

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Statistical Mechanics and Dynamics of Liquids in and out of Equilibrium

Permalink

<https://escholarship.org/uc/item/15v1r70c>

Author

Hudson, Alexander

Publication Date

2016

Peer reviewed|Thesis/dissertation

Statistical Mechanics and Dynamics of Liquids in and out of Equilibrium

by

Alexander Hudson

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Chemistry

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor David Chandler, Chair

Professor Phillip Geissler

Professor Jamie Cate

Professor Oskar Hallatschek

Summer 2016

Statistical Mechanics and Dynamics of Liquids in and out of Equilibrium

Copyright 2016
by
Alexander Hudson

Abstract

Statistical Mechanics and Dynamics of Liquids in and out of Equilibrium

by

Alexander Hudson

Doctor of Philosophy in Chemistry

University of California, Berkeley

Professor David Chandler, Chair

Liquids display an astonishing array of phenomena that, at first glance, seem to have little in common. In this thesis, we study two of these phenomena. The first is the hydrophobic effect, the well-known tendency of oil and water to demix and the driving force for biological assembly. The second is the glass transition, the process by which liquids cooled to very low temperatures become increasingly viscous until succumbing to structural arrest and displaying the rigidity of a solid with none of its microscopic ordering.

In the first part of this thesis, we expand upon our prior understanding of the hydrophobic effect. We study the behavior of water in the confined volume created by a pair of RNA helices at the inter-subunit interface of the bacterial ribosome, which provides an opportunity to apply our contemporary understanding of the hydrophobic effect, first detailed by Lum, Chandler, and Weeks almost twenty years ago, to a context that is less idealized than previously studied. We also discuss and improve upon computational models of hydrophobic solvation, expanding their applicability to more complicated solute geometries.

In the second part of this thesis, we study the glass transition from the perspective of the East model, a kinetically constrained lattice model of glass formers. We review and derive a number of results for East model glasses that we expect to emerge in more realistic, atomistic models of glass formers. We choose one such model and drive it out of equilibrium by cooling, forming a glass. We then check whether our East model predictions are correct. Of particular interest to us is whether inter-excitation correlations emerge out of equilibrium, a key prediction of the East model. Our atomistic results suggest that they do not, forcing us to reconsider our conceptual mapping between the highly abstract East model and more realistic models of glass formers.

The common thread in these seemingly disparate parts of the thesis is the importance of collective fluctuations in liquid-state phenomena. The hydrophobic effect is a physical consequence of liquid-vapor interfacial fluctuations near extended hydrophobic surfaces, while dynamics in supercooled liquids, the precursors to glasses, are dominated by small, collective motions of particles in an otherwise jammed material. The importance of fluctuations in these and other liquid-state phenomena make the tools of statistical mechanics particularly

suitable to their study. In this thesis, we employ these tools in the service of furthering our understanding of the liquid state.

Contents

Contents	i
1 Structure and dynamics of water and RNA at a ribosomal contact	1
1.1 Introduction	1
1.2 Dynamics and thermodynamics of water at bridge B6	3
1.2.1 Simulation setup: Bridge B6 only	3
1.2.2 Intermittent dynamics of solvent in the interhelical groove	4
1.2.3 Thermodynamics of dewetting in the interhelical groove	7
1.3 Dynamics of RNA at bridge B6	12
1.3.1 Simulation setup: Bridge B6 and its neighborhood	14
1.3.2 Long-timescale dynamics are evident in coarse-grained movies	14
1.3.3 Principal component analysis reveals collective motion	17
2 Improved model of hydrophobic solvation at small lengthscales	21
2.1 Introduction	21
2.2 Lum-Chandler-Weeks Theory	23
2.2.1 Defining a solvent density field	23
2.2.2 Decomposition of the solvent density field	24
2.2.3 Choosing the large length-scale Hamiltonian	25
2.2.4 Integrating out Gaussian fluctuations	26
2.3 Generalization of the one-basis approximation	27
2.4 Superiority of the generalization of the one-basis approximation	29
2.5 Necessity of the Gaussian approximation to accuracy of the model	35
3 Signatures of glasses formed by cooling	38
3.1 Introduction	38
3.1.1 Super-Arrhenius behavior of supercooled liquids	39
3.1.2 Glasses are liquids out of equilibrium	39
3.2 Kinetically constrained models of glass formers	41
3.2.1 Introduction to the East model	41
3.2.2 The East model out of equilibrium	46
3.2.3 Correspondence of the East model with atomistic models of glass formers	50

3.2.4	Formation of East model glasses	51
3.2.5	Inter-excitation correlations in East model glasses	54
3.3	The Wahnström model: An atomistic model of glass-forming liquids	62
3.3.1	Avoiding crystallization	63
3.3.2	Crystallization of the Wahnström model at modestly supercooled conditions	65
3.3.3	Rapid cooling prevents crystallization in the Wahnström liquid	68
3.4	Forming a glass from a Wahnström fluid	70
3.4.1	Calculating non-equilibrium averages	71
3.4.2	Iso-configurational averaging	73
3.4.3	Prediction 1: Crossover from parabolic to Arrhenius	74
3.4.4	Prediction 2: Concentration of soft spots plateaus	76
3.4.5	Prediction 3: Distribution of soft spots becomes correlated	77
3.4.6	Conclusions and discussion	83
3.A	Forming a glass from a two-dimensional fluid	84
	Bibliography	86

Chapter 1

Structure and dynamics of water and RNA at a ribosomal contact

1.1 Introduction

The ribosome is the cellular machinery responsible for translating messenger RNA (mRNA) from the nucleus into the polypeptide chains that fold into functional proteins. It is a huge macromolecular structure comprising multiple chains of RNA and dozens of proteins. At the highest level of organization, it consists of two asymmetric subunits that associate during the translation process. In the bacterial ribosome, these subunits are called 50S (the “large” subunit) and 30S (the “small” subunit).¹ Together they form the 70S bacterial ribosome,² as shown in figure 1.1.

The large interface that forms between the subunits during association consists of dozens of contacts, or *bridges*, between them. In the work that follows we will be primarily concerned with bridge B6, which consists of fragments of two RNA helices — one from each subunit, as shown in figure 1.1 — that form a pocket large enough to accommodate a monolayer of water [32, 6, 23]. Because RNA contains strongly hydrophilic chemical groups, this region is expected to be strongly solvated. However, results of X-ray crystallography show little to no evidence of strongly bound water molecules [32].

One possible explanation for this observation is poor ensemble averaging in the crystal. Because the atomic-scale observations provided by X-ray crystallography are averaged over the large number of copies present in the sample, only regions with structure well-conserved between unit cells will be well-resolved. For this reason, solvent beyond the first one or two solvation cells is rarely if ever observed in crystallography. If the positions of the solvating water molecules at bridge B6 are not preserved between unit cells, they would be difficult to detect in the resulting diffraction pattern.

¹The “S” in the names of the subunits is an abbreviation for Svedberg, a nonlinear unit of mass, hence why the 50S subunit is the large subunit and the 30S subunit is the small subunit.

²The eukaryotic ribosome is larger at 80S.




Figure 1.1: Bridge B6 of the bacterial ribosome [32, 6, 23]. (a) Rendering of the full 70S ribosome, with the small subunit in red/blue and the large subunit in gray/yellow. RNA in blue/gray, protein in red/yellow. (b) The location of bridge B6 on the two subunits. Fragments of each subunit that are part of B6 are rendered in full color. The rendering of the two subunits is produced by taking the perspective in part (a) and then applying the following operations: rotate 90° about the x -axis (so that 30S is occluded by 50S), rotate 90° about the y -axis (so that 30S is to the left and 50S to the right), and then open the ribosome “like a book” (i.e., rotate the two subunits about the y -axis in opposite directions). Intermediate stages of this transformation are rendered in part (c). Green boxes in part (c) indicate the approximate location of B6 with respect to the whole 70S ribosome.

Alternatively, the chemical environment at B6 could be unexpectedly unfavorable to solvation, leading to spontaneous dewetting of the region under the right conditions. This scenario is enabled by the proximity of water at ambient conditions to a liquid-vapor phase transition, which has been identified as the key physical principle underlying the hydrophobic effect [19, 4, 5, 27, 9, 10, 20, 22, 43, 26, 42, 24, 25, 36, 44, 45, 38]. Although liquid water is stable in a bulk environment, the presence of a hydrophobic surface can destabilize the liquid in favor of the vapor phase, often to dramatic effect [11]. In reference [27], Patel et al. show how biological molecules take advantage of this sensitivity to modulate their structure and function. Inspired by the physical principles first detailed by Lum, Chandler, and Weeks in reference [19], Katira et al. propose an “order-phobic” effect, analogous to the hydrophobic effect, in lipid bilayers, which they suggest might partly explain the organization of bilayer-embedded proteins [13].

Our aim here is not so ambitious, but to see to how our understanding of the hydrophobic effect might apply to the inter-subunit contact picture in figure 1.1(b). It might not seem obvious why a theory of hydrophobicity should apply in an environment that is appropriately described as hydrophilic, but sometimes hydrophobicity can be found lurking in unexpected places [18, 31]. In reference [11], Hummer et al. show how an environment that is both hydrophobic and highly confined can drive dramatic wetting and dewetting transitions. Although the carbon nanotube studied by Hummer and co-workers is an idealized setup, it will prove to be a useful example as we consider the hydration of bridge B6 in the next section.

1.2 Dynamics and thermodynamics of water at bridge B6

In this section, we explore solvation of bridge B6 via molecular dynamics simulations. Because of the ribosome’s enormous size, simulation of the whole macromolecular assembly is not feasible. We are thus forced to limit our simulations to relatively small cutouts of the ribosome that include the fragment of interest. Our simplest set of simulations include only the residues that make up bridge B6, restrained by harmonic potentials so that the correct configuration is maintained during simulation, while our largest simulation includes a much larger chunk of the ribosome surrounding B6, restrained only at the outer edges so that the residues in and immediately around B6 can move freely. We will see later that as we include more of the ribosome in our simulations, long time-scale dynamics begin to emerge.

1.2.1 Simulation setup: Bridge B6 only

We begin with the simplest possible simulation setup, in which we include only the residues belonging to bridge B6. This relatively small cutout includes 35 residues and 752 heavy atoms, and is shown in figure 1.1b. To prevent this fragment, which consists of two unattached pieces, from coming apart, it is necessary to apply position restraints to some of the atoms.

In the simulations that follow, we apply the harmonic potential of the form

$$U(\Delta\mathbf{r}_i) = \frac{1}{2}k|\Delta\mathbf{r}_i|^2$$

to a subset of atoms. Here, k is the strength of the biasing³ and $\Delta\mathbf{r}_i$ denotes the deviation of atom i from its position in the crystal structure, or $\Delta\mathbf{r} \equiv \mathbf{r}_i - \mathbf{r}_i^{\text{xtl}}$. If pr denotes the set of atoms that are position restrained, then the total biasing potential added to the system is

$$U_{\text{pr}} = \sum_{j \in \text{pr}} U(\Delta\mathbf{r}_j).$$

Many possible choices of pr are possible. In this section, we choose pr to be the set of all heavy (i.e., non-hydrogen) atoms, so that the RNA remains very near its crystal configuration. This choice of restraints precludes long-timescale, string-like vibrations in the RNA chains that would be possible if only the ends of the chains were restrained. Later, when we expand our simulations to include a much larger cutout of the ribosome and employ a more gentle set of restraints, we will see that this string-like motion is an important component of the dynamics over tens and hundreds of nanoseconds. However, keeping the atoms near their initial positions allows us to construct a static probe volume in the space between the helices, which will be an important tool in analyzing the water density fluctuations in this confined volume.

Details of the simulation

We solvate our B6-only cutout in a box of 6058 TIP4P/2005 water molecules, plus a small number of K^+ and Cl^- ions in order to balance the charge on the RNA. Using the Gromacs 4.5.5 molecular dynamics package [30], we simulate this box of RNA-plus-solvent for 10+ ns, employing the velocity-rescaling method of Bussi et al. [2] to maintain the temperature at $T = 298$ K. To maintain constant pressure, and to ensure that we are near liquid-vapor phase coexistence, we employ a slab setup in which our RNA-plus-solvent is buffered in the $\pm z$ directions by fluctuating liquid-vapor interfaces, similar to the method employed in reference [25], though our interfaces span the entire xy -plane. This setup ensures that the large density fluctuations that accompany dewetting are not artificially suppressed by a finite but fixed volume. The effective pressure under these conditions is $p \approx 0$.

1.2.2 Intermittent dynamics of solvent in the interhelical groove

Figure 1.3 shows configurations of water molecules in the inter-helical groove of bridge B6 at several points during the simulation. This sequence of configurations shows the collective process by which a chain of hydrogen bonded water molecules (blue in figure 1.3) ejects one of its members (the topmost in figure 1.3) from the groove back into the bulk.

³Note: Need to look up what value I chose. It's something like $1k_B T/\text{\AA}^2$.




Figure 1.2: Setup of the B6-only simulations. The two renderings (top and bottom) show the same snapshot after > 100 ns of equilibration. The cutout is solvated in a slab of water of thickness w in the z -direction, with the box size L_z in that direction approximately three times the slab thickness. This setup puts liquid water into coexistence with its vapor phase, and the buffering interfaces on either side allow the effective volume of the slab of water to fluctuate, mimicking constant pressure conditions with $p \approx 0$. This setup is prepared by equilibrating the solvated cutout at constant volume conditions, with $L_z = w$ and only liquid present, then instantaneously changing L_z to $3w$ and equilibrating. The cohesive forces of water, combined with the suddenness of the change in volume, keep the slab intact.

As shown in the top row of snapshots (frames 1-3), the chain of four water molecules maintains its initial hydrogen-bonded structure for hundreds of nanoseconds. Following this long period of inactivity, the chain pulls upward (frame 4) and the topmost hydrogen bond breaks (frame 5), releasing one of the water molecules into the bulk (frame 6). This process takes less than 20 ps to complete. The intermittent dynamics on display here — characterized by long periods of very little movement followed by sudden, collective, and dramatic motion — is representative of the solvent dynamics in this inter-helical groove and is typical of the dynamics of water in confined geometries.

