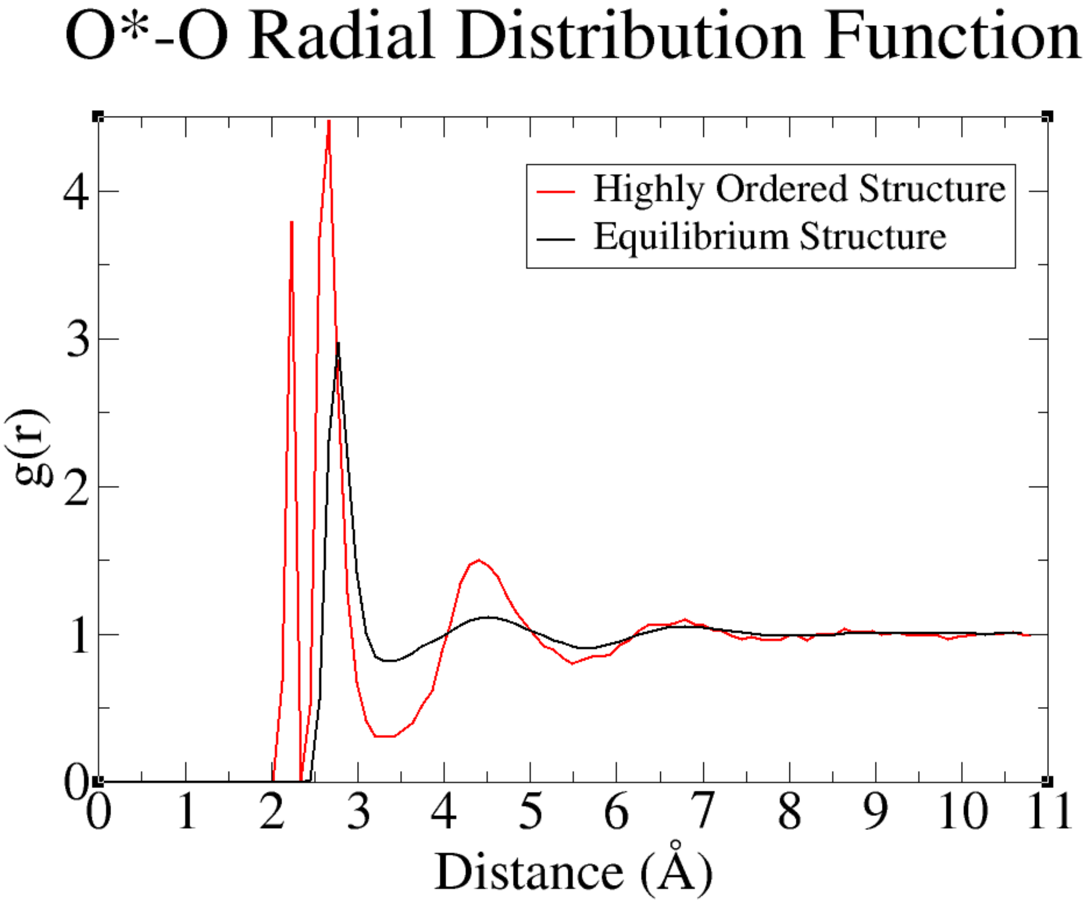

Figure 4.11: Radial distribution function (RDF) of oxygen atoms around the oxgen atom covalently bonded to the tagged hydrogen atom in trajectories with a harmonic bias around electric fields of 25 V/Å. This RDF quantitatively describes the configuration seen in the right panel of Figure 4.8.

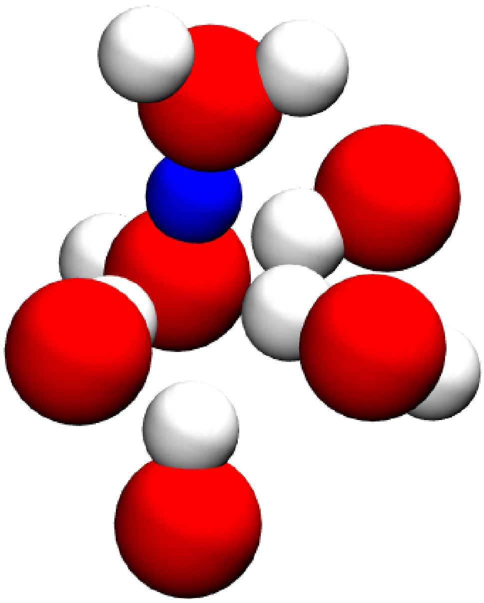

Figure 4.12: Representative configuration from a trajectory with a bias on the electric field on the blue hydrogen atom towards 25 V/Å. In this configuration, the water molecule with the hydrogen atom being pulled away strongly is accepting four hydrogen bonds, a large fluctuation from the equilibrium average of 1.8.

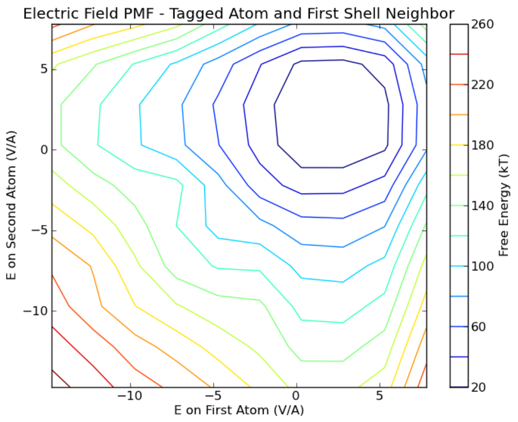

Figure 4.13: The potential of mean force for electric field fluctuations on two hydrogen atoms whose water molecules are within each other's first solvation shell.

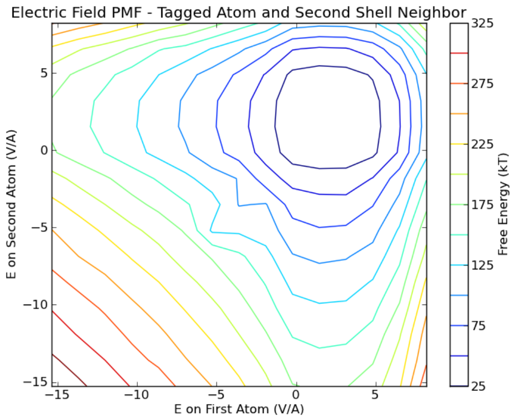

Figure 4.14: The potential of mean force for electric field fluctuations on two hydrogen atoms whose water molecules are within each other's second solvation shell. When a central atom and an atom in the second solvation shell are biased towards large positive fields, a third atom forms a bridge between the two tagged atoms and remains there on the order of 1 ns, giving a path for a proton to diffuse 3 hydrogen bonds away from its initial position by the Grotthuss mechanism.

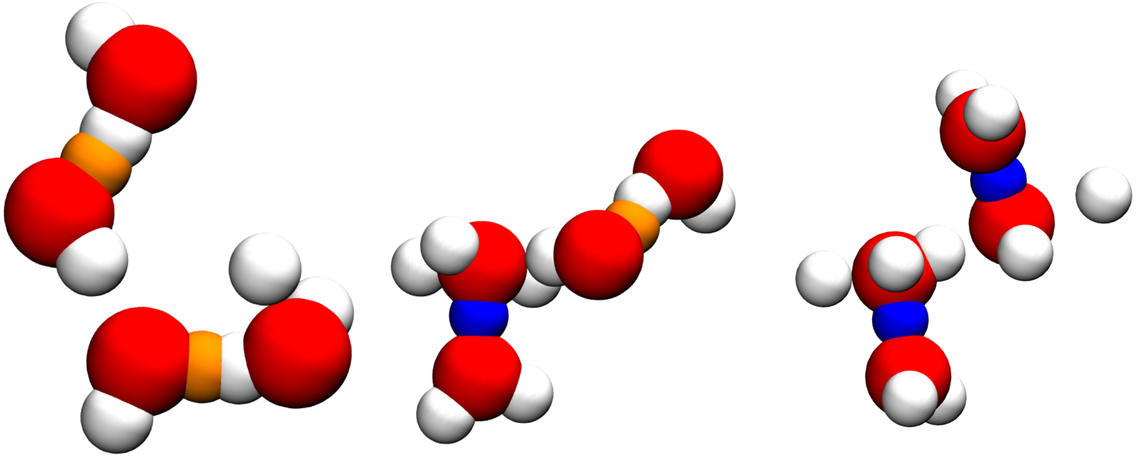

Figure 4.15: Representative configurations from trajectories with two molecules in each other's first shell biased to have large electric fields on one of their hydrogen atoms. Hydrogen atoms in blue are having their covalent OH bond stretched. Hydrogen atoms in orange are having their covalent OH bond compressed.

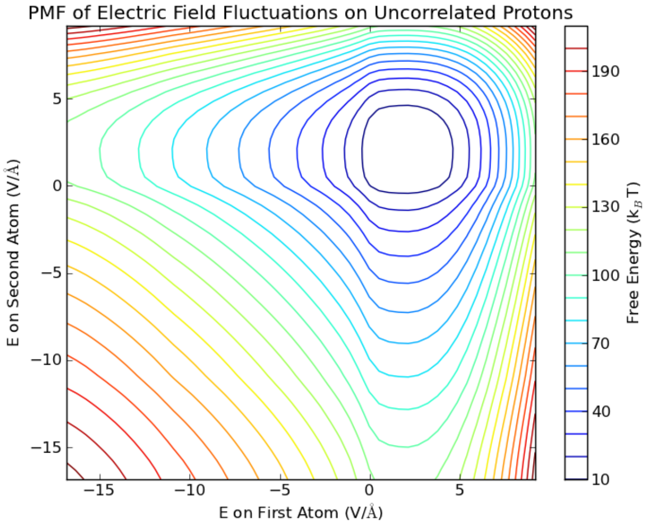

Figure 4.16: This plot shows what the 2D PMF of the electric field of two tagged hydrogen atoms in each other's first solvation shell would look like if their fluctuations were uncorrelated.