In reference [11], Hummer et al. discover similar dynamics in a hydrophobic carbon nanotube. The nanotube in their setup is large enough to comfortably accommodate a single row of water molecules, and under most conditions it is fully occupied by water. However, under the right conditions, they observe the nanotube spontaneously dewetting. They confirm that the nanotube is bistable, supporting a fully occupied “wet” state and a completely empty “dry” state. Moreover, transitions between these two states are highly intermittent, involving the sudden and collective movement of a one-dimensional chain of water molecules into




Figure 1.3: Dynamics of water molecules in the inter-helical groove of bridge B6. The graphics here show the process by which a water molecule is ejected from the groove back into the bulk. This process is highly collective, involving the concerted motion of four water molecules (blue), and intermittent, occurring very rapidly (~ 10 ps) after a very long period of quiescence (~ 100 ps). Note that the bottom-most three water molecules in the chain each replaces the position of its immediate neighbor.

or out of the nanotube, and occurring between periods of quiescence that are much longer than the time needed to complete the transition. Hummer et al. attribute the dynamics to the confined geometry of the nanotube and the lack of favorable interactions between water and the nanotube, which limit water molecules inside the nanotube to hydrogen bond only with their two immediate neighbors in front and behind. Upon dewetting, the entire chain is ripped out of the nanotube to avoid breaking this chain of hydrogen bonds.

Unlike the carbon nanotube studied by Hummer and coworkers, the inter-helical groove we consider in this section is very hydrophilic, forming strong hydrogen bonds with the water molecules occupying the groove. Consequently, there is no large free energetic penalty to partial occupancy of the confined region as there is in reference [11].⁴ However, despite the clear difference in thermodynamics, the two cases share clear similarities in their dynamics. This suggests that, in confined geometries where water molecules have at most two neighbors, the dynamics is determined by the preservation of precious hydrogen bonds. In both cases, this manifests as a dramatic collective movements in which strings of water molecules move in concert to maintain intra-chain hydrogen bonds. In the case of the nanotube, the result is a complete emptying of the volume; in the case of bridge B6, the result is the far less dramatic rearranging of solvent within the volume.

1.2.3 Thermodynamics of dewetting in the interhelical groove

Straightforward molecular dynamics simulation of our B6-only cutout demonstrates that, similar to the carbon nanotube case, the dynamics of solvent reorganization in the inter-helical groove is collective and intermittent. In contrast to the nanotube, however, no spontaneous dewetting occurs; indeed, the groove remains nearly fully occupied over the course of the whole trajectory. The presence of numerous hydrogen bonds between water molecules and the base groups of the surrounding RNA further supports the notion that the bridge B6 prefers to remain strongly solvated. In order to confirm that this is the case, and to see if there is any possibility of a dewetted state under any conditions, we perform importance sampling on the number of water molecules occupying the interhelical groove.

To do so, we must first define an appropriate probe volume. To that end, we construct a $10\text{\AA} \times 7\text{\AA} \times 12.4\text{\AA}$ cuboidal probe volume encompassing the space between the helices. The order parameter of interest is the number of water molecules in this volume, denoted N_v . The size, position, and orientation of the probe volume is chosen such that N_v counts only those water molecules that occupy the interhelical volume of interest. Although this probe volume, which we denote v , is larger than the actual volume of interest, the part of v that lies outside the interhelical groove is filled by the excluded volume of the RNA fragments and therefore inaccessible to solvent. The number of water molecules in our cuboidal probe volume is therefore an accurate measure of the number of water molecules occupying the interhelical volume.

⁴Indeed, we will see later that the statistics of solvent fluctuations in this region is monostable, strongly favoring a fully solvated state. Moreover, deviations from this state are sub-Gaussian, reflecting how deeply unfavorable it is to dewet the groove.




Figure 1.4: Illustration of the cuboidal probe volume (yellow) used to analyze density fluctuations in the interhelical groove of bridge B6. Each picture is one frame in a longer movie in which the helices and probe volume are rotated in space to give a sense of perspective. The space between the two helices is clearly visible in frames 2 and 3, which show the helices from above and below, i.e., the line-of-sight is parallel to the z -axis. Near the end of the movie, the 50S fragment fades away (frame 5) so that, after additional rotating and zooming (between frames 5 and 6), we can peer into the interhelical space through the 50S fragment (frame 6). This perspective is convenient for observing the activity in the probe volume, and will be used in later figures. The full movie is available online at <https://youtu.be/9KgUQHauvZs>.




Figure 1.5: Left: Number of water molecules N_v (red), and its coarse-grained counterpart $\tilde{N}_v$ (blue), in the probe volume depicted in figure 1.4. Right: Configurations of water molecules (oxygens in red) occupying the probe volume (yellow) at the times marked in the time-series plot (green and yellow stars). The perspective in the configuration snapshots is the same as in frame 6 of figure 1.4.

Because N_v is a discrete quantity, any umbrella potential based on it will be a discontinuous function of particle coordinates and, consequently, introduce impulsive forces.⁵ To circumvent this issue, we perform indirect umbrella sampling as detailed in reference [24, 25], applying biasing potentials that are based on a continuous analog of N_v . To construct such a quantity, we first note that N_v may be written as

$$N_v \equiv \int_{\mathbf{r} \in v} \rho(\mathbf{r}). \quad (1.1)$$

Here, $\rho(\mathbf{r})$ is the solvent density field, which may be expressed as

$$\rho(\mathbf{r}) = \sum_i \delta(\mathbf{r} - \mathbf{r}_i), \quad (1.2)$$

where $\delta(\mathbf{r})$ is the three-dimensional Dirac delta function characterized by the properties

$$\delta(\mathbf{r}) = 0 \text{ for } \mathbf{r} \neq \mathbf{0} \quad \text{and} \quad \int_{\mathbf{r} \in V} \delta(\mathbf{r}) = 1,$$

⁵In practice, the force computed for such a potential will always be zero, since the potential is constant away from the discontinuities.

for any volume V that contains the origin. The combination of equations 1.1 and 1.2 yields another expression for the occupation number,

$$N_v = \sum_i \Theta_v(\mathbf{r}_i), \quad (1.3)$$

where $\Theta_v(\mathbf{r})$ is an indicator function, defined by the integral

$$\Theta_v(\mathbf{r}) \equiv \int_{\mathbf{r}' \in v} \delta(\mathbf{r}' - \mathbf{r}),$$

which is equal to one if $\mathbf{r}$ is located in v , and zero otherwise. Equation 1.3 is an intuitive formula that expresses N_v as a simple count of the particles within v . It should be immediately clear that the delta functions in the expression for $\rho(\mathbf{r})$ in equation 1.2 are the source of the discontinuity in N_v ; its integral, though well-defined, is a discontinuous function. Rectifying the discontinuities in N_v thus requires replacing the delta functions in equation 1.2 with functions that are better behaved, i.e., that spread or coarse-grain the density of a particle over a finite region of space. That is, we wish to turn $\rho(\mathbf{r})$ into $\rho_{cg}(\mathbf{r})$, where ρ_{cg} is a sum of coarse-graining functions $f(\mathbf{r})$,

$$\rho_{cg}(\mathbf{r}) = \sum_i f(\mathbf{r} - \mathbf{r}_i). \quad (1.4)$$

The function $f(\mathbf{r})$ should take its maximum value at the origin and decrease monotonically as the distance from the origin increases. It should satisfy $f(\mathbf{r}) \geq 0$ for every point $\mathbf{r}$, and, like the delta function it replaces, its integral over all space should equal one. Unlike $\delta(\mathbf{r})$, $f(\mathbf{r})$ should be finite and well-defined everywhere. In references [24, 25], Patel et al. choose a truncated, shifted Gaussian for their coarse-graining function. We do the same in the work that follows.

Figure 1.5 compares N_v and $\tilde{N}_v$ for our probe volume over the course of an unbiased simulation. As expected, the two quantities are closely correlated, with large, persistent changes in one quantity reflected in the other quantity. Note that N_v displays larger short-time variation than $\tilde{N}_v$, owing to the fact that small changes in position can cause discrete jumps in N_v as particles move across the boundary of v . Figure 1.5 also shows two example configurations of the water molecules occupying v , one with a relatively high occupation (top right) and one with a relatively low occupation (bottom right), both occurring in the course of the natural fluctuations of the system at equilibrium.

With our order parameter, N_v , and our closely related biasing parameter, $\tilde{N}_v$, suitably constructed, we can perform umbrella sampling and construct the potential of mean force $F(N_v) = -k_B T \ln P(N_v)$ all the way down to $N_v = 0$. We apply the umbrella potential

$$U = \frac{1}{2} \kappa \left(\tilde{N}_v - \tilde{N}_0 \right)^2$$




Figure 1.6: Log probability of finding N_v water molecules within the probe volume v . The red and blue circles show the results of our umbrella sampling calculations. Both sets of data were computed from 10 separate simulations, or windows, with a particular value of the umbrella parameter $\tilde{N}_0$ taken from the set $\{-4, -2, 0, 2, 4, 6, 8, 10, 12, 14\}$. To generate the red data set, we sampled configurations every 10 ps for 100 ns, following 20 ns of equilibration. To generate the blue data set, we sampled every 10 ps for 80 ns, following 10 ns of equilibration. Initial conditions for the red data set were drawn from the equilibrium ensemble with no umbrella potential applied. Initial conditions for the blue data set were taken from the red data set. Specifically, the final configuration of the $\tilde{N}_0 = 14$ window of the red data set was the initial condition for the five “lowest” windows (i.e., $\tilde{N}_0 \in \{-4, -2, 0, 2, 4\}$), and the final configuration of the $\tilde{N}_0 = -4$ window of the red data set was the initial condition for the five “highest” windows ($\tilde{N}_0 \in \{6, 8, 10, 12, 14\}$). The close agreement between the two data sets indicates that the error bars are smaller than the size of the plot markers, and that there are no ill effects from hysteresis. All dashed lines are Gaussian fits to the red data set. In the main plot, the Gaussian with the larger variance is produced by fitting only the data in the range $N_v \in [6, 10]$, while the other Gaussian is fit to the full data set. The inset shows the log probability and full Gaussian fit after coupling a constant force βf to N_v , where $\beta f = 6.798$ and is the minimum force needed to make $N_v = 3$ the most probable value.

to our system for $\kappa = 5$ kJ/mol (approximately $2k_B T$ at 298 K) and several values of $\tilde{N}_0$, sufficient to sample all relevant values⁶ of N_v . We then use the multistate Bennett acceptance ratio (MBAR) method to construct the potential of mean force from our separate simulations [33]. The result is shown in figure 1.6, for two sets of calculations (see the caption for more detail).

The distribution displays a noticeable linearity at low N_v . Whether this is characterized as a “fat” or “skinny” tail — i.e., whether $P(N_v)$ is sub- or super-Gaussian — depends on the Gaussian used for comparison. In figure 1.6, we show two Gaussian fits to the data (gray dashed lines). Fitting data near the mean yields the Gaussian with the larger variance, while fitting the full data set yields the one with the smaller variance. Fluctuations in N_v very close to the mean are likely the result of small fluctuations in the positions of water molecules near the edges of the probe volume, while fluctuations much further from the mean are likely the result of larger, more disruptive rearrangements.⁷ Thus, it is not surprising that a fit to the full data set results in a narrower Gaussian.

Compared to the narrower Gaussian fit, the data in figure 1.6 displays a small fat tail at low N_v . The physical significance of this feature is illustrated in the inset, where we show $\ln P(N_v)$ after coupling a small force, $\beta f = 6.798$, to N_v that favors expulsion of water from the probe volume. The strength of the force we choose is the minimum required to make $N_v = 3$ the most probable value.⁸ With this added force, the distribution becomes nearly flat over the range $N_v \in [0, 5]$, the free energy varying by less than $2k_B T$. The full distribution is highly asymmetric and non-Gaussian. Under application of this small force (which could be mimicked by tuning the pressure or chemical potential), the interhelical volume will frequently and spontaneously dewet. Figure 1.7, and the video associated with it, suggests that this process would, like the process by which the carbon nanotube in reference [11] dewets, involve a collective motion of several water molecules.

We have shown that, despite the interhelical volume being a far less idealized, more complicated, more heterogeneous environment than the carbon nanotube studied by Hummer et al., many of the qualitative features observed in that study seem to be generic to confined volumes where hydrogen bonding — whether due to geometric, chemistry, or both — is constrained.

1.3 Dynamics of RNA at bridge B6

In the last section, we characterized the structure, dynamics, and thermodynamics of the water solvating bridge B6. However, we did so under conditions where the RNA itself was forced to remain near its configuration in the crystal structure. In this section, we turn to a

⁶To reach $N_v = 0$ required choosing $\tilde{N}_0 < 0$. See the caption of figure 1.6 for more details.

⁷For low N_v , this means removing water molecules from a hydrophilic environment; for high N_v , this means packing water molecules in a confined volume that (partly due to our position restraints on the RNA) cannot easily expand.

⁸Note that, because we have chosen the *minimum* such force, $P(3) = P(4)$.




Figure 1.7: Snapshots showing the process by which water molecules are ejected from the probe volume by the umbrella potential $U = (\kappa/2)(\tilde{N}_v - \tilde{N}_0)^2$ for the choice of $\kappa = 5$ kJ/mol and $\tilde{N}_0 = -4$. Both snapshots are taken from a longer movie, spanning 1 ns of simulation, which can be found online at <https://youtu.be/5-F835Gsw5o>.

separate but related question: To what extent is this configuration, observed in experiment and studied in detail in the last section, representative of the structure that prevails under ambient conditions?

Because the crystal structure is the end product of a cryo-cooling protocol, it is, in principle, possible that it differs in small but significant ways from the structure that would be observed under more normal, physiologically relevant conditions. We also know that biological molecules are highly dynamic, and that rather than adopting a single static structure, they fluctuate among a large ensemble of similar structures. What does this ensemble of structures look like? What do the dynamics, the transitions between these structures, look like?

To probe these questions, we resort again to molecular dynamics simulations. In what follows, we not only relax the restrictive set of position restraints that we used in the previous section; we also include a much larger portion of the surrounding ribosome. By restraining only the outer edges of this large cutout, we can much more faithfully capture the dynamics of the helix-helix contact.

1.3.1 Simulation setup: Bridge B6 and its neighborhood

The cutout we consider includes all residues — RNA and protein — with at least one heavy atom within $30\text{\AA}$ of bridge B6 in the crystal structure. Mathematically, this condition can be expressed in terms of an indicator function $h(R)$, defined by

$$1 - h(R) \equiv \prod_{i \in R} \prod_{j \in \text{B6}} \Theta(|\mathbf{r}_i - \mathbf{r}_j| - a),$$

where $a = 30\text{\AA}$. The expression on the right-hand-side is a double product, first over all heavy atoms i that belong to residue R , then over all heavy atoms j that are part of bridge B6. We include all atoms R for which $h(R) = 1$. The resulting cutout is shown in figure 1.8 from numerous viewing angles. The structure of the RNA chains in the neighborhood of bridge B6 is shown in figure 1.9. Note that the 30S fragment of B6 is part of a much larger RNA double helix that sits at the interface of the two subunits, while the 50S fragment is part of an RNA hairpin loop, a common RNA structural motif.

Details of the simulation

We solvate this cutout in a box of 42,623 water molecules and an excess concentration of KCl.⁹ As before, we maintain constant temperature, $T = 298\text{ K}$, using the stochastic velocity rescaling method [2]. In contrast to our B6-only simulations, we use the Parrinello-Rahman barostat to maintain constant pressure ($p = 1\text{ atm}$) conditions in lieu of the slab setup, owing to the larger box size and quantity of water required for the latter.¹⁰ We perform extensive equilibration, annealing the system several times over the course of 500 ns. Following this, we observe the dynamics of the system for an additional $1\text{ }\mu\text{s}$.

1.3.2 Long-timescale dynamics are evident in coarse-grained movies

Short-time vibrational motion (roughly on the order of 1–10 ps), can obscure more important structural dynamics on much longer timescales (10–100 ns or longer). To remove this vibrational motion, we coarse grain our trajectories in time. Specifically, if $\mathbf{r}_i(t)$ denotes the raw coordinate of atom i at time t , we compute

$$\bar{\mathbf{r}}_i(t) \equiv \frac{1}{M\Delta t} \sum_{i=0}^{M-1} \mathbf{r}_i(t - i\Delta t), \quad (1.5)$$

⁹Note that we include an “excess” concentration of ions, so as to provide a more realistic ionic environment. We also include 32 Mg^{2+} ions from the crystal structure.

¹⁰The slab setup requires a large buffer of water in the direction perpendicular to the slab so as to avoid interference between the solute and the liquid-vapor interface. This extra buffer, and the large vacuum on either side of the slab, is eliminated with a barostat, allowing the simulation to run faster.




Figure 1.8: Illustration of the cutout surrounding bridge B6. 30S in gray, 50S in blue. Each rendering is a frame (timestamp given in lower left) taken from a movie in which the cutout is rotated to give a sense of perspective. Portions of the cutout not part of B6 are gradually faded out so that the position of B6 with respect to the rest of the cutout can be understood. Frames 6 and 7 exclude most but not all of the cutout and are meant to show what the immediate surroundings look like around the B6 contact. See figure 1.9 for an alternate rendering of frame 6. Full movie available online at <https://youtu.be/oGMRrB2pgHU>.




Figure 1.9: Frame 6 of figure 1.8 with the phosphate backbone highlighted and all non-phosphorus atoms faded. In this representation, it is clear that the 30S fragment of B6 is part of a much larger double helix structure, while the 30S fragment of B6 is part of a hairpin loop, a common DNA/RNA structural motif in which a single strand of DNA/RNA loops around to base pair with itself.

where Δt is the time resolution and $\tau \equiv M\Delta t$ is the coarse-graining time. In what follows, we choose $\Delta t = 10$ ps and $\tau = 1$ ns. We construct a long movie of this 1- μ s trajectory, rendering frames every 10 ps of the structure of bridge B6. Figure 1.10 shows several illustrative frames from this movie. (The full movie is available at the URL listed in the caption.) Several features of the dynamics are evident in this movie. With the short-time vibrational motion removed, the torsional dynamics dominate the dynamics over relatively short times. Over much longer times, clear conformational changes are evident. Approximately 400 ns into the trajectory (see frames 2, 3, and 4), the helices separate noticeably, dramatically increasing the volume of space between them. After several hundred nanoseconds, the helices come back together.




Figure 1.10: Snapshots of bridge B6 during the simulation, using coarse-grained coordinates as in equation 1.5, with $\Delta t = 10$ ps and $\tau = 1$ ns. The 50S and 30S fragments are on the left and right, respectively. Successive frames are connected by arrows to guide the eye. Each snapshot is identified in figure 1.11 by the colored star in the upper left corner. A notable feature of the dynamics over the timescale of observation ($1 \mu s$) is the opening and closing of the interhelical volume (compare frames 1 and 3). These snapshots are individual frames in a much longer movie, which is available in full at <https://youtu.be/AmMIOIHVPxo>.

1.3.3 Principal component analysis reveals collective motion

To quantify the collective motion observed in figure 1.10 (and associated movie), we perform a principle component analysis. We construct the covariance matrix¹¹

$$C_{\alpha\beta} \equiv \langle \delta x_{\alpha} \delta x_{\beta} \rangle = \langle (x_{\alpha} - \langle x_{\alpha} \rangle) (x_{\beta} - \langle x_{\beta} \rangle) \rangle, \quad (1.6)$$

where α, β index the positional degrees of freedom (e.g., the y coordinate of the i th atom), and the brackets denote a simple average over a set of configurations. In this case, the set

¹¹In constructing this matrix, we use the coordinates of the 752 heavy atoms belonging to bridge B6, excluding all others atoms in the cutout. This makes $C_{\alpha\beta}$ a 2256×2256 matrix.




Figure 1.11: Pseudo-free-energy plots in the p_1-p_2 (left) and p_1-p_3 (right) principal component spaces. The main figures show iso-density surfaces of $\ln N(p_i, p_j)$, where $N(p_i, p_j)$ is the number of times the bin indexed by (p_i, p_j) is sampled during the course of our simulation. The contour spacing is 0.5. The inset shows the color-coded density of $\ln N$, with deep red indicating frequently visited regions, and blue indicating the opposite. White regions are never visited during the simulation. The bin spacing is 0.25 nm^2 for the contour plots (main figure) and 0.10 nm^2 for the density plots (inset). The colored symbols mark the projections of the configurations in figure 1.10 onto these 2d spaces. The plus sign marks the projection of the initial crystal configuration.

of configurations comprises those sampled during the simulation, so that

$$\langle x_\alpha \rangle = \frac{1}{T} \sum_{t=0}^T x_\alpha(t),$$

where T is the length of the simulation. The eigenvectors of $C_{\alpha\beta}$ correspond to collective modes of motion, and their eigenvalues give the contribution of each to the total mean-square displacement. Typically the first few eigenvectors — those with the largest eigenvalues — account for most of the total motion and correspond to slow, long-wavelength motion. These are called the principal components or “essential dynamics” [1]. Projecting configurations onto the first two or three of these eigenmodes greatly simplifies the description of the system and, in this case, allows us to follow the dynamics of the system in terms of a small number of collective variables.




Figure 1.12: Motion corresponding to the first principal component. The snapshots show what the structure of bridge B6 looks like with $p_i = 0$ for $i \neq 1$ and p_1 set to the most extreme values visited during the simulation. Large negative values correspond to the helices being in close contact, while large positive values correspond to the helices being open. A video interpolating between these extremes can be found at <https://youtu.be/KconFdw-IrA>, where links to similar videos for the second and third principle components can also be found.

We project our 1- μ s-long trajectory onto the first three principal components, which together account for over half of the total mean-square displacement. We denote the projection of a configuration onto the i th principal component by p_i . Figure 1.13 plots p_1 and p_2 as functions of time. Figure 1.11 shows pseudo-free-energy plots for the 2-dimensional p_1 - p_2 and p_1 - p_3 spaces, with the configurations in figure 1.10 marked by symbols.

In the top-three principal component space, basins of attraction emerge, corresponding to metastable structural conformations or states. The long-time dynamics is also clarified by this projection. As shown in figure 1.13 (top), the dynamics over hundreds of nanoseconds is characterized by transitions between the basins in figure 1.11. Figure 1.12 demonstrates what motion along the first principal component looks like; the rendered structures correspond to extreme values of p_1 visited during the simulation (while setting $p_i = 0$ for $i \neq 1$). For p_1 , this difference is noticeable, corresponding to the dramatic separation of the helices we identified by eye in figure 1.10. A video interpolating between these two extremes is available online at the provided link. In figure 1.13, this separation of helices corresponds to the large jump in p_1 between 0.4 and 0.6 μ s, and is predictably accompanied by a noticeable increase in the




Figure 1.13: Time series of the first two principal components (top) and the number of water molecules, $N_{5\text{\AA}}$, within $5\text{\AA}$ of bridge B6 (bottom). The opening of the helices observed between 0.4 and $0.6 \mu\text{s}$ corresponds to a clear jump in the first principle component, and is manifested in a large increase in N_v as the solvent-accessible volume also increases.

number of water molecules within $5\text{\AA}$ of bridge B6 (bottom plot) as the solvent-accessible volume increases.

These results strongly suggest that the experimentally observed crystal structure of bridge B6 is well-represented at equilibrium under ambient conditions. Figure 1.11 shows that the crystal structure of bridge B6, when projected on the first three principal components, falls in a region of the phase space that is frequently visited. Indeed, as the time series in figure 1.13 shows, the system finds itself close to the initial configuration at $t = 0$ (following a long annealing/equilibration protocol) and, after wandering away from that region of the phase space, eventually returns. We can therefore be reasonably confident that this structure is not an artificial product of the cryo-cooling protocol.

Chapter 2

Improved model of hydrophobic solvation at small lengthscales

2.1 Introduction

The importance of the hydrophobic effect has been recognized as far back as Kauzmann, who attributed the stability of proteins to the clustering of hydrophobic residues in the interior of proteins [14]. Its everyday manifestation, the demixing of oil and water, betrays the subtlety of its microscopic origin. Even the term “hydrophobic” is misleading, as the tendency for oil and water to demix stems not from any unfavorable interaction between the two, but rather from the strong cohesive forces present in bulk water. The formation of hydrogen bonds in water makes its self-interaction very favorable, and small nonpolar molecules cannot make up for the hydrogen bond interactions they disrupt. Consequently, there appears to be an effective repulsion between water and hydrophobic solutes, which is mediated through the solvent as it rearranges itself to minimize the disruption imposed by the solutes. Understanding this solvent response is key to understanding the hydrophobic effect.

In what follows, we will study ideal hydrophobes that interact with water strictly through excluded volume. That is, their presence excludes any water molecules from entering a specified region. The effect of weak attractions — e.g., dispersion forces — is merely perturbative compared to excluded volume interactions. Of particular interest to us will be the free energy of solvation, ΔG , for an ideal hydrophobe of a given shape and size. From statistical mechanics, we know that the free energy cost of imposing a constraint is directly and simply related to the probability that the constraint is satisfied in its absence. In this case,

$$\Delta G_v = -\ln P_v(0).$$

Above, v denotes the excluded volume of the hydrophobe and $P_v(N)$ is the probability that, in the absence of any constraints, the volume v contains N water molecules. That is, the free energy to accommodate a solute with excluded volume v in bulk water is the negative




Figure 2.1: Qualitative sketch of the solvation free energy, βF , per unit surface area for spherical solutes of radius R . For small solutes ($R \lesssim 1$ nm), the solvation free energy grows approximately with the volume, while for large solutes ($R \gtrsim 1$ nm) it grows with the surface area (with a proportionality constant that is approximately the surface tension γ). In reality, the crossover from one regime to the other is considerably less abrupt than illustrated here.

logarithm of the probability that v spontaneously dries in the absence of the excluded volume constraint. Although $P_v(0)$ is of special significance given its intimate relation to the solvation free energy, the full distribution encodes important information about the response of water to perturbations in its vicinity.

In reference [19], Lum, Chandler, and Weeks (LCW) argue, and confirm with computer simulations, that hydrophobic solvation is qualitatively different at small and large length scales. For small solutes with spatial extent $R \lesssim 1$ nm, ΔG grows with the excluded volume v of the solute, while for larger solutes, ΔG grows with the surface area. Figure 2.1 depicts this in terms of $\Delta G/R^2$, which grows linearly with R until about 1 nm and then plateaus to a constant value.

The behavior of ΔG for large solutes is consistent with the common understanding of the hydrophobic effect, which identifies the key driving force as the reduction in solvent-exposed surface area that accompanies aggregation of hydrophobic solutes. At the molecular level, this effect can be understood in terms of hydrogen bonding. Water molecules at an extended hydrophobic surface lose hydrogen bonds, so a reduction in the area of this surface produces a favorable gain in enthalpy as water molecules are transferred to the bulk. LCW expand on this simple picture, arguing that water at an extended hydrophobic surface behaves similarly

to water at the phase boundary between liquid and vapor.¹ In particular, water at the surface of a large hydrophobic solute forms a genuine liquid-vapor interface that obeys the expected capillary scaling, as confirmed by subsequent computer simulation studies. Thus, for large R , LCW theory is based on the physics of fluctuating interfaces.²

For very small solutes, a qualitatively different mechanism is at play. Water molecules in the vicinity of a small hydrophobic solute can maintain their hydrogen bonding network by carefully organizing themselves around the solute. In doing so, they avoid the large penalty of losing hydrogen bonds, but sacrifice freedom of movement. Consequently, there is an entropic — rather than enthalpic — cost to solvation, which scales roughly with the excluded volume of the solute. This scaling can be understood by appealing to bulk thermodynamics: the free energetic cost to remove bulk water from a volume v is, to a very simple approximation, the pressure³ times the volume, or pv . A more accurate treatment of small-length-scale solvation is based on the empirical observation that, for small volumes, density fluctuations in bulk water are remarkably Gaussian. That is, $P_v(N)$ for small v has the Gaussian form

$$\ln [P_v(N)] = -\frac{1}{2\sigma_v^2} (N - \bar{N}_v)^2 + \text{const.}, \quad (2.1)$$

where $\bar{N}_v$ is the mean number of water molecules in the volume v and $\sigma_v^2 \equiv \langle (\delta N)^2 \rangle$ is the variance or compressibility.⁴ With this in mind, the solvation free energy is given by

$$\Delta G_v = -\ln P_v(0) = \bar{N}_v^2 / 2\sigma_v^2,$$

which scales roughly with v since σ_v^2 grows approximately as $\bar{N}_v$. In a later section, we will show that equation 2.1 is a limiting case of a more general Gaussian field theory, the latter of which forms the small- R component of LCW theory.

2.2 Lum-Chandler-Weeks Theory

In the introduction, we briefly discussed the length-scale dependence of hydrophobic solvation at the heart of the LCW theory of hydrophobic solvation. In this section, we present a brief primer on the theory.