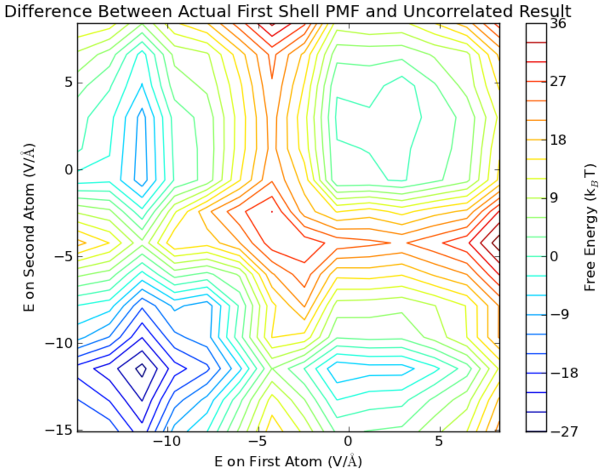

Figure 4.17: The difference between the actual 2D PMF of the electric field of two tagged hydrogens in each other's first solvation shell and the result that would have appeared for uncorrelated two hydrogen atoms individually following the 1D distribution.

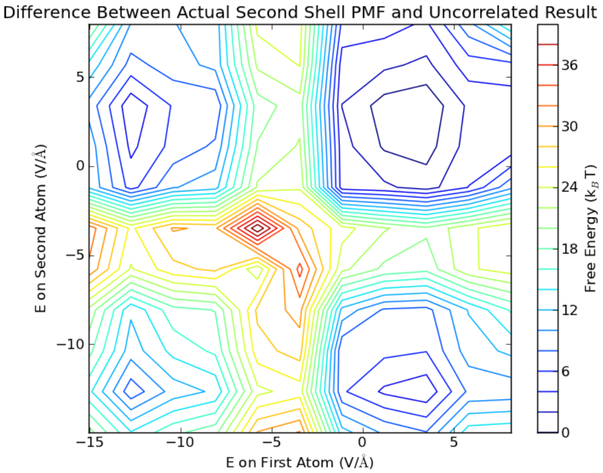

Figure 4.18: The difference between the actual 2D PMF of the electric field of two tagged hydrogens in each other's second solvation shell and the result that would have appeared for uncorrelated two hydrogen atoms individually following the 1D distribution.

4.7 Conclusion

Autoionization is an essential property of water. Understanding how it happens is important both for its direct relevance to water's interaction with materials and biological molecules and because it can serve as an example for any other type of ion dissociation. On a macroscopic scale, autoionization is very common. A glass of distilled water will contain a very large number of hydroxide and hydronium molecules. On a microscopic scale, though, autoionization is very uncommon. In a computationally tractable brute force simulation, it would be impossible to observe many autoionization events. From previous work by Geissler and co-workers [29], we know that autoionization is caused by the fleeting appearance of a very large electric field. In this chapter, I have studied how those large electric fields come about using purely classical simulations. Although such simulations cannot show the actual process of bond breaking, they can show the origin of those large electric fields. Even with a purely classical force field, collecting many samples of the field large enough to induce autoionization in a brute force simulation is intractable. By developing a new fix in LAMMPS, I have used umbrella sampling to collect much better statistics for large fluctuations of the electric field. Applying a bias directly on the electric field allows the system to find the lowest energy configurations which lead to a given electric field. This setup does not build in any information about the hypothesis that the electric field comes primarily for short-ranged interactions. From the resulting trajectories, I have calculated radial distribution functions for configurations conditioned on them having certain values of the electric field.

These calculations confirm the finding from Dellago and co-workers [32] that the contributions to large electric fields are overwhelmingly from short-ranged interactions and add structural detail. The configurations with large electric fields stretching a tagged covalent OH bond have ice-like order with an unusually high number of hydrogen bonds donated to the tagged molecule. As the electric field is pushed to grow larger and larger, the nearest neighbor molecules pull closer and the average number of hydrogen bonds donated to the tagged molecule grows.

Chapter 5

Hydrogen Bonds and Ions

The solvation of ions in water underpins two incompletely explained phenomena of great practical and scientific interest: the Hofmeister series and the attraction of ions to interfaces. I will briefly describe these in turn. The Hofmeister series, first discovered by Franz Hofmeister in 1888, is an ordering of ions based on their propensity to cause proteins to denature - meaning their native structure falls apart, inhibiting their biological function - or remain stable. Protein surfaces are covered in many charged or polar amino acids. The exact interaction between a specific ion and a specific protein depends on the sequence of amino acids in the protein and the structure of the folded protein. Reasonable intuition, then, would suggest that how much more likely a protein is to denature in the presence of a given ion would depend substantially on which protein was being denatured. This intuition is partially correct - the exact value of a protein's solubility in the presence of a given ion does vary by protein - but misses a significant trend: lists of ions ordered by their effect on the solubility of different proteins almost all match. That is, for nearly all proteins the ion CO_3^{2-} significantly enhances stability whereas I^- promotes protein denaturation. Scientists are often engaged in the game of trying to find and explain universal phenomena, those that depend on only a few basic features of an experiment and not the details. In this case, we infer that the molecular mechanism by which specific ions stabilize or denature any protein is independent of the details of the protein. The Hofmeister series has been thoroughly documented experimentally and the same ordering has been found in experiments on enzyme activity, protein crystallization, and bacterial growth [36].

An explanation of the molecular mechanism behind the Hofmeister series independent of the exact details of any given protein has yet to be found. One crucial ingredient must be that ions are strongly attracted to the surface of the proteins. Without a strong interaction, the ions would have little effect on the proteins. Significant progress has been made by others using Poisson-Boltzmann theory, in particular Nadine Schwierz, Dominik Horinek, and Roland Netz. Their approach involves treating a protein as a surface composed of fractions of hydrophobic, positively charged, and negatively charged patches and calculating the approximate interaction of each protein with these idealized surfaces [37] [38]. In this chapter I will take a different approach, namely searching for an order parameter that we can use to develop a Gaussian field theory of ion solvation. I will explain this approach more below.

Another phenomenon closely related to the Hofmeister series is the attraction of ions to liquid-vapor interfaces. Ions of either positive or negative charge are attracted to liquid-vapor interfaces, although the effect is highly asymmetrical. In simulations, an anion can have 1,000 times the surface affinity of a cation [39]. This result originally came as a surprise. If you approximate an ion in water as a point charge in a semi-infinite continuous dielectric medium, you find that its potential goes to infinity at the interface, meaning that the surface charge density would be zero. Experimental results that showed ions are attracted to the interface contradicted somewhat more detailed models that predicted ions would be present at the surface but at lower concentration than in the bulk [40]. This result is likely intimately related to the Hofmeister series: ions are unexpectedly attracted to interfaces between water and anything that isn't water. The presence of ions at the interface lowers the surface tension, which, in the case of proteins, leads them to unfold. In the search for a general mechanism and explanation, we are free to do our experiments and simulations on liquid-vapor interfaces and we can simply check later that the same general mechanism applies to proteins, regardless of their exact structure.

Ultimately, we expect that a successful theory of ion solvation may resemble Lum-Chandler-Weeks (LCW) theory for hydrophobic solvation [41] because an ion is approximately a hydrophobe with the addition of a charge. An excellent introduction to LCW theory is given in David Chandler's review in *Nature* [42]. Hydrophobes show two qualitatively different sets of behavior in water depending on their size. An aqueous mixture of many small hydrophobes will stay well mixed. The presence of a small hydrophobe imposes an entropic cost as the surrounding water molecules are constrained in which directions they can point to form hydrogen bonds. This entropic cost grows with the volume of the hydrophobic cavity. Alternatively, a collection individual hydrophobes has the option of aggregating to form one large cavity. When the cavity is roughly 1 nm in diameter or larger it begins preventing some water molecules from forming any hydrogen bond at all, imposing an enthalpic cost but freeing them from the entropic cost of being forced to point in a particular direction. This enthalpic cost grows with the surface area of the hydrophobic cavity. A collection of large hydrophobes with fixed total volume can lower the surface area of their interfaces with the solvent by aggregating and forming one larger cavity. LCW theory predicts the crossover between these regimes with quantitative accuracy, bringing together two limits that were previously only understood separately and qualitatively [42].

LCW theory is based on the Gaussian statistics of fluctuations in the density of water over short length scales [41]. Consider the molecular density field $\rho(\vec{r})$ with average $\langle \rho(\vec{r}) \rangle \equiv n(\vec{r})$ where the origin can be taken as the center of a tagged water molecule. We can write the free energy of a given configuration of the density field as

$$F_0[\rho(\vec{r})] = \int d\vec{r} \left[w(\rho(\vec{r})) + \tfrac{1}{2} m |\nabla \rho(\vec{r})|^2 \right],$$

where $w(\rho(\vec{r}))$ is a local free energy density and m gives an energy scale to fluctuations in the density. The local free energy is given by

$$w'(n_s(\vec{r})) = m\nabla^2 n_s(\vec{r}) + 2a[\bar{n}(\vec{r}) - n_s(\vec{r})].$$

The quantity $n_s(\vec{r})$ is the slowly varying component of the density field, so the difference in the second term accounts for the rapid fluctuations in density around an interface. The success of LCW relies crucially on the fact that the quantity $\delta\rho(\vec{r}) = \rho(\vec{r}) - n_s(\vec{r})$ is a Gaussian random field. The covariance matrix of the multivariate Gaussian describing the fluctuations in the density is defined by $\chi(\vec{r}, \vec{r}') = \langle \delta\rho(\vec{r})\delta\rho(\vec{r}')\rangle$.

LCW falls short of explaining interactions between ions and water or between proteins and water because it includes no information about charge or polarity. An ion excludes volume, just like a hydrophobe, but its effects are not captured by the density field. An ion or a dipole also affects the orientation of water molecules around it, disturbing the hydrogen bonding network separately from the density. We would like to build on LCW by adding a second order parameter that accounts for these other effects of ions that go beyond density. To maintain the analytical tractability of LCW theory, the new order parameter should be Gaussian. Capturing the salient features of ion solvation with a Gaussian order parameter is a tall order. As an illustration, I'll show one relevant order parameter that fails before moving on to more promising candidates.