2.2.1 Defining a solvent density field

Because solvation is intimately related to the natural density fluctuations of the solvent, we begin by considering $\rho(\mathbf{r})$, the solvent density at a point $\mathbf{r}$. If the solvent consists of N point

¹Water at this phase boundary similarly loses hydrogen bonds.

²Although modern LCW theory heavily emphasizes the importance of *fluctuations* in the liquid-vapor interface, the original LCW paper provides only a mean-field treatment.

³Recall that pressure has units of energy density.

⁴Intuitively, a more compressible fluid allows for larger density fluctuations.

particles with positions $\{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N\}$, then we may express the density as a sum of delta functions:

$$\rho(\mathbf{r}) = \sum_{i=1}^N \delta(\mathbf{r} - \mathbf{r}_i). \quad (2.2)$$

In the above, $\delta(\mathbf{r})$ is the three-dimensional Dirac delta function, characterized by the properties

$$\delta(\mathbf{r}) = 0 \text{ for } \mathbf{r} \neq \mathbf{0} \quad \text{and} \quad \int_{\mathbf{r}} \delta(\mathbf{r}) = 1,$$

where $\mathbf{0}$ is the null vector $(0, 0, 0)$. With this definition, the integral of $\rho(\mathbf{r})$ over a volume v gives the number of solvent particles in the volume,

$$N_v = \int_{\mathbf{r} \in v} \rho(\mathbf{r}),$$

as required. Furthermore, for a box of pure solvent, the equilibrium ensemble average of $\rho(\mathbf{r})$ returns the bulk density, $\langle \rho(\mathbf{r}) \rangle = \bar{\rho}$, where $\bar{\rho} = N/V$ for a box with fixed volume V .

Although it is possible to work in terms of the discrete particle coordinates $\{\mathbf{r}_i\}$, we will find it convenient to work in terms of the field $\rho(\mathbf{r})$, which — via equation 2.2 above — fully specifies the configuration of the system. In thermal equilibrium, the system will fluctuate according to a Hamiltonian $H_0(\mathcal{C})$, where $\mathcal{C}$ denotes a configuration of the system. Since a configuration implies a particular density field $\rho(\mathbf{r})$, we write this as $H_0[\rho(\mathbf{r})]$ to emphasize that a particular realization of the field is associated with some value of H_0 .

Because $\rho(\mathbf{r})$ is field defined at every point in the system, H_0 technically depends on an infinity of variables. We can think of real space as consisting of a finite number of grid points, each associated with a small volume Δv , with ρ assuming values of 0 or $1/\Delta v$ at each point depending on whether a solvent particle occupies the volume associated with that grid point. The field $\rho(\mathbf{r})$ can then be expressed as a finite list of binary variables $\{\rho_1, \rho_2, \dots, \rho_M\}$, where ρ_i is the density at grid point i , given by

$$\rho_i = \int_{\mathbf{r} \in i} \rho(\mathbf{r}),$$

and M is the total number of grid points. With the grid defined in this way, $H_0[\rho(\mathbf{r})]$ is replaced by the simpler, more intuitive $H_0(\rho_1, \rho_2, \dots, \rho_M)$. The continuum case is obviously recovered in the limit $\Delta v \rightarrow 0$ or, equivalently, $M \rightarrow \infty$.

2.2.2 Decomposition of the solvent density field

The observation by LCW that hydrophobic solvation is physically different at small and large length scales motivates the following decomposition of the density field into small- and large-wavelength components:

$$\rho(\mathbf{r}) = \bar{\rho}n(\mathbf{r}) + \delta\rho(\mathbf{r}). \quad (2.3)$$

Here, $n(\mathbf{r})$ is a field that varies slowly in space and time, while $\delta\rho(\mathbf{r})$ is quickly varying field with Gaussian statistics. We may think of the field $\bar{\rho}n(\mathbf{r})$ as the density that remains after coarse-graining the full density $\rho(\mathbf{r})$ over a lengthscale λ that is large enough to remove the Gaussian fluctuations in $\delta\rho(\mathbf{r})$.⁵ That is, $\bar{\rho}n(\mathbf{r})$ is the part of the density field that remains after integrating out the Gaussian fluctuations, which persist only over short length- and time-scales. With this decomposition, the Hamiltonian can be written as

$$H_0[\rho(\mathbf{r})] = H_L[n(\mathbf{r})] + H_S[\delta\rho(\mathbf{r})] + H_I[n(\mathbf{r}), \delta\rho(\mathbf{r})],$$

where H_S has the bilinear form

$$H_S = \frac{1}{2} \int_{\mathbf{r}} \int_{\mathbf{r}'} \delta\rho(\mathbf{r}) \chi^{-1}(\mathbf{r}, \mathbf{r}') \delta\rho(\mathbf{r}'),$$

H_I gives the coupling between the short- and long-wavelength fields (an issue which we will briefly revisit later), and H_L governs the coarse-grained density field $\bar{\rho}n(\mathbf{r})$. Consistent with empirical observations of large length-scale solvation, H_L should support a thermodynamic transition between a low-density vapor phase ($n \approx 0$) and a high-density liquid phase ($n \approx 1$). In reference [19], LCW choose a square-gradient Hamiltonian with the parameters chosen to reproduce water's experimentally known surface tension.

The physical picture implicit in 2.3 and subsequent equations is as follows: over relatively large length- and time-scales, the field $n(\mathbf{r})$ varies smoothly between high- and low-density phases (corresponding to liquid and vapor, respectively), while on shorter length- and time-scales, $\delta\rho(\mathbf{r})$ fluctuates about $\bar{\rho}n(\mathbf{r})$.

2.2.3 Choosing the large length-scale Hamiltonian

As previously mentioned, LCW choose a square-gradient Hamiltonian for H_L , which has the form

$$H_L = \int_{\mathbf{r}} w[n(\mathbf{r})] + \frac{1}{2} m |\nabla n(\mathbf{r})|^2. \quad (2.4)$$

The first term, $w[n(\mathbf{r})]$, is a free-energy density with a functional form that supports a first-order transition between a low-density phase at $n \approx 0$ and a high-density phase at $n \approx 1$. It depends implicitly on an external parameter, such as pressure or chemical potential, that tunes the relative stability of the two phases, e.g.,

$$w[n(\mathbf{r})] = w[n(\mathbf{r}); \mu].$$

The second term in equation 2.4 is a gradient term, which penalizes rapid spatial changes in the density and thus favors a homogeneous $n(\mathbf{r})$.

At coexistence, with no boundary conditions imposed on the system, there will be two solutions that minimize the free energy: $n(\mathbf{r}) = 1$ and $n(\mathbf{r}) = 0$. Imposing certain boundary

⁵The coarse-graining length should be of the order of the size of a water molecule, $\lambda \sim 3.5 \text{ \AA}$.

conditions, such as $\lim_{z \rightarrow -\infty} [n(\mathbf{r})] = 1$, will break this degeneracy in favor of one solution. Other boundary conditions may render it impossible to minimize the terms in H_L simultaneously. In this case, the resulting competition between $w(n)$ and the square gradient term will give rise to a nontrivial, inhomogeneous solution.

As an example of this competition, consider an infinite system at coexistence, upon which we impose the following set of boundary conditions:

$$\lim_{z \rightarrow -\infty} n(\mathbf{r}) = 1 \quad \text{and} \quad \lim_{z \rightarrow \infty} n(\mathbf{r}) = 0.$$

These boundary conditions preclude a single homogeneous phase. Both phases will be present in the system, and there will be a transition between them. The first term in H_L , the free-energy density, favors a discontinuous transition between homogeneous phases: to the left of the phase boundary, $n = 1$ everywhere; to the right, $n = 0$ everywhere. On the other hand, the square gradient term in H_L strongly disfavors such abrupt transitions, preferring a transition as broad as possible. In other words, the free energy density w favors an interface with width $\lambda = 0$, while the square-gradient term favors one with infinite width or $\lambda = \infty$. The competition between these terms in H_L produces a density profile $n(z)$ that varies gradually from $n = 1$ to $n = 0$ over an interfacial region of finite thickness, $0 < \lambda < \infty$.

The simple square-gradient Hamiltonian in equation 2.4 is thus sufficient to produce the salient features of large-length scale solvation: a first-order transition and interfaces between coexisting phases. The preceding analysis, though easy to perform, is based on a mean-field approximation. Accounting for the fluctuations that have been neglected is a nontrivial task. Simulation is a natural tool for this situations, but H_L as written in equation 2.4 does not lend itself easily to simulation. Prior work has focused on discretizing the long-wavelength field to produce an Ising-like Hamiltonian more amenable to simulation [45, 38]. References [36, 37] demonstrate that a simple lattice gas is minimally sufficient to provide quantitatively accurate results, provided that the parameters of the model are carefully chosen⁶ and that a Gaussian contribution is included to penalize overlap between occupied (or “wet”) cells and the excluded volumes of solutes.

2.2.4 Integrating out Gaussian fluctuations

It is convenient to integrate out the fluctuations in $\delta\rho(\mathbf{r})$ to yield an effective Hamiltonian for the slowly varying field $n(\mathbf{r})$. Because $\delta\rho(\mathbf{r})$ has Gaussian statistics, this operation is straightforward by considering $\delta\rho(\mathbf{r})$ as a vector indexed by $\mathbf{r}$. Discretizing space as we did in the previous section, dividing the volume of the system into M small subvolumes, turns $\delta\rho(\mathbf{r})$ into an M -dimensional Gaussian vector. Integrating out the fluctuations in $\delta\rho(\mathbf{r})$ then requires evaluating an M -dimensional Gaussian, a straightforward process which, in the limit $M \rightarrow \infty$, recovers the continuum case.

⁶In particular, the temperature of the lattice gas must be high enough to avoid falling below a temperature T_R , where a so-called *roughening transition* occurs. For $T < T_R$, the fluctuations of the lattice become quiescent, and no longer produces a correct capillary spectrum [36].

Prior work to create models that are easy to simulate have eschewed a full field-theoretic calculation in favor of the so-called “one-basis approximation” [45, 38, 36]. Physically, this approximation corresponds to constraining the mean density in an excluded volume v to be zero, rather than forcing the density to vanish at each and every point in v . This approximation has proven very effective in producing quantitatively accurate solvation free energies for solutes of various shapes and sizes [36]. In the next section, we derive a generalization of this approximation that is applicable to collections of independent solutes with non-overlapping excluded volumes.

2.3 Generalization of the one-basis approximation

In this section, we demonstrate how clever approximation can greatly simplify the small-lengthscale component of an LCW-based model, making it computationally convenient without sacrificing (much) accuracy [37]. We begin by considering the effect of a solute’s excluded volume on the surrounding solvent density. Recall that a solvent-excluding volume v introduces a constraint on the total solvent density field,

$$\rho_l n(\mathbf{r}) + \delta\rho(\mathbf{r}) = 0, \quad (2.5)$$

at each point $\mathbf{r} \in v$. Integrating out small-wavelength Gaussian fluctuations $\delta\rho(\mathbf{r})$ in the presence of this constraint requires evaluating

$$Z_v = \int \mathcal{D}\delta\rho(\mathbf{r}) e^{-\beta H_s} \prod_{\mathbf{r} \in v} \delta(\rho_l n(\mathbf{r}) + \delta\rho(\mathbf{r})). \quad (2.6)$$

The product of delta functions enforces the constraint in equation 2.5, and H_s is the unconstrained Hamiltonian for $\delta\rho(\mathbf{r})$,

$$\beta H_s = \frac{1}{2} \int_{\mathbf{r}} \int_{\mathbf{r}'} \delta\rho(\mathbf{r}) \chi^{-1}(\mathbf{r}, \mathbf{r}') \delta\rho(\mathbf{r}').$$

A formal expression for Z_v can be easily derived, but numerical evaluation is impractical due to the infinite product of delta functions. For this reason, prior work on coarse graining water has replaced the pointwise constraint in equation 2.6 by a single constraint on the average density in v , as in equation (4). The resulting expression for Z_v is much more manageable and has been applied with great success to systems containing only a single solute with a simple contiguous geometry. However, for systems with many independently moving solutes that may be separated by large distances, the average constraint in equation (4) introduces spurious correlations between solutes that never vanish. If v is the union of m non-intersecting regions,

$$v = v^{(1)} \cup v^{(2)} \cup \dots \cup v^{(m)},$$

a simple solution is to constrain the average density in each region $v^{(\alpha)}$ separately. The infinite product of delta functions in 2.6 then becomes a product of m delta functions,

$$\prod_{\mathbf{r} \in v} \delta(\rho_l n(\mathbf{r}) + \delta\rho(\mathbf{r})) \rightarrow \prod_{\alpha=1}^m \delta\left(\int_{\mathbf{r} \in v^{(\alpha)}} \rho_l n(\mathbf{r}) + \delta\rho(\mathbf{r})\right). \quad (2.7)$$

Using the Fourier representation of the delta function, the integral in equation 2.6 becomes

$$\int \mathcal{D}\delta\rho(\mathbf{r}) e^{-\beta H_s} \prod_{\alpha=1}^m \int d\psi_\alpha \exp\left(-2\pi i \psi_\alpha \int_{\mathbf{r} \in v^{(\alpha)}} \rho_l n(\mathbf{r}) + \delta\rho(\mathbf{r})\right). \quad (2.8)$$

Defining

$$N_\alpha = \int_{\mathbf{r} \in v^{(\alpha)}} \rho_l n(\mathbf{r}) \quad \text{and} \quad \Phi(\mathbf{r}) = \begin{cases} 2\pi i \psi_\alpha, & \mathbf{r} \in v^{(\alpha)}, \\ 0, & \text{else,} \end{cases}$$

and rearranging the order of the integrals, we obtain

$$\prod_{\alpha=1}^m \int d\psi_\alpha e^{-2\pi i N_\alpha \psi_\alpha} \int \mathcal{D}\delta\rho(\mathbf{r}) \exp\left(-\beta H_s - \int_{\mathbf{r}} \Phi(\mathbf{r}) \delta\rho(\mathbf{r})\right). \quad (2.9)$$

Evaluating the inner integral over $\delta\rho(\mathbf{r})$ results in

$$Z_v = Z_0 \prod_{\alpha=1}^m \int d\psi_\alpha e^{-2\pi i N_\alpha \psi_\alpha} \exp\left(\frac{1}{2} \int_{\mathbf{r}} \int_{\mathbf{r}'} \Phi(\mathbf{r}) \chi(\mathbf{r}, \mathbf{r}') \Phi(\mathbf{r}')\right).$$