5.1 Tetrahedral Order Parameter

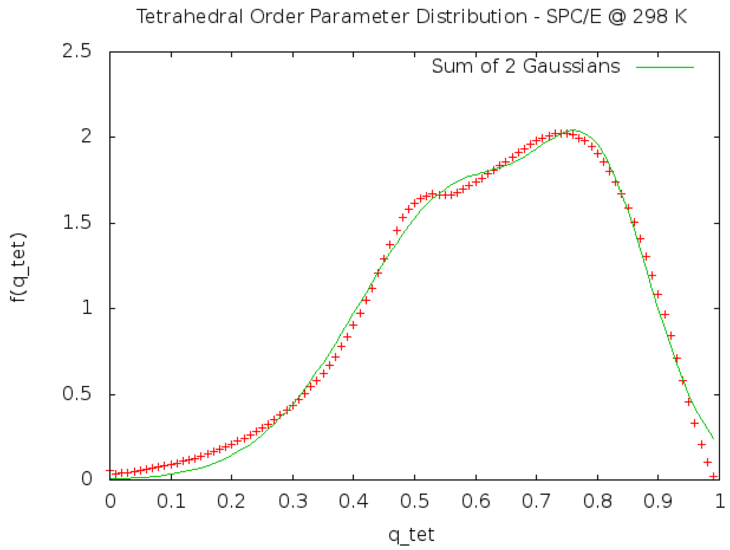

Figure 5.1: The tetrahedral order parameter does not fit a Gaussian. Although it looks bimodal, it also does not fit the sum of two Gaussians. These data come from a simulation of bulk SPC/E water at 298 K.

One order parameter that reveals information about the presence of an ion is the tetrahedral order parameter, q_{tet}, plotted in Figure 5.1. In perfectly ordered ice, the four nearest neighbors of each molecule form a tetrahedron. In liquid water, such configurations will emerge frequently but fleetingly. The distribution of the tetrahedral order parameter allows us to measure to what extent a system follows tetrahedral ordering. It is defined as

$$q_{\text{tet}} = 1 - \tfrac{3}{8}\sum_{i}^{4}\sum_{j=i+1}^{4}\left(\cos\theta_{ij} + \tfrac{1}{3}\right)^2,$$

where θ_{ij} is the angle between the vectors pointing from the central molecule to the i^{th} and the j^{th} nearest neighbors. The average value for an ideal gas is 0. The average value for a perfectly tetrahedral system is 1 [43] [44]. A reasonable guess for the distribution of q_{tet} in liquid water would be that it is a Gaussian centered around some value significantly below 1. Alternatively, there could be two basic states of water, one with tetrahedral order and one without tetrahedral order. In that case, the distribution might be the sum of two Gaussians, one with low and one with high values of q_{tet}. Both guesses, though, are wrong. Bounded between -1/3 and 1 but with a wide variance, the data do not remotely fit either a single Gaussian or the sum of two Gaussians. Although the tetrahedral order parameter is very useful for measuring the amount of order in a system, its poor fit to a Gaussian suggests that even a locally defined variant will not be useful for building a Guassian field theory of ion solvation.

Our intuition is that the presence of an ion or dipole distorts the hydrogen bonding network largely by causing rotations, not necessarily by forcing translation of water molecules. The tetrahedral order parameter will miss that information because it is based entirely on the locations of water molecules - typically defined as their oxygen centers - and not their orientations. The angle θ_{ij} in the tetrahedral order parameter does not describe the orientation of individual water molecules - the positions of the hydrogen bonds do not appear in the calculation of q_{tet} - but instead the angle between sets of three water molecules. A more promising route may be the density field of hydrogen bonds.

5.2 Fluctuations of the Hydrogen Bond Density Field

To begin, we need a precise definition of a hydrogen bond. I will follow the example of Alenka Luzar and David Chandler and define two water molecules as hydrogen bonded if and only if the distance between their oxygen centers is less than 3.5 Å and the O-H...O angle is less than 30° [45]. The precise definition of a hydrogen bond used is not important; small changes in either the angle or distance cutoff will not make a qualitative difference. Although some people have speculated that hydrogen bonds may have two states, that speculation has been demonstrated to be false. If there were two states of hydrogen bonds, we would have to worry that a small change in our bond definition may exclude one of the states and qualitatively change our results. Fortunately, the statistics of hydrogen bonds are well explained by fluctuations around a single state [46].

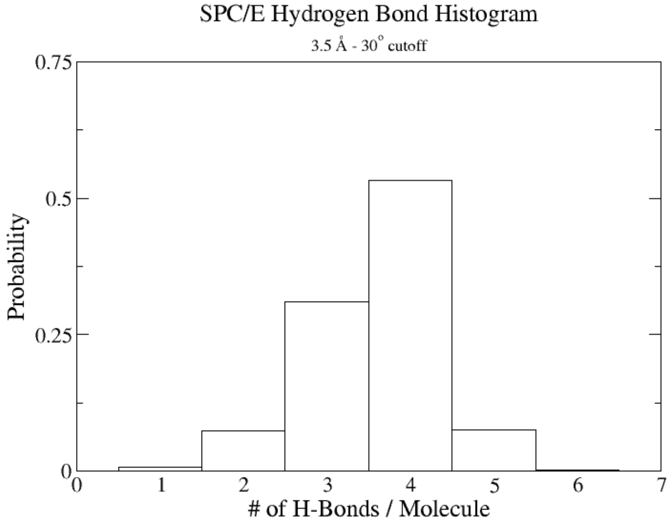

Figure 5.2: The distribution of the number of hydrogen bonds to a single water molecule. This distribution is found in a simulation of SPC/E water where a hydrogen bond is defined as in the main text. The average number of water molecules per molecule is 3.60 with a non-Gaussian distribution.

With hydrogen bonds precisely defined, the simplest measurement is the distribution of the number of hydrogen bonds on a single water molecule. See Figure 5.2. The distribution is, again, non-Gaussian with an average of 3.6 bonds per molecule. Rather than building a theory around the number of water molecules on a given molecule, though, we would like to use the density field of hydrogen bonds in space. Instead of tracking individual water molecules, I will track cubic probe volumes with a volume chosen to give an average occupation of one molecule per cell. To attribute a hydrogen bond to a particular probe volume, I define its position as the midpoint between the oxygen atoms of the water molecules donating and receiving the hydrogen bond. See Figure 5.3. By construction, the average of this distribution is the same as the average of the distribution for single water molecules. However, the distribution is now Gaussian.

The nearest neighbors of a tagged molecule are not independent of each other, so the number of hydrogen bonds on a single molecule is not the sum of independent random variables, giving us no reason to expect it to be Gaussian. It may be surprising that the hydrogen bond density fluctuations in a cell with only one molecule could be Gaussian if the distribution for a fixed molecule is highly non-Gaussian. This result is consistent with the density fluctuations of water molecules themselves, which are Gaussian down to surprisingly small probe volumes - even probe volumes that hold less than one molecule on average [47]. I demonstrate this effect in both spherical and cubic probe volumes in Figure 5.4.

A key feature of ion solvation is the asymmetry between anions and cations. Flipping the

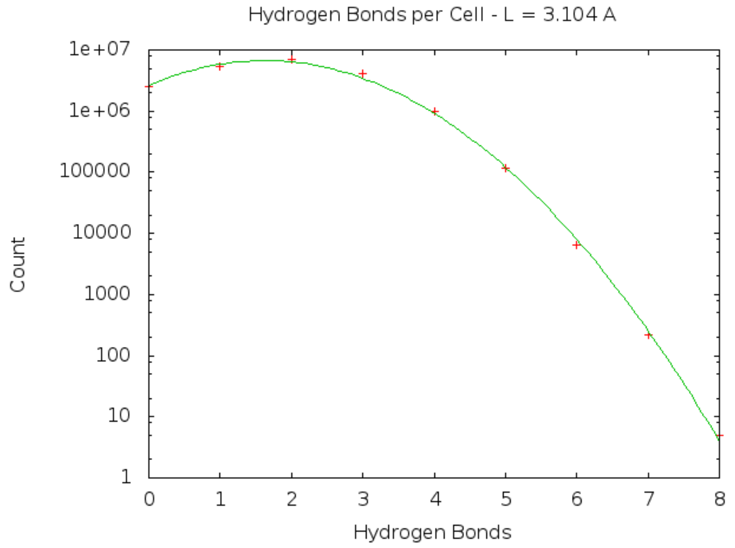

Figure 5.3: The distribution of the number of hydrogen bonds in a cell with sidelength $L = 3.104$ Å, which on average holds 1 water molecule.

sign of the charge on an ion in a simulation may have a very subtle effect on the density field of hydrogen bonds. Separately tracking the density field of donated and accepted hydrogen bonds would show a very dramatic change induced by swapping the test charge. A cation will act as a source of hydrogen bonds and an anion will act as a net sink of hydrogen bonds. If I define the position of a donated hydrogen bond as the oxygen center of the donating molecule and the position of an accepted hydrogen bond as the oxygen center of the accepting molecule, I can track density fluctuations in probe volumes as before. See Figure 5.5. Across several probe volumes, the distributions for donated and accepted hydrogen bonds are highly Gaussian.

To confirm that the density field of hydrogen bonds has Gaussian spatial correlations - just like the water molecules themselves - I can check that the Fourier modes are Gaussian. Given a vector of hydrogen bond positions, $\vec{X} = \{\vec{x}_1, ..., \vec{x}_N\}$, the Fourier transform is given by $\hat{\rho}_{\mathrm{HB}}(k_n) = \sum_{j=1}^{N} \mathrm{Re}\left[\exp[ik_n\hat{x} \cdot \vec{x}_j]\right]$, where $k_n = \frac{2\pi n}{L}$. I can get a total of six samples from a single configuration by taking the imaginary component, $\hat{\rho}_{\mathrm{HB}}(k_n) = \sum_{j=1}^{N} \mathrm{Im}\left[\exp[ik_n\hat{x} \cdot \vec{x}_j]\right]$ and by replacing $\hat{x}$ with $\hat{y}$ and $\hat{z}$. For each wavevector, k_n, I collect many samples of $\hat{\rho}(k_n)$ over a trajectory. If density fluctuations at different points in space follow a multivariate Gaussian distribution, then the distribution of $\hat{\rho}(k_n)$ should also be Gaussian. As you can see in Figure 5.6, they are, in fact, highly Gaussian. The density field of hydrogen bonds should therefore be an acceptable order parameter for a Gaussian field theory of ion solvation.

Now that we're sure the fluctuations of the hydrogen bond density are Gaussian, it would

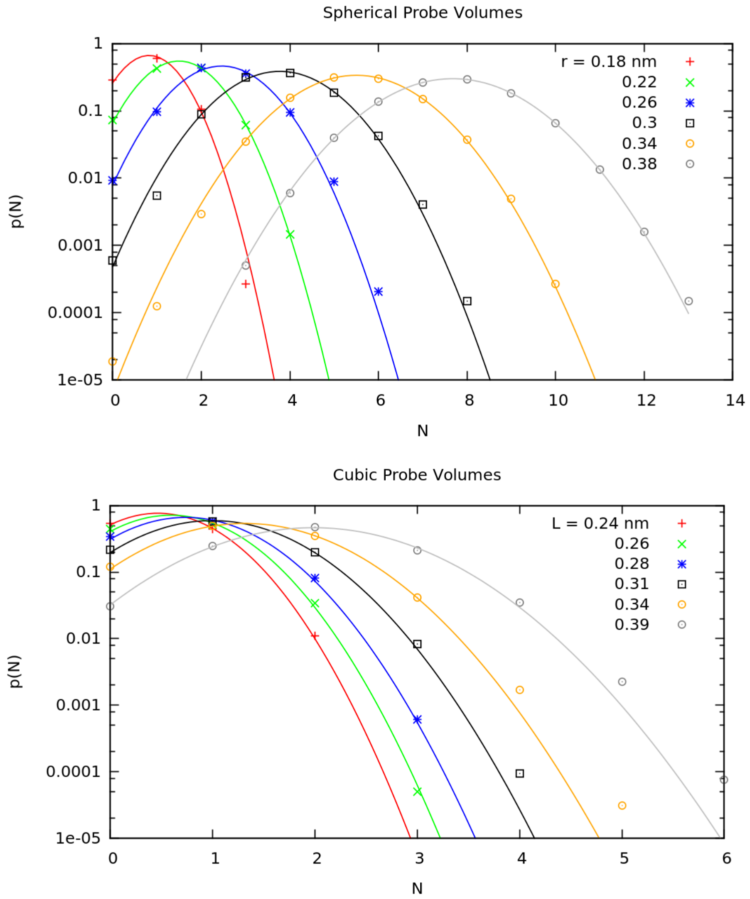

Figure 5.4: Fluctuations in the number of water molecules in a spherical probe volume (top) or a cubic probe volume (bottom). Lines are Gaussians with the same mean and variance as the data. Both cubic and spherical probe volume distributions are Gaussian to a high degree even in small volumes that hold an average of less than one molecule.

be helpful to know what the mean of that Gaussian is. Those mean fluctuations are given by the radial distribution function, seen in Figure 5.7. Hydrogen bonds effectively have an excluded volume a little over 1 Å and show a separation between a first and second solvation shell around 3 Å.

For a Gaussian field theory based on hydrogen bonds to predict charge asymmetry, the radial distribution function of hydrogen bonds around cations and anions should be meaningfully different. In fact, they are; see Figure 5.8. The density of hydrogen bonds shows a narrower, taller peak around anions than around cations. Hydrogen atoms are apparently more tightly constrained around an anion than around a cation. This basic physics will reappear below when I show the behavior of hydrogen bond orientations around an ion instead of hydrogen bond positions.

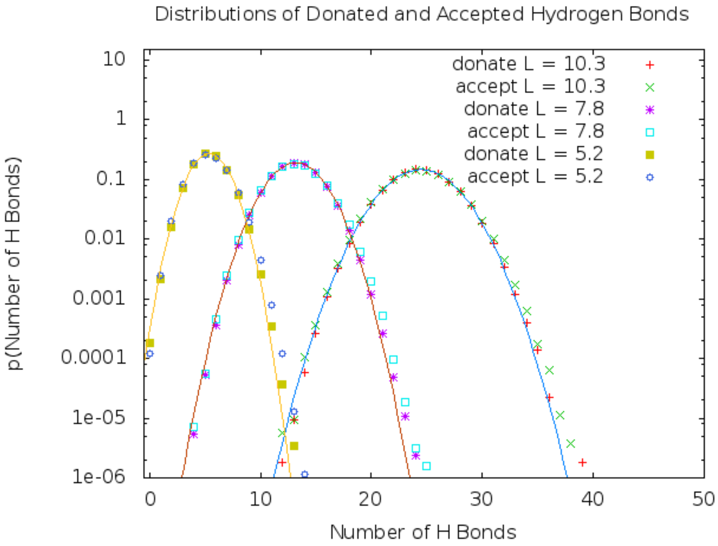

Figure 5.5: Donated and accepted hydrogen bond density fluctuations in three sizes of cubic probe volumes. Lines are Gaussians with the same mean and variance as the data.

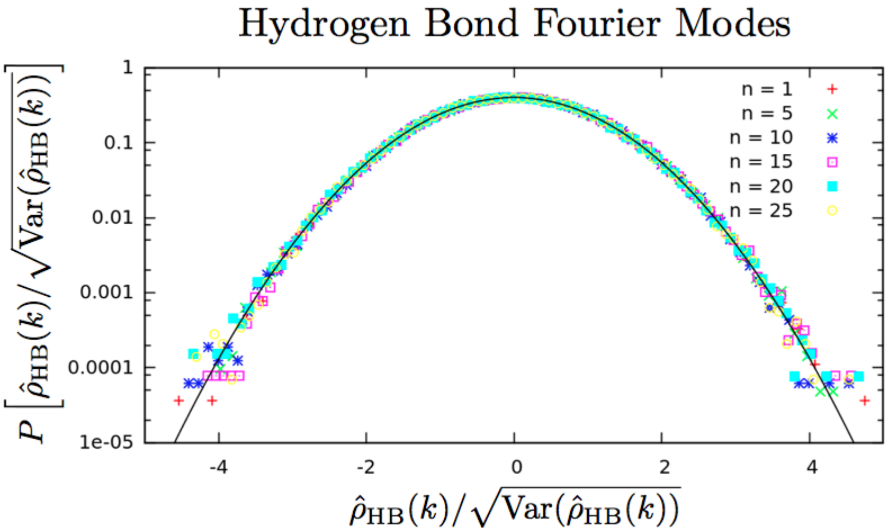

Figure 5.6: Distribution of the normalized Fourier modes of the hydrogen bond density field. Each individual sample of $\hat{\rho}_{\mathrm{HB}}(k_n)$ has been divided by $\sqrt{\mathrm{Var}(\hat{\rho}_{\mathrm{HB}}(k_n))}$. The solid line is a Gaussian of mean 0 and variance 1.

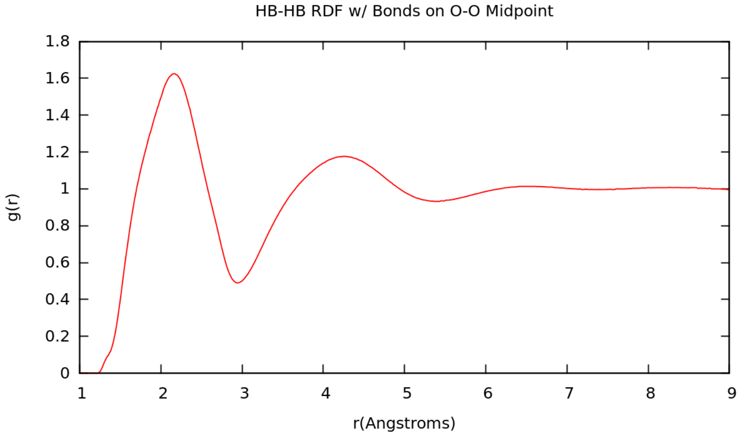

Figure 5.7: Radial distribution functions of hydrogen bonds located at the midpoint between the Oxygen centers of the donating and accepting molecule..

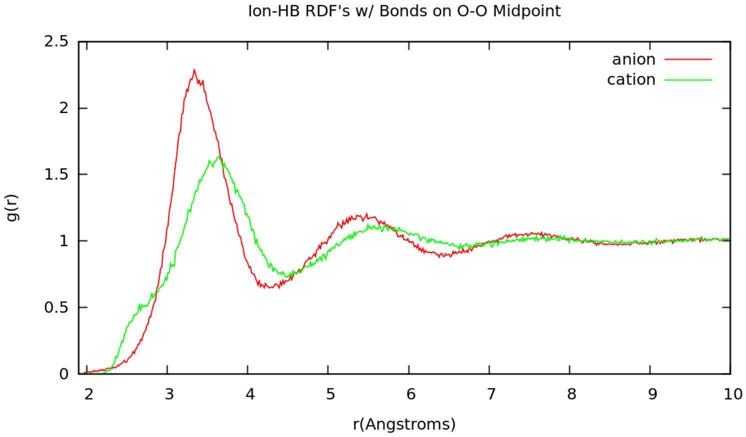

Figure 5.8: Radial distribution functions of hydrogen bonds located at the midpoint between the Oxygen centers of the donating and accepting molecule around a central ion.

5.3 Structure of Hydrogen Bond Network Around an Ion

Besides studying the positions of hydrogen bonds, we can also investigate the orientations of hydrogen bonds around a central molecule, either a water molecule or an ion. A charged molecule will have many more hydrogen bonds in its first solvation shell than a hydrophobic molecule, but looking only at the locations of hydrogen bonds will not show a stark difference between cations and anions. Studying how hydrogen bond orientations change in response to a cation or anion should give us greater physical insight into the molecular underpinnings of charge solvation asymmetry.

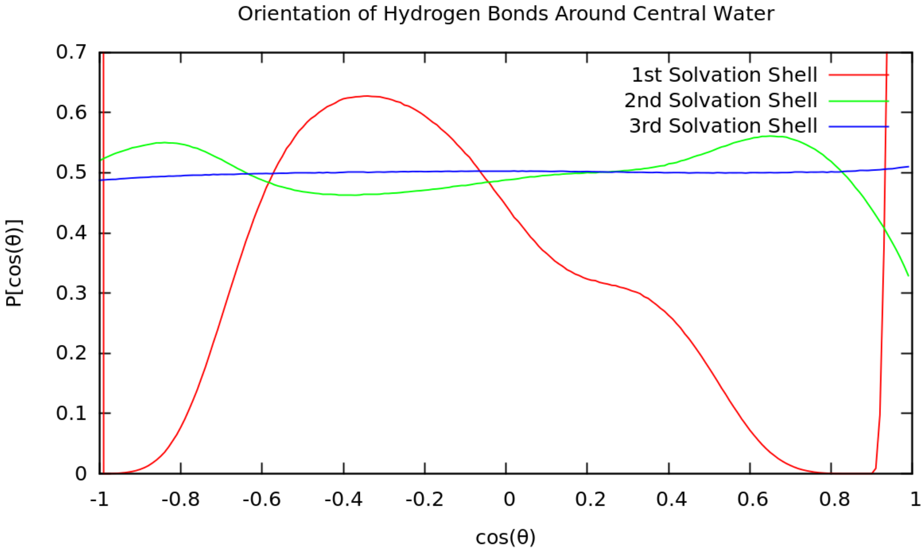

Figure 5.9: Statistics of hydrogen bonds between water molecules. Each bond is tagged with a unit vector located at the O-O midpoint pointing from the donating hydrogen atom to the accepting oxygen atom.