The term in the exponent evaluates to

$$-\frac{1}{2} (2\pi)^2 \sum_{\alpha=1}^m \sum_{\beta=1}^m \psi_\alpha \sigma_{\alpha\beta} \psi_\beta,$$

where $\sigma_{\alpha\beta}$ is given by the double integral

$$\sigma_{\alpha\beta} = \int_{\mathbf{r} \in v^{(\alpha)}} \int_{\mathbf{r}' \in v^{(\beta)}} \Theta(\mathbf{r}) \chi(\mathbf{r}, \mathbf{r}') \Theta(\mathbf{r}').$$

With this, Z_v can be expressed compactly in matrix notation,

$$Z_v = Z_0 \prod_{\alpha=1}^m \int \left(\frac{d\psi_\alpha}{2\pi}\right) \exp\left(-\frac{1}{2} \boldsymbol{\psi}^\top \boldsymbol{\sigma} \boldsymbol{\psi} - i \mathbf{N}^\top \boldsymbol{\psi}\right),$$

where $\boldsymbol{\psi}$ is a column vector with elements $2\pi\psi_\alpha$, $\mathbf{N}$ is a column vector with elements N_α , and $\boldsymbol{\sigma}$ is an $m \times m$ matrix with elements $\sigma_{\alpha\beta}$. This integral is easily evaluated to give

$$Z_v = Z_0 \frac{1}{\sqrt{\det(2\pi\boldsymbol{\sigma})}} \exp\left(-\frac{1}{2} \mathbf{N}^\top \boldsymbol{\sigma}^{-1} \mathbf{N}\right).$$

The free energetic contribution of the small-wavelength field, relative to the unconstrained case, is then given by

$$F_v - F_0 = -T \ln \left(\frac{Z_v}{Z_0} \right) = \frac{T}{2} \left(\ln(\det(2\pi\boldsymbol{\sigma})) + \mathbf{N}^\top \boldsymbol{\sigma}^{-1} \mathbf{N} \right), \quad (2.10)$$

which are the final two terms in equation (5). For two solutes α and β , the off-diagonal matrix element $\sigma_{\alpha\beta}$ gives the degree to which Gaussian fluctuations in regions $v^{(\alpha)}$ and $v^{(\beta)}$ are coupled. If the solutes are separated by much more than the correlation length of liquid water, this term is zero, and Gaussian fluctuations in one region have no effect on the other. In the limit of infinite separation between all solutes, $\boldsymbol{\sigma}$ becomes a diagonal matrix and the free energy in equation 2.10 reduces to a simple sum of uncoupled terms,

$$F_v - F_0 = \frac{T}{2} \sum_{\alpha=1}^m \left(\ln(2\pi\sigma_{\alpha\alpha}) + \frac{N_\alpha^2}{\sigma_{\alpha\alpha}} \right). \quad (2.11)$$

This expression is a natural generalization of equation 6 in [PRL] for many well-separated solutes. The one-basis approximation is recovered by setting $m = 1$. In the next section, we explore several examples where this generalization is necessary to produce quantitatively accurate and physically reasonable results.

2.4 Superiority of the generalization of the one-basis approximation

In this section, we present several simple examples, featuring multiple hydrophobic solutes, in which the standard one-basis approximation fails and its generalization is necessary to cure its major ills. In what follows, we often consider two solutes separated by much more than a lengthscale λ_χ , which we define to be the distance over which the response function $\chi(\mathbf{r}, \mathbf{r}')$ decays to zero. In practice, when we say that two solutes with excluded volumes v and v' are separated by a distance $\gg \lambda_\chi$, we mean that the solutes are sufficiently far apart that, for all pairs of cells i that overlap partially with v and cells j that overlap partially with v' , the matrix element $\chi_{ij} = 0$, where

$$\chi_{ij} = \int_{\mathbf{r} \in i} \int_{\mathbf{r}' \in j} \chi(\mathbf{r}, \mathbf{r}').$$

In the m -basis approximation, this corresponds to vanishing cross terms in the matrix $\boldsymbol{\sigma}$.

Example 1

Consider two identical solutes of radius $R = 8\text{\AA}$, with excluded volumes v and v' , separated by a distance much larger than λ_χ , as depicted in figure 2.2. We examine fluctuations in the




Figure 2.2: Probability of observing n_v , the number of filled cells that completely overlap with the excluded volume v , when there is an identical excluded volume v' in the box but far away from v . Red curves correspond to calculations using the one-basis approximation, while blue curves correspond to calculations using its multi-basis generalization. The excluded volumes are spheres of radii $8\text{\AA}$, and are separated by a distance that is sufficiently large that $\chi(\mathbf{r}, \mathbf{r}') = 0$ for all $\mathbf{r} \in v$ and $\mathbf{r}' \in v'$. For solutes of this size, the maximum value of n_v is 210. The main plot (left) shows how this distribution differs between the one-basis and m -basis approximations. The behavior near the mean is similar, but there are significant differences in the low-probability, large- n_v tail of the distribution, with the one-basis approximation vastly overestimating the free energetic cost of filling the excluded volume v . The two small plots (right) show how $P(n_v)$ changes if v' is constrained to be completely filled ($n_{v'} = 210$). Because v' is so far away from v , constraints on $n_{v'}$ should have no effect on n_v . The multi-basis approximation correctly produces this result, while the one-basis approximation shows a significant, spurious difference: with v' completely filled, large- n_v fluctuations become much less likely (top right, red dashed line).

quantity

$$n_v = \sum_i \delta(v_i - 1)n_i,$$

where $\delta(x)$ is the Kronecker delta: $\delta(0) = 1$ and $\delta(x) = 0$ for $x \neq 0$. This quantity counts the number of filled cells that fully overlap with v . It is a measure of how much the long-wavelength field “penetrates” the excluded volume of the solute. Large values of n_v minimize the energetic contribution of the lattice gas Hamiltonian⁷ at the expense of a larger Gaussian contribution.⁸

We examine the distribution of n_v , subject to constraints on $n_{v'}$. That is, we consider what happens to fluctuations in the long-wavelength density field within the excluded volume v when we “push” on the long-wavelength field within another excluded volume very far away. The main plot of figure 2.2 (left) shows $P(n_v)$ for both the one-basis and multi-basis approximations when $n_{v'}$ is left unconstrained. The two distributions look very similar at low n_v , but differ significantly at high n_v . This can be reasoned as follows: Large values of n_v correspond to deep penetration of the long-wavelength field into the excluded volume v . With the excluded volume “wet,” the Gaussian terms dominates the Hamiltonian, and the spurious correlations introduced by the one-basis approximation become very important. By contrast, low values of n_v correspond to evacuation of the excluded volume by the long-wavelength field, which reduces the contribution of the Gaussian model and its approximations in favor of the lattice gas and its capillary fluctuations.

The small plots in figure 2.2 (right) show what happens when we constrain the other excluded volume, v' , to be completely filled. The large separation between v and v' should ensure that constraints on v' have no effect on the fluctuations in v , but as the top right plot shows, the one-basis approximation fails in this regard, with large- n_v fluctuations becoming *much* less likely when v' is filled.

Example 2

Consider a solute of radius R . What is the solvation free energy of this solute as a function of R , given the presence of another solute of radius r a distance $\delta R \gg \lambda_\chi$ away? As before, the presence of the second solute should have no effect. Figure 2.3 shows that this is not the case for the one-basis approximation (red lines). Compared to the m -basis approximation (blue), which produces the correct result, the one-basis approximation underestimates the solvation free energy. For small r , this underestimate is most significant at small and large lengthscales, while for large r it is most pronounced in the crossover region.

Note that the solvation free energy is largely unaffected for large R when r (right plot). When r and R are both large, solvation is dominated by the lattice gas component of the Hamiltonian. Capillary fluctuations, rather than Gaussian fluctuations, are the important contribution. The approximation invoked for the latter matters less. However, a curious

⁷There are fewer nearest neighbor cells that differ in their occupancy.

⁸The value of N_v in the Gaussian term $N_v^2/2\sigma_v$ is larger, while σ_v is smaller.




Figure 2.3: Solvation free energy, βF , per unit surface area of spherical solutes versus radius R of the excluded volume. Results in red correspond to one-basis calculations, while results in blue correspond to m -basis (in this case, $m = 2$) calculations. The setup is shown at lower right. We compute the solvation free energy of a spherical solute of radius R with another spherical solute, of radius r , elsewhere in the box. The separation between the two solutes is such that $\chi(\mathbf{r}, \mathbf{r}') = 0$ for all points $\mathbf{r}$ and $\mathbf{r}'$ in the two excluded volumes, respectively. The left plot shows the results for $r = 5\text{\AA}$, while the right plot shows the results for $r = 10\text{\AA}$. For both choices of r , the use of the one-basis approximation in a system with more than one solute gives rise to quantitative disagreements with the more physically motivated m -basis result. Interestingly, the lengthscales where the one-basis and m -basis results diverge most depends on the choice of r . The inset compares the m -basis results for the two choices of r , confirming that, as expected, there is no difference when the two solutes are very far apart.




Figure 2.4: Normalized cross-correlation coefficient, $C_{01}/\sqrt{C_{00}C_{11}}$, for two solutes, where $C_{\alpha\beta} \equiv \langle \delta N_\alpha \delta N_\beta \rangle$ and $N_\alpha \equiv \rho \sum_i n_i v_i^{(\alpha)}$. This quantity measures the degree to which fluctuations of the large-lengthscale field in the excluded volumes of the two solutes are correlated. One-basis results are shown in red, m -basis results are shown in blue. Left shows the results for two well-separated spherical solutes of radius R , as R varies from $5\text{\AA}$ (small lengthscale) to $12\text{\AA}$ (large lengthscale). Not surprisingly, there is no correlation ($C_{01} = 0$) in the m -basis model, as expected and required. In contrast, the one-basis model suffers from spurious correlations, which are worst when R is small and the Gaussian component of the model dominates. Right show the results for two spherical solutes of radius $R = 10\text{\AA}$ as the distance separating them, δR , varies. Correlations in the m -basis model decay to zero around $\delta R = 5\text{\AA}$, whereas they plateau to a negative value for the one-basis model, indicating the presence of anti-correlations in N_0 and N_1 . Note that N_0 and N_1 are almost already anti-correlated in the one-basis model, except near the peak at $\delta R = 3\text{\AA}$, where spontaneous dewetting occurs between the solutes. In contrast, the m -basis model exhibits positive correlations for $\delta R < 5\text{\AA}$.

feature of the results in figure 2.3 is that a single, small solute ($r = 5\text{\AA}$ in the left plot) can affect the solvation free energy of the other solute even when R is large. This suggests that the accuracy of the model, even at large lengthscales, depends intimately on the small-lengthscale component. We will briefly return to this issue in the next section.

Example 3

We now consider correlations between N_v and $N_{v'}$, where

$$N_v \equiv \rho \sum_i n_i v_i,$$

and similarly for $N_{v'}$, for two identical solutes with excluded volumes v and v' . First consider, as in the previous examples, the case where these two solutes are very far apart. How does the correlation between N_v and $N_{v'}$, as measured by the correlation coefficient

$$C \equiv \frac{\langle \delta N_v \delta N_{v'} \rangle}{\langle (\delta N_v)^2 \rangle^{1/2} \langle (\delta N_{v'})^2 \rangle^{1/2}},$$

where $\delta N_v \equiv N_v - \langle N_v \rangle$, change as the radius R of these two solutes changes? Our expectation, of course, is that $C = 0$ for all values of R , since δR , the distance separating the solutes, satisfies $\delta R \gg \lambda_\chi$. Figure 2.4 (left) confirms that the multi-basis approximation produces no correlations between N_v and $N_{v'}$, while the one-basis approximation introduces significant negative correlations between these quantities. That is, in the one-basis approximation, an increase in N_v tends to be offset by a decrease in $N_{v'}$, and vice versa, even when the solutes are far apart. This anti-correlation is most significant for small R , but it is still nonzero at larger lengthscales.

Now consider fixing R at $R = 10\text{\AA}$, and consider what happens to C as δR is varied. Not surprisingly, while the multi-basis result tends to zero for large δR , the one-basis result tends to a negative value. Interestingly, for both approximations, the correlation spikes around $\delta R = 3\text{\AA}$, which is likely the result of a dewetting event in which the cells between the two solutes spontaneously evacuate. This evacuation event would tend to decrease both N_v and $N_{v'}$, so that there is a large positive correlation between them. Interestingly, the multi-basis approximation produces non-negative correlations for all values of δR , while the one-basis approximation produces negative correlations for all R , except at the value of R that produces a dewetting event.

This example demonstrates more vividly than the other two examples how a strictly one-basis treatment of multiple solutes produces strange, unphysical (and entirely expected) results. The results shown in figure 2.3 (right) might suggest that the choice of approximation is largely irrelevant as long as the solutes under consideration are sufficiently large that the capillary fluctuations dominate the physics. But however accurate the one-basis approximation may be in producing estimates of the solvation free energy, figures 2.4 and 2.2 both demonstrate that this accidental accuracy hides serious inaccuracies in the fluctuations, and that whatever accuracy it achieves in the mean does not extend to the wings of the probability distributions. The fact that the one-basis approximation produces negative correlations for nearly all δR , as shown in figure 2.4 (right), suggests that the one-basis approximation samples a very different set of configurations⁹ than it should as the solutes come together.

⁹In particular, the binary field n_i probably resembles something like a dumbbell-shaped balloon as it is squeezed from either end.




Figure 2.5: Solvation free energy, βF , per unit surface area for spherical solutes of radius R , computed using the models in equation 2.11 (red) and equation 2.12 (blue). Although both models produce clear crossover behavior, and at around the same lengthscale, the quantitative details of this crossover differ significantly between the two. In particular, the simpler model, which contains only a linear (rather than bilinear) term in the overlap volumes, produces a swift crossover from small- to large-lengthscale solvation. This is in stark contrast to the broad, nontrivial crossover produced by the more accurate model in equation 2.11 and observed in atomistic models of water. Additionally, even at large lengthscales (e.g., $R > 1 \text{ nm}$) the two models differ significantly in their estimates of the solvation free energy, despite the lack of any differences in the large-lengthscale component of the Hamiltonian. This demonstrates that, even at relatively large lengthscales, where the physics is dominated by capillary fluctuations, the small-lengthscale, Gaussian contribution is still important.

2.5 Necessity of the Gaussian approximation to accuracy of the model

In the previous section, we saw how an accurate treatment of the small-lengthscale component of the model is important to the quantitative accuracy of the model and, in particular, that it is necessary to generalize the one-basis approximation in order to remove unphysical fluctuations from the model when multiple independent solutes are present. We also saw, in figure 2.3 (left), how a poor treatment of small-lengthscale solvation can noticeably affect the accuracy of the model even at large lengthscales where capillary modes dominate and

we might imagine that inaccuracies in the small-lengthscale component are inconsequential. In this section, we expand upon this latter point.

Figure 2.5 shows the solvation free energy for two models that feature identical large-lengthscale components but different treatments of small-lengthscale approximation. The red curve corresponds to the model used in this chapter and in reference [36]. The blue curve corresponds to the model used in reference [44], with the parameters of the lattice gas component adjusted to match those in reference [36]. This model, reproduced here,

$$H_{tW} = -\epsilon \sum_{\langle ij \rangle} n_i n_j - \mu \sum_i n_i + \sum_i c v_i n_i, \quad (2.12)$$

has a very simple small-lengthscale component. The parameter c in the last term is chosen to match experimentally known solvation free energies of small hydrophobic solutes, and has the value $c = 0.06 k_B T / \text{\AA}^3$ [44]. The solvation free energies produced by these two models differ significantly in their quantitative details, even as they both produce crossover behavior around the same lengthscale. In particular, the model in equation 2.12 produces none of the subtle details at small values of R that the more complex, and more accurate, model of reference [36] does. The model of equation 2.12 predicts that the solvation free energy per unit surface area increases perfectly linearly for small R before very abruptly plateauing. Clearly, a more accurate treatment of small-lengthscale solvation is necessary to capture with quantitative accuracy the subtle details of the broad crossover.

Perhaps more surprising than the lack of quantitate accuracy at small and intermediate values of R is that the blue curve in 2.5 plateaus to a noticeably lower value than the red curve, suggesting that the ills of equation 2.12's treatment of small-lengthscale solvation have consequences even at larger lengthscales, an effect we already noted in figure 2.3 (right) when considering the differences between one- and multi-basis treatments of small-lengthscale solvation.

Why is there such a large discrepancy even when the lattice gas component should dominate the physics? A possible explanation is that, even when a solute is large and induces drying in its vicinity, the long-wavelength field will always penetrate the excluded volume of the solute to some extent. Consequently, N_v and σ_v are always nonzero and there is some Gaussian contribution to the energy. The results in figure 2.5, and to a lesser extent those in figure 2.3, suggest that an accurate treatment of this contribution is necessary to produce quantitatively accurate results in the large- R regime.

Whether the simple model in equation 2.12 is good enough, or a more accurate treatment, like the one presented here, is necessary, depends greatly on the particular application. The simplicity of 2.12 makes it computationally much easier to simulate, owing to the lack of any explicit coupling between cells that are not nearest neighbors. Furthermore, the more accurate small-lengthscale treatment, for all its merit, relies on a discretization of the underlying field theory that introduces lattice artifacts that are challenging to correct. For static applications, like computing solvation free energies or computing potentials of mean force between multiple hydrophobic solutes, these artifacts are unimportant. For dynamic

applications, however, they tend to pin solutes at points on the lattice, an effect that can complicate interpretation of the results. In future work, it might prove fruitful to explore whether it is possible to construct a new model that combines the best of these qualities: minimal lattice artifacts, computational convenience, and quantitative accuracy.

Chapter 3

Signatures of glasses formed by cooling

3.1 Introduction

A glass may be described colloquially as a liquid that has been cooled until it hardens into an amorphous solid. In contrast to most solids, which order their constituent particles into regular patterns, glasses lack any clear structural order at the molecular level. If one were to compare the arrangement of particles in a glass with that in a liquid, it would be difficult to discern the difference. Snapshots of molecular simulations of glasses reveal no obvious change in the structural ordering of particles as the temperature is lowered below T_g , the glass transition temperature.

Dynamics, rather than structure, distinguishes a glass from its liquid precursor. As a glass-forming liquid is cooled to very low temperatures, the dynamics of its individual particles slows considerably, until, near and below T_g , most particles become trapped near their initial positions and can do little more than vibrate in local potential energy basins. This dynamical slowing down is associated with rapidly growing structural relaxation times, reaching the second or minute timescale more typical of everyday life than the microscopic rearrangements of atoms and molecules.

Over timescales typical of experimental observation, such a material no longer appears to flow like a normal liquid, instead possessing an apparent rigidity more typically associated with crystalline solids. The solid-like quality of a glass is thus a consequence of rapidly growing timescales of relaxation in the liquid, which eventually exceed the finite timescale of observation, rather than due to a thermodynamic order-disorder transition. The challenge of understanding glasses and how they form is therefore the challenge of understanding the dynamical behavior of liquids cooled to very cold temperatures.

3.1.1 Super-Arrhenius behavior of supercooled liquids

The dynamics of supercooled liquids — that is, liquids cooled to low temperatures, which we define as temperatures well below a material-dependent onset temperature T_o — is markedly different from the dynamics of liquids at high temperatures ($T > T_o$). Relaxation in high-temperature liquids obeys the Arrhenius form,

$$\tau(T) = \tau_0 e^{E/T}, \quad (3.1)$$

where τ is the relaxation time of the liquid at temperature T , τ_0 is a microscopic timescale, and E is the activation energy for relaxation. In the simplest physical picture corresponding to equation 3.1, the liquid must overcome an energy barrier E in order to relax. In this picture, τ_0 is the fundamental timescale for liquid dynamics, and may be interpreted as the typical time between successive attempts to overcome the barrier.¹ Crucially, for $T > T_o$, the energy barrier E does not depend on temperature and, consequently, growth of the relaxation time with inverse temperature is muted.² In particular, $\ln \tau$ grows linearly with $\beta = 1/T$.

By contrast, relaxation in supercooled liquids is super-Arrhenius: the growth of $\ln \tau$ with β is faster than linear. Although widespread agreement on the exact functional form of the dependence of $\ln \tau$ on β is lacking, a vast array of experimental and simulation data strongly suggests that the proper form is parabolic, so that $\ln \tau$ grows as β^2 . The natural consequence of this parabolic law is that, even at finite temperatures (but temperatures still well below T_o , above which the Arrhenius form in equation 3.1 holds), the relaxation time may become extremely large. In the barrier-crossing picture presented above, a parabolic law implies a temperature-dependent energy barrier $E \propto T^{-1}$. We will see later that this increasing energy barrier corresponds to a length scale that increases as temperature is lowered.

3.1.2 Glasses are liquids out of equilibrium

Consider what happens when, as in the laboratory preparation of a glass, we cool a liquid initially in equilibrium at a high temperature to some much lower temperature³ at a finite rate ν_c . The relaxation time sets the timescale of the system's response to external perturbation, and as T decreases, the time for the system to equilibrate to the changing temperature increases dramatically. Eventually we reach a point at which the system responds too slowly

¹Alternatively, and relatedly, we may think of τ_0 as a commitment or plateau time, which is the time it takes the system to complete the relaxation process — i.e., to commit to the new basin — once it has made it to the top of the barrier.

²The choice of “muted” to describe the exponential growth of τ with $\beta = 1/T$ might seem inappropriate, but if equation 3.1 were to hold at all temperatures, a glass transition would never occur.

³Unless otherwise specified, the *temperature of the system* always refers to the temperature of a heat bath coupled to the system. We call this, interchangeably, the bath temperature or working temperature. The distinction between the working temperature and the internal temperature of the system is important out of equilibrium, when the temperature associated with the system's configuration may differ from the temperature of the bath. The latter is always under control, whereas the former — the so-called *fictive temperature*, T_f — is determined by the system's physical response to external manipulation.




Figure 3.1: Rough illustration of the crossover from super-Arrhenius to Arrhenius below the glass transition temperature. The black dashed line corresponds to the equilibrium relaxation time, while the solid colored lines correspond to relaxation times for systems cooled at finite rates. β_o denotes the inverse onset temperature: $\beta_o \equiv 1/T_o$. For slower cooling rates, the crossover back to Arrhenius relaxation occurs at lower temperatures, and the resulting slope is steeper due to the larger entrenched energy barrier $E(T_g)$.

to equilibrate on the timescale, roughly $|\nu_c|^{-1}$, of significant changes in temperature. That is, the system falls out of equilibrium and its configurations are no longer representative of the equilibrium ensemble at the working temperature T . In what follows, we define the glass transition temperature T_g as the temperature at which the system falls out of equilibrium. Specifically, T_g is the temperature that satisfies the approximate relation⁴

$$\frac{d\tau}{dT} \approx \frac{1}{\nu_c}. \quad (3.2)$$

From equation 3.2 and the physical argument that precedes it, it is clear that T_g is a function of the cooling rate ν_c , and that a slower cooling rate produces a lower T_g — the liquid remains in equilibrium to lower temperatures.

Because the system is unable to equilibrate for $T < T_g$, the system is essentially locked into its configuration at T_g , and therefore the energy barrier to structural relaxation is fixed

⁴We expect this relation to hold with a possibly rate-dependent proportionality constant. In section 3.2.2 we explore the validity of this relation for the East model when it is forced out of equilibrium by cooling.

at $E(T_g)$. Consequently, the relaxation time becomes Arrhenius again for $T < T_g$. Figure 3.1 shows this cooling-rate-dependent crossover for a material with energy scale J .

The glass transition is thus an out-of-equilibrium phenomenon, made possible by a relaxation process that becomes increasingly frustrated as temperature is lowered. Understanding the microscopic nature and origin of this frustration in supercooled liquids is key to understanding how they transform into glasses.

3.2 Kinetically constrained models of glass formers

A popular class of models for understanding the nature of dynamics in glass-forming liquids are the so-called *kinetically constrained models (KCMs)*, simple non-interacting lattice models with constrained dynamics. In this section, we review the most popular such model, the East model [12]. Despite its simplicity, the East model has highly nontrivial, extraordinarily rich dynamics at low temperatures, sharing many characteristics with deeply supercooled liquids. This combination of simplicity and physical richness makes it ideally suited for studying the glass transition.

3.2.1 Introduction to the East model

The East model consists of sites (or spins) on a one-dimensional lattice, which may be excited ($n = 1$) or unexcited ($n = 0$). An external field (or chemical potential) favors unexcited sites, with excited sites (“excitations”) costing energy J_0 . The Hamiltonian is then given by

$$H = \sum_i n_i, \quad (3.3)$$

where the zero of energy corresponds to a system containing no excitations, and H is written in units of J_0 , the only energy scale in the system. Due to the lack of spin-spin interactions in the Hamiltonian, the thermodynamics of the East model is trivial and can be described completely by the density or concentration of excitations as a function of temperature:

$$c(T) = \frac{1}{1 + e^{1/T}}. \quad (3.4)$$

Dynamics in the East model is facilitated by excitations. Specifically, an excitation facilitates changes of state in the site immediately to its right, so that the presence of an excitation at site i allows for the birth or death of an excitation at site $i + 1$.⁵

⁵This is why it’s called the East model. Note that the choice of direction is irrelevant and we could equally well choose the direction of facilitation to be left. Although this might technically be called the West model, it is identical to the East model. The key is that facilitation only happens in one direction, so that excitations may not facilitate their own diffusion through the lattice, as they do in the one-spin facilitated Fredrickson-Andersen (FA) model, a KCM with qualitatively different dynamics than the East model.

the leftmost excitation facilitating the creation of new excitations between them.⁶

Naively, we might expect this relaxation process to consist of the creation of an unbroken string of excitations that spans the domain, which may then relax as excitations are destroyed from right to left. This process is possible but highly unlikely for the relatively large domains typical at low temperatures, since it represents the maximum-energy pathway to relaxation, with the barrier scaling linearly with the domain size ℓ .⁷

It is clear that lower-energy pathways exist.⁸ The more realistic process depicted in figure 3.2, for example, lowers the energy at intermediate stages of relaxation by breaking the string of excitations, leaving an isolated “spring-board” excitation in the middle of the domain (at $x = 5$) that may facilitate the relaxation of the smaller subdomain that remains to its right. Once this subdomain has finished relaxing, the intermediate excitation must be destroyed via relaxation of the subdomain to its left. The end result is a relaxation process characterized by a bubble-like structure in spacetime, with the intermediate excitation forming the eastern boundary of this bubble.

Spacetime bubbles in the East model, noticeable in equilibrium simulations, are thus associated with metastable intermediates within a larger⁹ relaxation process. They form as the system attempts to lower its energy at intermediate stages of relaxation. As a consequence of their presence in the dynamics, we expect the energy barrier to relaxation to scale more favorably with ℓ than the linear relation associated with the naive process. To find how $E(\ell)$ scales with ℓ , note that figure 3.2 implies a recursive relationship in which the original 4-site domain is decomposed into two 2-site domains, at the cost of inserting a single excitation. The energy barrier may thus be written as

$$E(\ell) = 1 + E(\ell/2).$$

The “1” on the right-hand side is the energetic cost of the intermediate excitation, while the term $E(\ell/2)$ represents the division of the original problem in half.¹⁰ Decomposing a second time, we obtain

$$E(\ell) = 1 + (1 + E(\ell/4)) = 2 + E(\ell/4).$$

⁶If the boundaries of this system are non-periodic, the leftmost excitation will never de-excite since there are no excitations to its left. If the boundaries are periodic, the rightmost excitation could facilitate the destruction of the leftmost excitation through the same process we describe here, but the domain separating them across the periodic boundary is larger than the one that does not cross the boundary, so this is unlikely to happen.

⁷More accurately, $E(\ell) = \ell - 1$ for this process. The domain size ℓ includes the excitation at the east end of the domain.

⁸The arguments presented in this section, and illustrated in figure 3.2, were made previously by Sollich and Evans in reference [34].

⁹By *larger*, we mean at longer length- and timescales.

¹⁰Technically, it should read $E(\ell) = 1 + \max[E_L, E_R]$, where $E_{L/R}$ is the energy barrier to relaxing the subdomain to the left/right of the intermediate excitation. However, this is only relevant if ℓ is not divisible by two. We assume here, without loss of generality, that ℓ is a power of two, so that all subsequent divisions result in equal-sized left and right domains.