For reference, let's look at the orientations of hydrogen bonds around a water molecule. See Figure 5.9. I measure the angle, θ, between the radial vector pointing away from the central molecule and the vector connecting the hydrogen atom donating a bond to the oxygen atom accepting a bond. If there were no correlations between these vectors, the probability distribution plotted versus $\cos(\theta)$ would be flat. When we restrict ourselves to looking at hydrogen bonds located in the third solvation shell, a flat distribution is what we see. That tells us that hydrogen bonds in a water molecule's third solvation shell are nearly unaffected by its presence. Hydrogen bonds in the second solvation shell show only a slight preference for pointing towards or away from the central molecule.

In the first solvation shell, we see large peaks at -1 and 1, corresponding to bonds directly to or from the tagged molecule. The peak at -1 is a delta function with an integral of roughly 1.8, corresponding to accepting 1.8 water molecules. The integral of the peak on the right must then also be 1.8 since I am averaging over a box of water molecules and each accepted water molecule corresponds to one donated water molecule. The peak on the right is not a delta function because I am defining the direction of the hydrogen bond based on the donating hydrogen atom's position and not the position of the donating molecule's oxygen.

We see a wider peak around $\cos(\theta) \approx -0.35$ with a distinct shoulder around $\cos(\theta) \approx +0.35$. These middle peaks correspond to bonds between molecules in or near the first solvation shell. Their bonds are mostly perpendicular to the radial vector with a bias to pointing more in towards the central molecule.

In comparison, let's look at the orientations of hydrogen bonds around a cation or an anion in Figure 5.10. In the third solvation shell, the almost perfectly horizontal distribution is replaced with a roughly linear distribution sloping slightly down for an anion and up for a cation. This matches the biasing that would be expected for hydrogen atoms in the presence of an ion, but even a mere 5.68 Å separation between the ion and the hydrogen bond already screens most of the influence of the ion. In the second solvation shell we see a subtle suppression of hydrogen bonds pointing towards a cation, but around an anion we see little suppression of hydrogen bonds pointing away, towards the outer solvation shells.

In the first solvation shell, a cation would show a large peak at +1 and an anion would show a large peak at -1 due to bonds directly to or from the ion. I will exclude those delta function peaks and instead look only at hydrogen bonds between water molecules. The anion shows almost complete suppression of nearby hydrogen bonds pointing away from the center. Nearly all of the water molecules in the first solvation shell are forced to point towards the anion. In contrast, the cation still allows for quite a few hydrogen bonds pointing not directly towards it but at least more towards it than away from it. The much greater similarity of the hydrogen bonding network in pure water to the cationic case than the anionic case gives us some insight into charge solvation asymmetry. Anions disturb the hydrogen bonding network substantially more than cations, imposing a much larger entropic cost, suggesting an explanation for their stronger surface affinity.

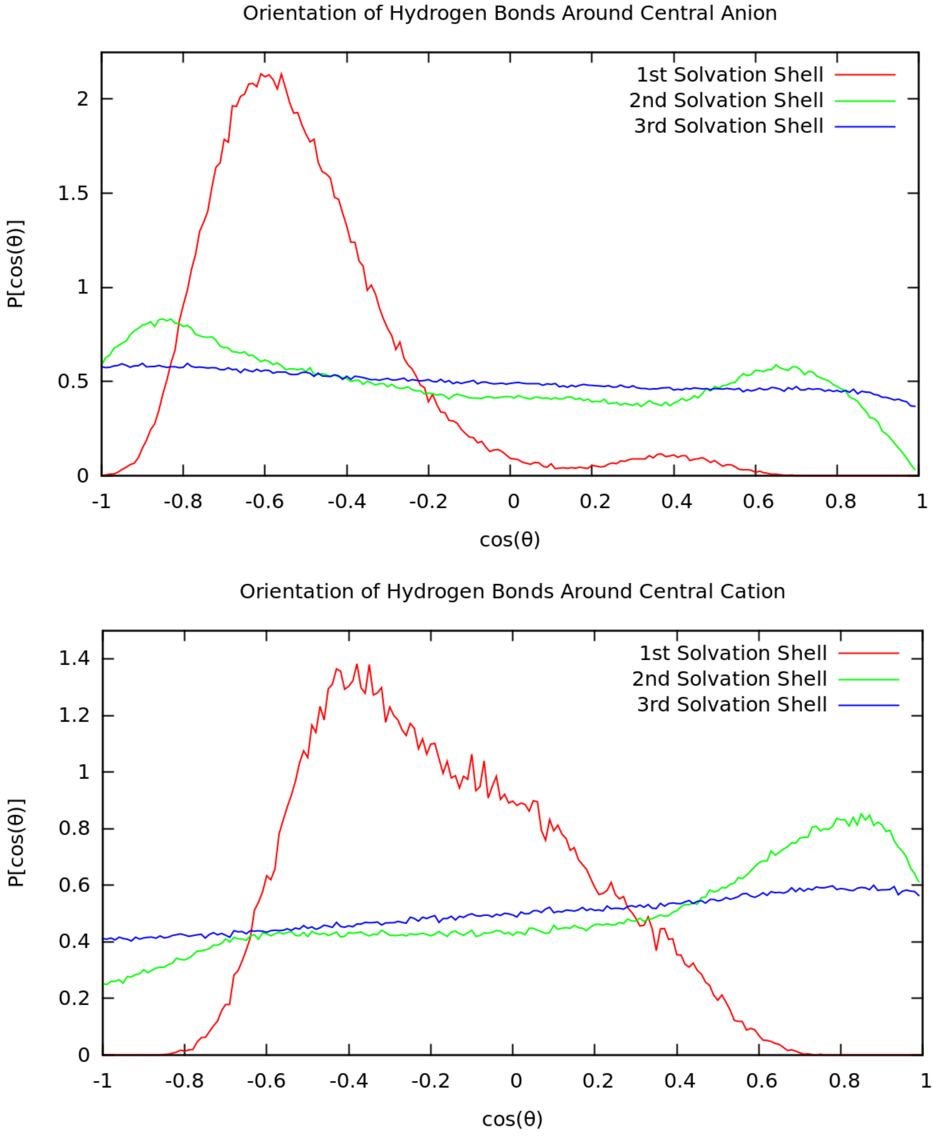

Figure 5.10: Statistics of hydrogen bonds tagged with unit vectors. The vector is located at the ion-O midpoint. For water molecules, the vector is located at the O-O midpoint. The vector points from the donating molecule to the accepting molecule.

5.4 Conclusion

A Gaussian field theory of ion solvation requires, by definition, a Gaussian order parameter that captures the effect of an ion on the surrounding solvent. A promising candidate is the density field of hydrogen bonds. The Fourier modes of the hydrogen bond density fit very well to a Gaussian, proving that the fluctuations at different points in space will fall on a multivariate Gaussian. One shortcoming of the density field of hydrogen bonds is that the response to an anion or cation is very similar. As an alternative, using two fields to track the positions of hydrogen bond donors and hydrogen bond acceptors could more clearly differentiate between the presence of an anion or cation at the cost of an additional parameter. The work in this chapter does not provide the necessary field theory but provides order parameters that could be used to develop a Gaussian field theory following the path outlined in David Chandler's 1993 *Physical Review E* paper [48].

Chapter 6

The Surface Potential of Water Varies with the Presence of an Ion

An ion can be approximated as a volume excluding cavity with a charge. Its solvation free energy is the free energy required to make room for the cavity and the free energy to add a charge to that cavity. Since the solvation energy of hydrophobes is already well understood thanks to Lum-Chandler-Weeks (LCW) theory [41], we draw on it for significant inspiration. After the first-order effect on dipoles in the surrounding solvent, the most important second-order feature of ion solvation is its asymmetry. Anions are attracted far more strongly to the air-water interface than cations of identical size are, an effect not captured by dielectric continuum theories (DCT) that depend solely on solvent polarization. Such a DCT assumes that the electric potential at the center of a cavity is Gaussian and that the free energy to solvate an ion is linear in the magnitude of the charge, which in part necessitates that the free energy of solvation is independent of the sign of the charge.

The most popular extension of DCT that attempts to account for charge asymmetry adds in a linear response to polarization at the air-water interface [49] [50] [51] [52]. This version of DCT, which I will refer to as DCT with an interface, calculates the cost of solvating a point charge in a dielectric medium near a polarized interface that is assumed to respond linearly to the presence of an anion or cation. Although it is a very clean hypothesis, it is, unfortunately, empirically wrong. Instead, the polarization at the interface responds non-linearly to the presence of an ion and its response varies with the diameter of the ion, an effect outside the scope of a model based on point charges.

6.1 Predictions of Dielectric Continuum Theory with an Interface

A solution with an excess charge on its surface would, of course, have an electrostatic interaction with any ions brought from the vapor phase into the liquid phase. Although a solution of pure water has no net charge, orientational bias near the interface results in an excess of hydrogen atoms pointing towards the vapor phase. A thin slice of roughly 2 Å at the interface thus shows an excess of positive charge, offset by an excess of negative charge slightly farther into the liquid phase [39]. This layer of positive charge could reasonably be expected to attract anions and repel cations. If the solvation energy of an ion could be decomposed as the sum of its cavitation energy and an electrochemical interaction with the polarization at the liquid-vapor interface, the surface interaction could account for charge asymmetry at the interface.

However, something must be amiss with this explanation. Orientational bias is not necessary to create a layer of excess charge at the interface. Consider a spherically symmetrical molecule with point charge -q at its center and charge +q spread uniformly over a sphere of radius r. Such a molecule, being spherically symmetrical, cannot have orientational bias and, being electrically neutral, cannot have electrostatic interactions with finite ions that do not penetrate its outer shell. And yet, such a molecule arranged on a cubic lattice would have an excess of charge over a layer with thickness less than r at an interface with vapor. The argument in the previous paragraph suggests that finite ions would have charge-asymmetrical surface adsorption in a solvent of this molecule, even though the ions have no electrostatic interaction with the solvent at all.

To get to the root of this contradiction, we need to make this explanation more precise. Define the surface potential as the difference between the average electrostatic potentials in liquid and vapor [39]:

$$\phi_{\mathrm{sp}} = 4\pi \textstyle\int_{z_v}^{z_l} z\rho(z)dz,$$

where z is the axis perpendicular to the interface, $\rho(z)$ is the average charge density at height z, and z_v and z_l are heights in the vapor and liquid phases. In simulation, the surface potential gives rise to an energy penalty of 50 kJ/mol to bring a point charge of -1 from the vapor into the liquid, and the opposite benefit for a charge of +1 [39]. As a naive approximation to the solvation free energy at the interface, one could add this energy to the free energy of creating a cavity, giving rise to a substantial asymmetry favoring anions at the interface more strongly than cations.