From the pattern so far, it is clear that repeating k times yields

$$E(\ell) = k + E(\ell/2^k).$$

Defining $n = \log_2 \ell$ and letting $k = n - 1$ gives $E(\ell) = (n - 1) + E(2) = (n - 1) + 1 = n$. Thus, the energy barrier for the minimum-energy relaxation pathway scales logarithmically with domain size, $E_m(\ell) = \log_2 \ell$, or in terms of the natural logarithm,

$$E_m(\ell) = (\ln 2)^{-1} \ln \ell. \quad (3.5)$$

Because there are many possible relaxation pathways with activation energies comparable to this minimum energy, the effective free energy barrier to relaxation is quantitatively different than this result. To account for the entropy of these pathways, we introduce a proportionality constant γ and generalize equation 3.5 to read

$$E(\ell) = \gamma \ln \ell. \quad (3.6)$$

If only the minimum-energy pathway is possible, $\gamma = (\ln 2)^{-1}$. In reality, with many pathways possible, the real value of γ for the East model is somewhat lower than this value.¹¹

Because the density of excitations decreases as temperature is lowered, the mean distance between them, $\ell_{\text{eq}}(T) = 1/c(T)$, increases. Thus, the barrier to relaxation grows with inverse temperature. Combining equation 3.6 with equation 3.4 gives

$$E(T) = \gamma \ln (1 + e^{1/T}).$$

For $T \ll 1$, this reduces to γ/T or $\gamma\beta$. Because the energy barrier increases linearly with inverse temperature, the East model obeys the parabolic law, with the relaxation time growing dramatically as temperature is lowered, but always finite for $T > 0$.

To confirm that the East model follows the parabolic law, we perform Monte Carlo simulations of the East model and compute the relaxation time over a wide range of temperatures. To do so, we introduce the persistence function

$$p_i(t) \equiv \prod_{t'=0}^t \delta[n_i(t') - n_i(0)], \quad (3.7)$$

which is 1 if site i has not changed its state up to time t , and 0 otherwise. Note that $p_i(t)$ is zero if the state of i differs from its initial state $n_i(0)$ at *any* point in the time interval, even if it changes back to the initial state at some point. This is intentional, as $p_i(t)$ is meant to indicate whether site i has been facilitated within time t .¹² From these considerations, it is

¹¹Chebloun, Faggionato, and Martinelli bound this value from below by $(2 \ln 2)^{-1}$. The actual value seems to lie somewhere in between: $(2 \ln 2)^{-1} < \gamma < (\ln 2)^{-1}$.

¹²Equivalently, we could say that $p_i(t)$ measures whether the initial barrier to relaxing site i has vanished within time t . If site i changes state twice within time t , this initial barrier has clearly vanished even though $n_i(t)$ is the same as $n_i(0)$.




Figure 3.3: Equilibrium relaxation times for the one-dimensional East model. The dashed red line is a parabolic fit to the computed data (red circles): $y = 0.923x^2 + 1.211x + 1.910$.

clear that $p_i(0) = 1$ and $p_i(\infty) = 0$, the former by definition and the latter because all sites will eventually have an opportunity to relax.¹³ The mean value of the persistence function,

$$\bar{p}(t) = \langle p_i(t) \rangle, \quad (3.8)$$

where the angle brackets denote an equilibrium average, gives the fraction of sites that have relaxed within time t . For the East model, we define the equilibrium relaxation time $\tau_{\text{eq}}(T)$ at temperature T to be the $1/e$ decay time of this mean persistence function.¹⁴

In practice, we perform the equilibrium average in equation 3.8 by preparing M equilibrium configurations, propagating the stochastic dynamics of each one up to time t , and then averaging the persistence function over all sites in all replicas. Mathematically, we compute

$$p_{\text{avg}}(t) = (ML)^{-1} \sum_{j=1}^M \sum_{i=1}^L p_i^{(j)}(t),$$

where L is the system size and $p_i^{(j)}(t)$ is the persistence function of site i in replica j . As $ML \rightarrow \infty$, the function above approaches the equilibrium average in equation 3.8, provided

¹³This assumes a system with periodic boundary conditions. The system in figure 3.2 obviously cannot relax any spin for $x < 4$ if the system is non-periodic, since there are no excitations to the left of the one at $x = 3$.

¹⁴Other definitions are certainly possible and give similar results. The definition given here is convenient to compute in computer simulations.




Figure 3.4: Relaxation times for the $d = 1$ East model when cooled at varying rates. The dashed red line is the parabolic fit from figure 3.3. Note that the relaxation times follow the equilibrium scaling to lower temperatures when the cooling rate is slower.

L is large enough to accommodate the typical distance between excitations, i.e., $L \gg \ell_{\text{eq}}(T)$. Typically we fix L at a value satisfying this condition and then increase M until we achieve satisfactory statistics.

Figure 3.3 shows the results of our numerical simulations, along with a parabolic fit to the data. The success of the fit strongly suggests that higher-order terms are negligible, so that the East model appears to follow the parabolic law.¹⁵ Furthermore, the measured curvature $J^2 = \gamma J_0^2$ implies that $\gamma \approx 0.9$, which falls within the range $(2 \ln 2)^{-1} < \gamma < (\ln 2)^{-1}$ that we expected before.

3.2.2 The East model out of equilibrium

Recall that equation 3.2 gives an approximate definition of the glass transition temperature, T_g , as the temperature satisfying

$$\frac{d\tau}{dT} \approx \frac{1}{\nu_c}.$$

In constructing figure 3.1, which illustrates the crossover from super-Arrhenius to Arrhenius behavior at the glass transition, we assume this relation holds with unit proportionality

¹⁵These results do not rule out the possibility of higher-order terms becoming important at even lower temperatures than the ones considered here, which are inaccessible to our simulations.

ν_c	T_g^a	T_g^b	C
10^{-3}	0.764		
10^{-4}	0.586		
10^{-5}	0.481		
10^{-6}	0.412	0.523	0.034
10^{-7}	0.363	0.407	0.123
10^{-8}	0.327	0.354	0.170
10^{-9}	0.299	0.320	0.168

Table 3.1: Glass transition temperatures for the East model when cooled at varying rates, computed using different methods. The second column lists T_g^a , the temperature that satisfies equation 3.2. The third column gives T_g^b , which is computed by fitting relaxation time data at low temperatures to an Arrhenius form. (See main text for details and discussion.) The last column gives the value of the proportionality constant in equation 3.9 implied by the value of T_g^b in the second column.

constant. In this section, we denote the glass transition temperature implied by this relation (with unit proportionality constant) as T_g^a in order to distinguish it from other definitions of T_g . Table 3.2.2 lists the values of T_g^a for the East model for a wide range of cooling rates.

In general, we expect the implied proportionality constant in equation 3.2 to assume a non-unit value, and to display some dependence on the cooling rate. That is, we expect T_g to satisfy the more general relation

$$\frac{d\tau}{dT} = C(\nu_c) \frac{1}{\nu_c}, \quad (3.9)$$

where $C(\nu_c)$ is the rate-dependent proportionality constant. To determine this proportionality constant for the East model, we compute T_g from cooling simulations. For each cooling rate listed in table 3.2.2, we compute the relaxation time over a range of temperatures sufficient to capture the super-Arrhenius-to-Arrhenius crossover. The resulting relaxation time curves are shown in figure 3.4. Not surprisingly, the expected equilibrium scaling (red dashed line) holds to lower temperatures for slower cooling rates.

From the data shown in figure 3.4, we can define a glass transition temperature, T_g^b , by fitting an Arrhenius, or linear, form to the low-temperature data.¹⁶ Figure 3.5(a) illustrates this procedure for the fastest and slowest cooling rates under consideration. The intersection of the line of best fit (blue dashed line) with the equilibrium parabola (gray line) defines T_g^b . Interestingly, for the three fastest cooling rates (10^{-3} to 10^{-5}), there is no such intersection and T_g^b is undefined.

¹⁶We try to limit this fit to temperatures beyond the crossover range. For the five fastest rates listed in table 3.2.2 (10^{-3} to 10^{-7}), we include all data points satisfying $\beta > 3$. For $\nu_c = 10^{-8}$, we only include points with $\beta > 3.2$, and for $\nu_c = 10^{-9}$ we limit ourselves to $\beta > 3.4$.




Figure 3.5: (a) Illustration of the procedure for determining the out-of-equilibrium energy barrier in the East model. The relaxation time data at low temperatures is fit to an Arrhenius (linear) form. The intersection of the resulting line with the equilibrium parabola gives T_g , while the slope of this line gives the out-of-equilibrium energy barrier. For the three fastest cooling rates, there is no intersection between the best-fit line and the equilibrium curve, so T_g is not defined. (b) Comparison of the energy barrier computed by various methods. See the main text for descriptions of these methods.

Table 3.2.2 lists the values of T_g^b we computed from the data. Note that $T_g^b > T_g^a$ for the cooling rates with well-defined T_g^b — the relaxation time data suggest higher glass transition temperatures than predicted by equation 3.2. With T_g^b computed from actual data, we can calculate the proportionality constant in equation 3.9. A simple rearrangement of that equation yields

$$C(\nu_c) = \nu_c \left(\frac{d\tau}{dT} \bigg|_{T=T_g^b} \right).$$

The resulting values of the proportionality constant are listed in table 3.2.2 for each cooling protocol with a well-defined T_g^b .

The slope of the best-fit line provides a measure of the effective energy barrier embedded in the East model after it falls out of equilibrium. Figure 3.5(b) compares this result with several other estimates of the energy barrier. Results in solid red and solid black show the energy barrier implied by the hierarchical relationship that we derived earlier, equation 3.6,

$$E(\ell) = \gamma \ln \ell,$$

where $\gamma = 0.923$ is taken from the curvature of the equilibrium relaxation time data in figure 3.3. To estimate the barrier, we need to know the value of ℓ in this equation. The solid red data show the results of using

$$\ell_{\text{ne}} = 1/c(T_g^a) = 1 + e^{1/T_g^a},$$

the equilibrium length at the glass transition temperature. The solid black data show the results of using the mean inter-excitation distance at the end of the cooling protocol ($T = 0$). This lengthscale is longer than $1/c(T_g)$ for any reasonable estimate of T_g , so it yields a higher energy barrier, as shown in the figure. Both sets of results severely underestimate the energy barrier implied by the slope from figure 3.5(a).

One reason for the large discrepancy in figure 3.5(b) is that the equilibrium curve, which is described numerically by the equation

$$\ln \tau = 0.923\beta^2 + 1.211\beta + 1.910,$$

has a term linear in β . This term has no clear physical meaning and is not predicted by the physical arguments in the preceding section. It can be interpreted as a kind of “mean field” contribution¹⁷ or as a consequence of a negative onset temperature.^{18,19} Either way, it reflects the fact that the minimum of the parabola does not coincide with $\beta = 0$.

We compensate for this linear term in one of two ways. In the first, we assume the linear term is analogous to the mean-field contribution in real liquids and, consequently, adds a constant energy barrier at all temperatures that is distinct from the hierarchical energy barrier in equation 3.6. Thus, to obtain the energy barrier we must subtract this contribution from the slope of the best-fit line in figure 3.5(a). This result is the dashed blue line in figure 3.5(b). In the second way, we replace β by $\beta - \beta_0$ (i.e., we measure all temperatures relative to an onset temperature $\beta_0 \neq 0$). This result is shown by the dashed red line in figure 3.5(b). Despite these attempted corrections, the energy barrier implied by the slope of the best-fit line still noticeably exceeds the predictions of the other methods.²⁰

Especially perplexing is the lack of agreement between the blue lines (solid and dashed) and the black line, the latter of which constitutes a direct measurement of the inter-excitation lengthscale that emerges far out of equilibrium. The failure of this direct measurement to accurately predict the out-of-equilibrium slope could suggest that the relationship between

¹⁷In real liquids this is a remnant of the high-temperature ($T > T_o$) Arrhenius behavior. It can be interpreted as the energy barrier that must be overcome for a particle to make a displacement even when that particle is located within a soft spot. It is not clear why the East model should have such a contribution.

¹⁸In the temperature-dependent relations we have derived, this would mean measuring the temperature relative to onset, so that β is replaced by $\beta - \beta_0$. It is not clear why $\beta_0 \neq 0$ for the East model, but it may reflect inadequacies in our mapping between the inter-excitation lengthscale ℓ and the relaxation time τ as we have defined it, inadequacies which are numerically rectified by introducing a nonzero reference point for the temperature.

¹⁹It could also be some combination of both, though in this section we assume one extreme or the other.

²⁰Note that we cannot directly compare the dashed lines in figure 3.5(b), since they result from mutually exclusive interpretations of the linear term.

the effective energy barrier and the distribution of lengths in the material may not be adequately captured by the simple relation $E(\ell) = \gamma \ln \ell$, with its single energy and lengthscale. Later, we will see that an East model glass is characterized by a very different distribution of inter-excitation distances than its equilibrium counterpart. The value of γ , which is related to the entropy of relaxation pathways and was computed for systems at equilibrium, may have a nontrivial dependence on the distribution of lengths and could therefore be different for materials that are out of equilibrium.

Another possible explanation for the observed discrepancies is that our definition of the relaxation time allows for significant aging of the material during the measurement. Since a large majority of spins in the system (roughly 63%) will relax in time τ , a system initially out of equilibrium will make significant progress towards equilibrium during the observation. Consequently, the relaxation time that we observe likely reflects a changing energy barrier. In particular, because the out-of-equilibrium configurations contain an excess of excitations relative to the temperature of the bath, this energy barrier will tend to increase during the observation, resulting in an out-of-equilibrium slope that is larger than expected. This is consistent with our observations.

3.2.3 Correspondence of the East model with atomistic models of glass formers

As described in the previous section, a consequence of the East model’s simple dynamical rules is a free energy barrier to relaxation that grows linearly with β . The East model thus offers a physical explanation for the parabolic law obeyed by real materials and atomistic models of glass formers. The question naturally arises whether the directional facilitation of the East model emerges in more realistic models.

Recent work by Keys et al. examining a wide range of atomistic models of glass formers suggests that directional facilitation becomes the dominant mode of transport at temperatures well below onset, with the movement of a particle highly correlated with subsequent motion, at later times, of the particles left in its wake. Movies of supercooled liquid dynamics illustrate the point dramatically, as motion emanates from mobile regions of particles in a quasi-one-dimensional, string-like fashion. Meanwhile, particle motion on short timescales²¹ is confined to localized “soft spots” no larger than a few particle diameters. These soft spots are analogues of the discrete excitations in the East model.