I will test this hypothesis by measuring the electric potential inside cavities with varying charges and at varying distances from an air-water interface. The charging free energy can be calculated from the full electric potential distribution as

$$F(q;z) = -k_B T \ln \int d\phi P_0(\phi;\sigma,z)\exp[-\beta q\phi],$$

where $P_0(\phi;\sigma,z)$ denotes the probability distribution of ϕ in a cavity with diameter σ, 0 charge, and a distance z from the interface.

DCT with an interface makes several predictions about solvation energies besides that there is charge asymmetry. First, because it is based on the interaction between a point charge and the surface, it predicts that the charging free energy is independent of the diameter of the ion. As seen in Figure 6.1, the distribution of ϕ varies considerably as a function of

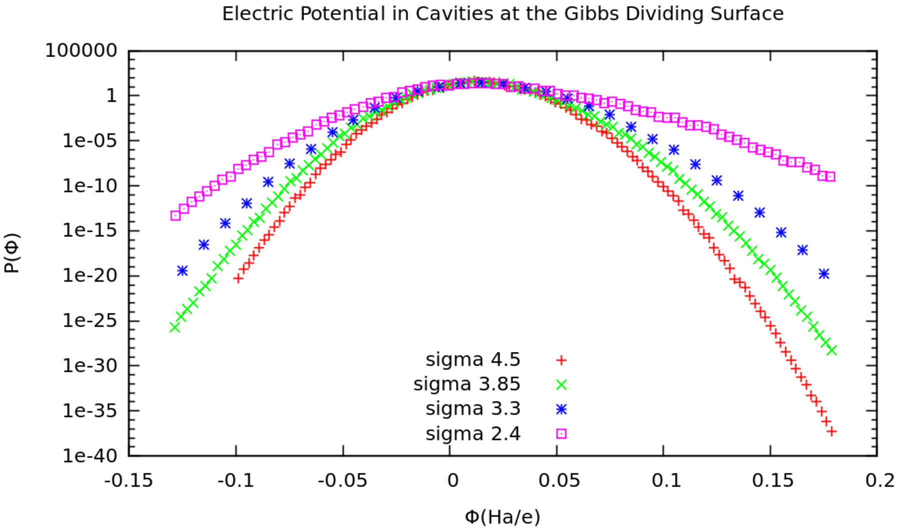

Figure 6.1: Comparison of the electric potential distributions measured at an air-water interface, where z=13.4 Å is the Gibbs Dividing Surface.

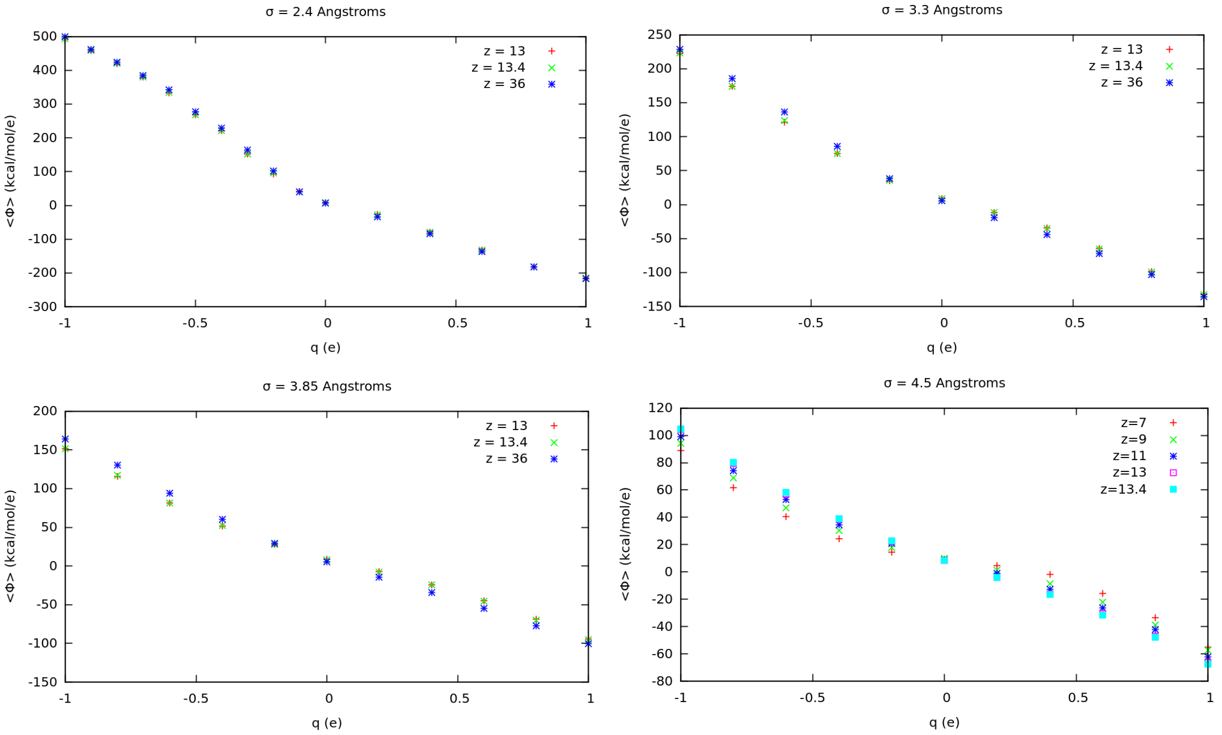

Figure 6.2: Average values of the electric potential, ϕ, for different cavity sizes and heights. The Gibbs Dividing Surface is at $z = 13.4$ Å, with bulk water above and vapor below.

cavity size. Second, the fundamental assumption of linear response theory is that $P_0(\phi;\sigma,z)$ can be approximated as a Gaussian, which gives

$$F(q;\sigma,z) \approx q\langle\phi(\sigma,z)\rangle - \frac{1}{2}q^2\langle(\delta\phi(\sigma,z))^2\rangle.$$

Asymmetry in the charging free energy is then

$$F(q;\sigma,z) - F(-q;\sigma,z) \approx 2q\langle\phi(\sigma,z)\rangle.$$

To check whether this approximation holds, I will calculate

$$\psi(q;\sigma,z) = \frac{F(q;\sigma,z) - F(-q;\sigma,z)}{2q}.$$

The fact that ion solvation is charge asymmetrical requires that $\psi(q;\sigma,z)$ is non-zero, but nothing more. DCT with an interface is essentially a minimal model built on this requirement, predicting that $\psi(q;\sigma,z)$ is non-zero but constant at $\langle\phi(\sigma,z)\rangle$. However, if ϕ were perfectly Gaussian, $\langle\phi(\sigma,z)\rangle_q$, where the subscript denotes the presence of a charge q at the center of the cavity, would be perfectly linear. As seen in Figure 6.2, that is not the case. Below, I will explore how meaningful that non-linearity is.

Finally, I will compare the asymmetry at the interface to the asymmetry in bulk via the quantity

$$\Delta\psi(q) = \psi(q,z_i) - \psi(q,z_b),$$

where z_i is the height of the Gibbs dividing surface and z_b is a point inside the bulk. The third prediction of DCT with an interface is that $\Delta\psi(q)$ is constant for a given cavity size. All three of these predictions are wrong.

6.2 How to Measure Surface Interactions

The surface potential is the dipole moment of the charge density in between z_b and z_l, so an intuitive explanation of the source of the surface potential is that it is the sum of the dipole moments of the molecules in the same region. However, a solvent composed of molecules with a spherically symmetrical charge distribution (and therefore no dipole moment) can give rise to a non-zero surface potential as well [39] [53]. This result even holds for solvent molecules that are both spherically symmetrical and net neutral and would therefore have no charge interaction with ions that can't penetrate into the solvent molecules. The surface potential explanation then makes the nonsensical prediction that there is charge solvation asymmetry in solvents that have no charge interactions.

These problems only prove that the surface potential is sensitive to model details that leave charging free energies unaffected, not that the surface potential is sensitive to the presence of an ion. We can put DCT with an interface on a more physically sound basis by

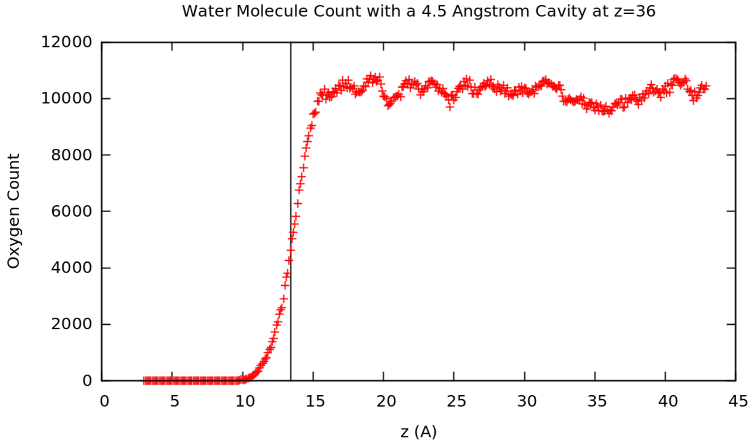

Figure 6.3: Counts of oxygen atoms as a function of height over the entire course of a simulation. The vertical line denotes the Gibbs dividing surface, which is used as the interfacial height, z_i.

testing ψ and $\Delta\psi$ as defined above. By taking the difference in charging free energies at the interface and in the bulk, we can calculate the cost of bringing a charge into the bulk. If that cost is different for cations and anions of equal charge magnitude, we have recovered charge asymmetry.

To measure the potential in a cavity at the interface, I need to define exactly where the interface is. First, I measure the density profile of water molecules in a simulation with a cavity very far from the interface. Using that density profile, I calculate the Gibbs dividing surface, z_{GDS}:

$$\int_0^{z_{\mathrm{GDS}}} \rho(z)dz = \int_{z_{\mathrm{GDS}}}^{z_{\mathrm{bulk}}} (\rho(z) - \rho_{\mathrm{bulk}})dz,$$

where ρ_{bulk} is the average density in the bulk. You can see a demonstration of this calculation for one simulation cell in Figure 6.3.

Even measuring the electrostatic potential correctly in a periodic simulation has pitfalls, in particular a subtle dependence on the size and shape of the simulation cell. In a periodic bulk simulation without the presence of vapor, the potential at the center of a cavity with a test charge will have a non-physical contribution from periodic images of the charge and from the neutralizing background charge. In the thermodynamic limit, there would be no periodic images or neutralizing background charge. To correct for these simulation artifacts, I calculate the potential on a test charge in an empty box with identical dimensions and subtract that value from potentials calculated in a simulation with the same test charge. With these corrections, the electric potentials at the center of cavities in any size or shape of periodically replicated cell collapse onto each other; see figure 6.4.

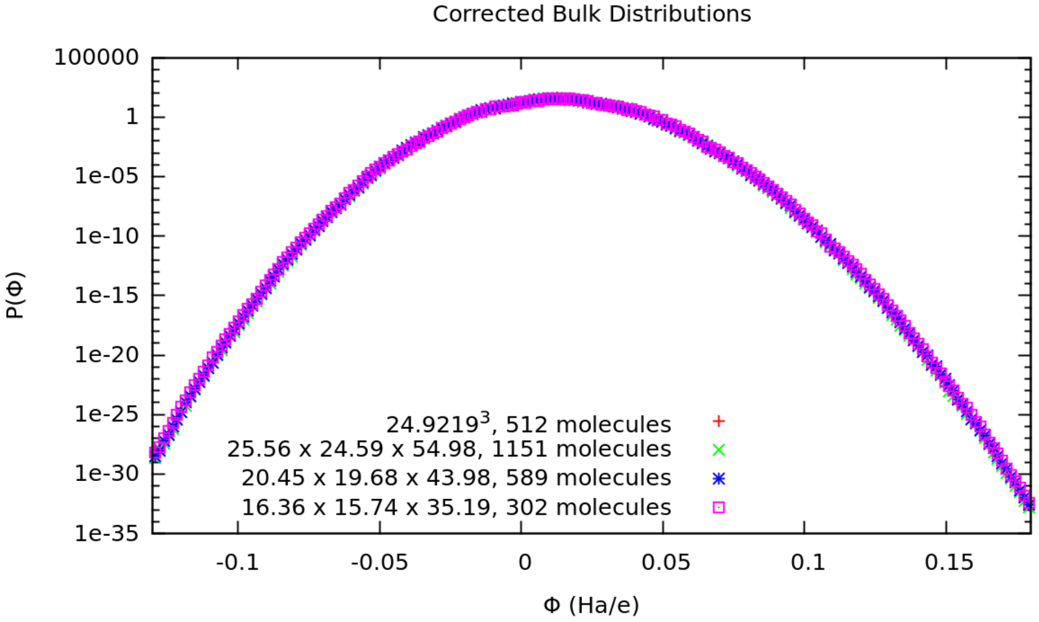

Figure 6.4: Electric potential distributions in the center of a cavity from simulations of bulk water with three sizes of cuboid simulation cells and one cubic cell all collapse onto each other with the corrections discussed in the text.

In a periodic simulation with an air-water interface, the finite size corrections are more complicated. Now, not only will the test charge have non-physical interactions with its periodic images and neutralizing background charge, the polarized surface will also interact with periodic images of itself. In the thermodynamic limit, there is only one interface. Fortunately, the interactions between the periodically replicated interfaces are independent of the sign of the charge and therefore cancel in ψ. I am unable to make rigorous statements about absolute values of $\langle\phi\rangle$ without a rigorous finite size correction (approximate expressions for such a correction is an active area of research in the group). Nevertheless, I am able to make rigorously correct statements about ψ, which is sufficient to prove that the surface potential depends on the presence of an ion. By defining $F_{\text{liq-vap}}(q;\sigma,z)$ as the charging free energy defined previously adjusted by a finite size correction, I end up with exactly the same $\psi(q;\sigma,z)$ as in the bulk case:

$$
\begin{aligned}
F_{\text{liq-vap}}(q;\sigma,z) &= F(q;\sigma,z) - F_{\text{FS}}(q;\sigma,z) \\
\psi(q,z) &= \frac{F(q;\sigma,z) - F_{\text{FS}}(|q|;\sigma,z) - (F(q;\sigma,z) - F_{\text{FS}}(|q|;\sigma,z))}{2q} \\
&= \frac{F(q;\sigma,z) - F(q;\sigma,z)}{2q}
\end{aligned}
$$

As explained above, DCT with an interface predicts that for a given cavity size and

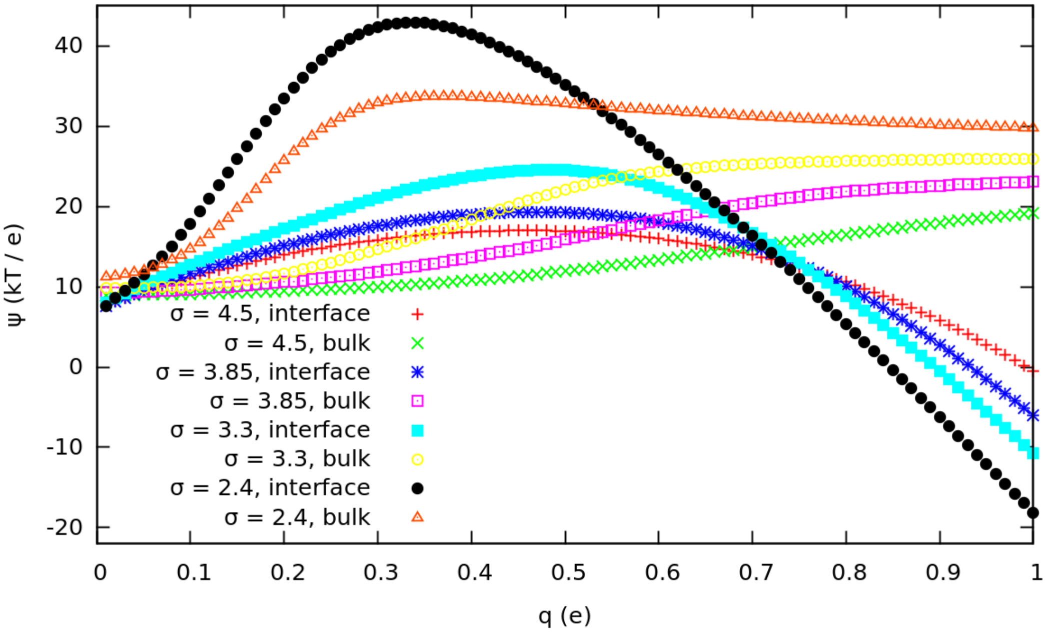

Figure 6.5: $\psi(q;\sigma,z) = \frac{F(q;\sigma,z)-F(-q;\sigma,z)}{2q}$ calculated for four sizes of cavities either at the interface - defined as the Gibbs dividing surface - or in the bulk. DCT with an interface predicts that each of these functions should be a constant.

distance into the bulk, this difference should be constant in the charge at the center of the cavity. If the polarization at the surface changes in response to the charge on the ion, ψ will not be constant. Figure 6.5 shows that ψ is not only not constant, but not even monotonic. DCT with an interface gives predictions that are both quantitatively and qualitatively incorrect.

Finally, I come to the quantity we care about most, $\Delta\psi(q;\sigma)$. This quantity is the important one because it determines the relative density of cations and anions at the interface. If each charge of ion has the same density in the bulk, then the ratio of their densities at the interface is

$$\frac{\rho(q;z_i)}{\rho(-q;z_i)} = \exp[-2q\Delta\psi(q)].$$

So the relative densities of cations and anions at the interface depends exponentially on $\Delta\psi$. Just like it does for ψ, DCT with an interface predicts that $\Delta\psi$ is constant in q. Again, as seen in Figure 6.6, not only is $\Delta\psi$ not constant, it is non-monotonic and changes sign.

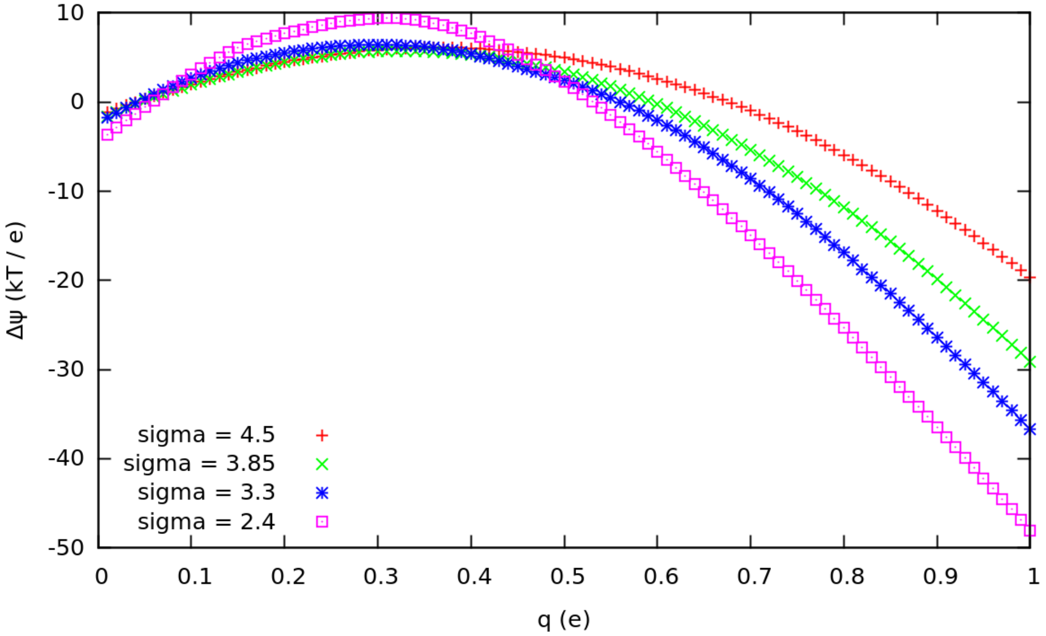

Figure 6.6: $\Delta\psi(q;\sigma) = \psi(q;\sigma,z_i) - \psi(q;\sigma,z_b)$ for four cavity sizes.

6.3 Conclusion

DCT with an interface, though reasonable, makes a number of incorrect predictions about ion solvation. First, it predicts that ion solvation asymmetry is independent of cavity size, which it is not. Second, it predicts that the difference in charging free energy between a cation and anion of identical charge magnitude grows linearly in that magnitude, which it does not. Third, it predicts that the difference of differences between the cation and anion in bulk and at the interface grows linearly in q, which it does not. These contradictions tell us that knowing $\langle\phi\rangle$ and $\langle(\delta\phi(\sigma,z))^2\rangle$ is insufficient to predict the solvation energy of an ion and that the surface potential responds non-linearly to the presence of that ion.

Bibliography

[1] K.G. Wilson. “Renormalization Group and Critical Phenomena. I. Renormalization Group and the Kadanoff Scaling Picture”. In: *Phys. Rev. B* 4 (9 1971).

[2] K.G. Wilson. “Renormalization Group and Critical Phenomena. II. Phase-Space Cell Analysis of Critical Behavior”. In: *Phys. Rev. B* 4 (9 1971).

[3] L.P. Kadanoff. “Scaling Laws for Ising Models Near Tc”. In: *Physics* 2.6 (1966), pp. 263–272.

[4] M.D. Ediger, C.A. Angell, and Sidney R. Nagel. “Supercooled Liquids and Glasses”. In: *J. Phys. Chem.* 100.31 (1996), pp. 13200–13212.

[5] T.R. Kirkpatrick, D. Thirumalai, and P.G. Wolynes. “Scaling Concepts for the Dynamics of Viscous Liquids Near an Ideal Glassy State”. In: *Phys. Rev. A* 40.2 (1989), pp. 1045–1054.

[6] P. Charbonneau et al. “Dimensional Study of the Caging Order Parameter at the Glass Transition”. In: *Proc. Natl. Acad. Sci.* 109.35 (2012), pp. 13939–13943.

[7] B. Charbonneau et al. “Dimensional Dependence of the Stokes-Einstein Relation and its Violation”. In: *J. Chem. Phys.* 139 (2013), p. 164502.

[8] Y.S. Elmatad, D. Chandler, and J.P. Garrahan. “Corresponding States of Structural Glass Formers”. In: *J. Phys. Chem. B* 113 (2009), pp. 5563–5567.

[9] Y.S. Elmatad, D. Chandler, and J.P. Garrahan. “Corresponding States of Structural Glass Formers. II”. In: *J. Phys. Chem. B* 114 (2010), pp. 17113–17119.

[10] D. Chandler and J.P. Garrahan. “Dynamics on the Way to Forming Glass: Bubbles in Space-Time”. In: *Annu. Rev. Phys. Chem.* 61 (2010), pp. 191–217.

[11] J. Jäckle and S. Eisinger. “A Hierarchically Constrained Kinetic Ising Model”. In: *Condensed Matter* 84 (1991), pp. 115–124.

[12] P. Sollich and M.R. Evans. “Glassy Time-Scale Divergence and Anamalous Coarsening in a Kinetically Constrained Spin Chain”. In: *Phys. Rev. Lett.* 83.16 (1999), pp. 3238–3241.

[13] P. Chleboun, A. Faggionato, and F. Martinelli. “Time Scale Separation in the Low Temperature East Model: Rigorous Results”. In: *J. Stat. Mech.* (2013), p. L04001.

[14] Y.J. Jung, J.P. Garrahan, and D. Chandler. "Excitation Lines and the Breakdown of Stokes-Einstein Relations in Supercooled Liquids". In: *Phys. Rev. E* 69 (2004), p. 061205.

[15] Y.J. Jung et al. "Comment on "Is There a Breakdown of the Stokes-Einstein Relation in Kinetically Constrained Models at Low Temperature?" by O. Blondel and C. Toninelli". In: *arXiv:1307.1651* (2013).

[16] S. Kim et al. "Study of the Upper-Critical Dimension of the East Model Through the Breakdown of the Stokes-Einstein Relation". In: (2017). arXiv: `cond-mat/1705.00095v1`.

[17] P. Ball. "The Borders of Order". In: *New Scientist* 222 (2966 2014), pp. 44–47.

[18] J.T. Lennon and S.E. Jones. "Microbial Seed Banks: The Ecological and Evolutionary Implications of Dormancy". In: *Nature Rev. Microbiol.* 9 (2011), pp. 119 –130.

[19] B.R. Parry and I.V. Surovtsev et al. "The Bacterial Cytoplasm Has Glass-like Properties and Is Fluidized by Metabolic Activity". In: *Cell* 156 (2014), pp. 183–194.

[20] J.P. Hansen and I.R. McDonald. *Theory of simple liquids.* Elsevier, 1990.

[21] A. Worcel, D.S. Goldman, and I.B. Sachs. "Properties and Fine Structure of the Ribosomes from Mycobacterium Tuberculosis". In: *Proc. Natl. Acad. Sci.* 61 (1968), pp. 122 –129.

[22] A.S. Keys et al. "Excitations are localized and relaxation is hierarchical in glass-forming liquids". In: *Phys. Rev. X* 1 (2 2011).

[23] W. Kob and H.C. Andersen. "Scaling Behavior in the β-Relaxation Regime of a Supercooled Lennard-Jones Mixture". In: *Phys. Rev. Lett.* 73.10 (1994).

[24] U. Bengtzelius, W. Gotze, and A Sjolander. "Dynamics of Supercooled Liquids and the Glass Transition". In: *J. Phys. C: Solid State Phys.* 17 (1984), pp. 5915–5934.

[25] D. Frenkel and B. Smit. *Understanding Molecular Simulation.* 2002.

[26] Charles Kittel. *Introduction to Solid State Physics.* 1996.

[27] J.D. Weeks, D. Chandler, and H.C. Andersen. "Role of Repulsive Forces in Determining the Equilibrium Structure of Simple Liquids". In: *J. Chem. Phys.* 54 (1971).

[28] F. Trovato and V. Tozzini. "Diffusion Within the Cytoplasm: A Mesoscale Model of Interacting Macromolecules". In: *Biophys. J.* 107 (2014), pp. 2579–2591.

[29] P.L. Geissler et al. "Autoionization in Liquid Water". In: *Science* 291 (2001), pp. 2121–2124.

[30] J.M.J. Swanson et al. "Proton Solvation and Transport in Aqueous and Biomolecular Systems". In: *J. Phys. Chem. B* 111.17 (2007), pp. 4300–4314.

[31] A.J. Ballard and C. Dellago. "Toward the Mechanism of Ionic Dissociation in Water". In: *J. Phys. Chem. B* 116.45 (2012), pp. 13490 –13497.

[32] B. Reischl, J Kofinger, and C. Dellago. “The Statistics of Electric Field Fluctuations in Liquid Water”. In: *Molecular Physics* 107 (2009), pp. 495–502.

[33] Patrick Stephen Varilly. “Fluctuations in Water and their Relation to the Hydrophobic Effect”. PhD thesis. University of California, Berkeley, 2011.

[34] P.P. Ewald. “Die Berechnung optischer und elektrostatischer Gitterpotentiale”. In: *Ann. Phys. (Berlin)* 369 (3 1921), pp. 253–287.

[35] N. Bjerrum. “Structure and Properties of Ice”. In: *Science* 115 (2989 1952), pp. 385–390.

[36] Y. Zhang and P.S. Cremer. “Interactions Between Macromolecules and Ions: the Hofmeister Series”. In: *Curr. Opin. Chem. Biol.* 10 (6 2006), pp. 658 –663.

[37] N. Schwierz, D. Horinek, and R.R. Netz. “Specific Ion Binding to Carboxylic Surface Groups and the pH Dependence of the Hofmeister Series”. In: *Langmuir* 31 (2015), pp. 215–225.

[38] N. Schwierz, D. Horinek, and R.R. Netz. “Reversed Anionic Hofmeister Series: The Interplay of Surface Charge and Surface Polarity”. In: *Langmuir* 26 (10 2010), pp. 7370–7379.

[39] Patrick Robert Shaffer. “Ion Solvation at Air-Water Interfaces”. PhD thesis. University of California, Berkeley, 2013.

[40] P.B. Petersen and R.J. Saykally. “Adsorption of Ions to the Surface of Dilute Electrolyte Solutions: the Jones-Ray Effect Revisited”. In: *J. Am. Chem. Soc.* 127 (44 2005), pp. 15446–52.

[41] K. Lum, D. Chandler, and J.D. Weeks. “Hydrophobicity at Small and Large Length Scales”. In: *J. Phys. Chem. B* 103 (1999), pp. 4570–4577.

[42] D. Chandler. “Interfaces and the Driving Force of Hydrophobic Assembly”. In: *Nature* 437 (2005), pp. 640–647.

[43] P.L. Chau and J. Hardwick. “A New Order Parameter for Tetrahedral Configurations”. In: *Mol. Phys.* 93 (3 1998), pp. 511 –518.

[44] J.R. Errington and P.G. Debenedetti. “Relationship Between Structural Order and the Anomalies of Water”. In: *Nature* 409 (2001), pp. 318–321.

[45] A. Luzar and D. Chandler. “Effect of Environment on Hydrogen Bond Dynamics in Liquid Water”. In: *Phys. Rev. Lett.* 76.6 (1996), pp. 928 –931.

[46] J.D. Smith et al. “Unified Description of Temperature-Dependent Hydrogen-Bond Rearrangements in Liquid Water”. In: *Proc. Natl. Acad. Sci.* 102.40 (2005), pp. 14171–14174.

[47] G. Hummer et al. “An Information Theory Model of Hydrophobic Interactions”. In: *Proc. Natl. Acad. Sci.* 93 (1996), pp. 8951–8955.

[48] D. Chandler. “Gaussian field model of fluids with an application to polymeric fluids”. In: *Phys. Rev. E* 48.4 (1993).

[49] Y. Levin. “Polarizable Ions at Interfaces”. In: *Phys. Rev. Lett.* 102 (2009).

[50] Y. Levin, A.P. dos Santos, and A. Diehl. “Ions at the Air-Water Interface: An End to a Hundred-Year-Old Mystery?” In: *Phys. Rev. Lett.* 103 (2009).

[51] A.P. dos Santos, A. Diehl, and Y. Levin. “Surface Tensions, Surface Potentials, and the Hofmeister Series of Electrolyte Solutions”. In: *Langmuir* 26 (2010), pp. 10778–10783.

[52] M.D. Baer et al. “Electrochemical Surface Potential Due to Classical Point Charge Models Drives Anion Adsorption to the Air-Water Interface”. In: *J. Phys. Chem. Lett.* 3 (2012), pp. 1565 –1570.

[53] M.A. Wilson and A. Pohorille. “Comment on “Study on the Liquid-Vapor Interface of Water. I. Simulation Results of Thermodynamic Properties and Orientational Structure””. In: *J. Chem. Phys.* 90 (9 1989).