For each atomistic model, it is possible to calculate analogues of the energy parameter J_0 and entropy parameter γ of the East model. Keys et al. succeed in showing that relaxation times, when suitably scaled by these parameters below the onset temperature, collapse onto a universal parabolic form. Taken together, these results suggest that the East model is the simplest possible model of glassy dynamics.

²¹By “short” we mean the time it takes to make a displacement on the order of a particle diameter σ .

3.2.4 Formation of East model glasses

With the East model introduced and its relevance to supercooled liquids established, we now consider what happens when we apply a cooling protocol to the East model. Previously we noted that a supercooled liquid, upon cooling, will eventually fall out of equilibrium, as the system's rapidly growing relaxation time renders it unable to keep up with changes in the temperature. With the East model, we can now understand the microscopic nature of this process.

As an East lattice is cooled, it adjusts to the decreasing temperature by eliminating excess excitations so that the concentration matches the equilibrium value in equation 3.4. Due to the kinetic constraint, eliminating an excitation requires relaxing the domain of unoccupied sites to its left. As temperature decreases, the typical distance between excitations — and by extension the typical size of domains that must relax to eliminate excitations — increases. At T_g , this domain size is too large to be relaxed on the timescale implied by the cooling rate,²² and the remaining excitations are locked into place. As temperature continues to decrease, the concentration plateaus, even as thermodynamics favors fewer excitations. Unable to rid itself of excess excitations, the system falls further out of equilibrium.

Figure 3.6 shows the concentration falling out of equilibrium as an East lattice is cooled at varying rates. Initially, the concentration follows the equilibrium line very closely,²³ but at sufficiently low temperatures, the remaining excitations are too far apart to relax in the given time. As a result, the concentration plateaus to a protocol-dependent value. Cooling more slowly allows the lattice to relax larger domains, so that more excitations are destroyed before the system falls out of equilibrium.

As we discuss at length later, in atomistic models the excitations are not observable except via their dynamical manifestations as particle displacements. While figure 3.6 presents a compelling picture of the glass transition, it is not something we can compute directly for atomistic systems. For this reason, we consider the analogy to particle displacements in the East model, spin flips, and instead of considering the concentration of excitations, we consider the *rate* of spin flips. We define a rate function for the East model that makes use of the persistence function defined in equation 3.7. Recall that the persistence function for site i , $p_i(t)$, the definition of which we reproduce here,

$$p_i(t) \equiv \prod_{t'=0}^t \delta[n_i(t') - n_i(0)],$$

measures whether, after an elapsed time t , site i has changed its state at least once. In terms of the mean persistence function, $\bar{p}(t)$, we define the rate as

$$r(T; t_{\text{obs}}) = \frac{1}{t_{\text{obs}}} (1 - \bar{p}(t_{\text{obs}})). \quad (3.10)$$

²²We could have equally well made this argument in terms of a growing energy barrier instead of a growing lengthscale. In the East model, the two are intimately related by the physical arguments provided in section 3.2 and mathematically expressed in equation 3.6.

²³The plotting is done so that the equilibrium relation in equation 3.4 gives a straight line with unit slope.




Figure 3.6: Concentration of excitations c versus working temperature T in the East model for various cooling rates. The straight line with unit slope (black, dashed) is the equilibrium result (achieved only for $|\nu_c| \rightarrow \infty$). The non-equilibrium curves display the same pattern, following the equilibrium line until plateauing to a constant value in the vicinity of the glass transition temperature T_g . Cooling more slowly depresses T_g , so that the concentration curve follows the equilibrium line to lower temperatures and plateaus at a lower value of c .

This is often called the *kink rate*.

The simplicity of the East model allows us to understand how $r(T)$ should behave under certain circumstances. Because spin flips only occur at facilitated sites next to existing excitations, for sufficiently short observation times r should be solely determined by the distribution of excitations at time zero rather than any ensuing dynamics. At low temperatures, excitations are sparse, so that facilitated sites are almost always unexcited. Thus, in order to flip, these sites must overcome the cost of creating an excitation. The rate at which a facilitated site flips must therefore decline as $\exp(-1/T)$. The total rate is then a product of this term and the concentration of excitations (i.e., facilitated sites),

$$r(T) \approx c(T)e^{-1/T}.$$

For systems in equilibrium, $c(T) \approx e^{-1/T}$ at low temperatures, so that

$$r_{\text{eq}}(T) \approx e^{-2/T}.$$

Thus, in equilibrium at low temperatures, $\ln(r)$ behaves as a straight line with a slope of 2. Figure 3.7 confirms this, and demonstrates that the result is independent of observation




Figure 3.7: Left: Rate of spin flips $r \equiv t_{\text{obs}}^{-1} L^{-1} \sum_{i=1}^L (1 - p_i(t_{\text{obs}}))$, where p_i is the persistence function defined in equation 3.7, for the East model in equilibrium. Different colors correspond to different choices of the observation time, which is in units of attempted spin flips per site. The dashed line is an approximate fit to the low temperature data, showing that the relevant energy scale is $2J_0$. Right: Rate of spin flips r for the East model cooled at constant rates $\nu_c = 10^{-3}$ (red), $\nu_c = 10^{-6}$ (orange), and $\nu_c = 10^{-9}$ (green). Note that the out-of-equilibrium curves approach a line with unit slope. The observation time was chosen to be $t_{\text{obs}} = 1$.

time, provided that the observation time does not exceed $t_{\text{obs}} = 1$.²⁴ Interestingly, at high temperatures, r deviates substantially from linear behavior, in contrast to $c(T)$, which remains perfectly linear at all temperatures. This deviation is due to the high concentration of excitations at high temperatures. Although r is proportional to c at all temperatures,²⁵ at high temperatures facilitated sites are very likely to be excited themselves, and therefore change state by de-exciting, a process with no energy barrier. Consequently, the slope at high temperature should be less than the low-temperature estimate of 2, which assumes a unit energy barrier to flipping. This smaller slope at high T can be clearly seen in figure 3.7.

In contrast to the equilibrium result, for systems cooled at constant rate ν_c , the concentration plateaus for $T < T_g$ at a rate-dependent value $c_0(\nu_c)$, so that

$$r_{\text{ne}}(T) \approx c_0(\nu_c) e^{-1/T}.$$

²⁴The chosen time unit here is Monte Carlo sweeps, or — equivalently — trial moves per site.

²⁵Provided that the observation time is sufficiently short.

Consequently, $\ln(r)$ approaches a straight line with unit slope for systems cooled below T_g , which can be clearly seen in figure 3.7.

Thus, with the East model we see that the source of frustration in supercooled liquids — which leads them to fall out of equilibrium as the glass transition temperature is approached from above — is a growing lengthscale between soft spots, which cost energy and are therefore annihilated as temperature decreases.

3.2.5 Inter-excitation correlations in East model glasses

The arguments offered in the previous section have focused on the “typical” or average length between excitations, $\ell_{\text{eq}}(T)$, which grows as temperature is lowered and excitations become less dense, and which becomes fixed at $\ell_{\text{eq}}(T_g)$ once $T < T_g$. This might lead one to assume that a glass is only out of equilibrium because this lengthscale does not match the expected equilibrium lengthscale at the instantaneous bath temperature — and that equilibrium can be trivially re-established by changing the bath temperature back to T_g .

Another way to articulate this assumption is with the concept of a *fictive temperature* T_f . To define the fictive temperature, we first note that a system in equilibrium at temperature T produces an ensemble of configurations; that is, there is a mapping from the bath temperature T to an ensemble of configurations. The fictive temperature is the reverse mapping: given an ensemble of configurations, the fictive temperature T_f is the bath temperature that would produce such an ensemble. Clearly not all ensembles have a fictive temperature, since not all ensembles can be produced by thermal equilibration with a heat bath at fixed temperature. If a system has a fictive temperature T_f , it may still be out of equilibrium if $T_f \neq T$, but it may immediately re-establish equilibrium by changing the bath temperature to T_f . On the other hand, systems lacking a fictive temperature must age before returning to equilibrium, regardless of the temperature.

The assumption articulated at the beginning of this section is that the ensemble of configurations produced by cooling the East model, or a supercooled liquid, below T_g belongs to the subset of non-equilibrium ensembles with a well-defined fictive temperature that does not match the bath temperature. Consequently, if the bath temperature is changed to match T_f , the system will immediately re-establish equilibrium without aging.

This simplistic picture is incorrect. Non-equilibrium ensembles produced by cooling have no equilibrium analogues, and thus lack a fictive temperature as defined above.²⁶ To illustrate this point, figure 3.8 compares configurations of the East model taken from a non-equilibrium ensemble — prepared by cooling from $T = 10$ down to zero in 10^4 Monte Carlo sweeps — and the equilibrium ensemble with the same average inter-excitation length. The configurations typical of each ensemble are noticeably different, with the configuration prepared by cooling

²⁶Keys, Garrahan, and Chandler generalize the fictive temperature concept by spatially resolving it. Each site i is characterized by its own fictive temperature $T_f(i)$, which is defined as the bath temperature at the last time the site relaxed. The simpler fictive temperature concept — with one system-wide temperature T_f — emerges from this more general definition in the mean-field approximation, i.e., when $T_f(i)$ is approximated as spatially homogeneous.




Figure 3.8: East model configurations drawn from equilibrium (A) and non-equilibrium (B) ensembles with the same average length between excitations. Each configuration is an 80-site snapshot of a 1000-site system. The non-equilibrium ensemble is prepared by cooling $T = 10$ to $T = 0$ in 10^4 sweeps. The average inter-excitation distance ℓ_{avg} is shown for comparison. Whereas the equilibrium configuration contains many small and large domains, relative to ℓ_{avg} , the non-equilibrium configuration displays a relatively even spacing between excitations, suggesting that the distribution of lengthscales in an East model glass is somewhat narrowly distributed around ℓ_{avg} .

exhibiting a relatively even spacing between excitations, while the equilibrium configuration features many more small ($\ell \ll \ell_{\text{avg}}$) and large ($\ell \gg \ell_{\text{avg}}$) domains.

The non-equilibrium structure that forms below T_g is thus distinct from the equilibrium structure that prevails above T_g , and a glass cannot be regarded as an equilibrium liquid with an “out-of-sync” fictive temperature. Rather, it has an entirely different structure altogether — one characterized by a completely different *distribution* of lengthscales between excitations than an equilibrium liquid.

To explore how this difference arises, we first consider the distribution of lengthscales — i.e., the distances between excitations, or domain sizes — in the East model at equilibrium. This distribution is the probability $P(\ell)$ that an excitation picked at random is separated from its nearest neighbor excitation by a distance ℓ .²⁷ The lack of site-site coupling in the Hamiltonian, equation 3.3, means this probability can be written as the product of ℓ terms,

$$P(\ell) = p_0(i+1)p_0(i+2) \cdots p_0(i+\ell-1)p_1(\ell),$$

where i is the site of the randomly selected excitation, $p_0(j)$ is the probability that site j is unoccupied, and $p_1(j)$ is the probability that site j is occupied. This expresses $P(\ell)$ as the probability that there is an excitation ℓ sites away from the excitation at i and none at any of the $\ell - 1$ sites between them. The translational invariance of the Hamiltonian means that

²⁷Note that $P(\ell)$ is a conditional average that requiring an excitation at the origin. The corresponding *unconditional* average is performed by picking *any* random site in the system, regardless of occupancy, and computing the probability that the closest excitation to that site is ℓ . Because equation 3.3 lacks explicit site-site coupling, the conditional and unconditional averages are identical.

p_0 and p_1 are independent of position, so that $P(\ell) = p_0^{\ell-1} p_1$. It is obvious that $p_1 = c_{\text{eq}}$ and $p_0 = 1 - p_1 = 1 - c_{\text{eq}}$, where c_{eq} is the equilibrium density of excitations, so that for the East model in equilibrium, the distribution of lengthscales is geometric and given by

$$P(\ell) = c_{\text{eq}}(1 - c_{\text{eq}})^{\ell-1}.$$

Rewriting in terms of $\ell_{\text{eq}} = 1/c_{\text{eq}}$,

$$P(\ell) = \frac{1}{\ell_{\text{eq}}} \left(1 - \frac{1}{\ell_{\text{eq}}}\right)^{\ell-1}. \quad (3.11)$$

In the continuum limit, $\ell_{\text{eq}} \gg 1$, the lattice spacing becomes irrelevant compared to the typical inter-excitation distance, and the discrete geometric distribution approaches a continuous exponential distribution, $P(\ell) \rightarrow \ell_{\text{eq}}^{-1} \exp(-\ell/\ell_{\text{eq}})$.²⁸ This exponential distribution is characteristic of particles that are energetically uncoupled and whose positions are, consequently, completely uncorrelated.²⁹

With the equilibrium form of $P(\ell)$, we can rationalize some of the qualitative features of the equilibrium configuration in figure 3.8. One noticeable feature of this configuration is the abundance of very small domains, which can now be understood as a natural consequence of ideal-gas statistics. Because the exponential distribution is monotonically decreasing in ℓ , the smallest domain size ($\ell = 1$) is always most probable, no matter how large ℓ_{eq} may be. In the equilibrium East model, we always expect to find clusters of closely spaced excitations.

In contrast, the non-equilibrium configuration in figure 3.8 lacks inter-excitation domains that are significantly smaller than the average length. The reduction in the probability $P(\ell)$ for $\ell < \ell_{\text{avg}}$ implied by this lack of small domains suggests that excitations develop negative pair correlations out of equilibrium.

Previously, we considered how the concentration of excitations, $c(T)$, falls out of equilibrium upon cooling, as shown in figure 3.6, by focusing on the typical distance $\ell(T) = 1/c(T)$ between excitations, which increases as temperature is lowered. We argued that once this typical distance exceeds what can be relaxed in the finite time allowed by the cooling protocol, the system no longer adjusts to the decreasing temperature, the configuration becomes fixed, and the system is no longer in equilibrium. We later noted that this single-lengthscale picture implied that equilibrium could be trivially re-established by changing the bath temperature T to the system's fictive temperature T_f .

We can now understand the shortcomings of this simplified picture. Rather than a single, temperature-dependent length $\ell_{\text{eq}}(T)$, which eventually grows so large that the system falls out of equilibrium, an East lattice in equilibrium instead has a broad distribution of lengths, including an abundance of short lengths $\ell < \ell_{\text{eq}}(T)$. These short domains relax more quickly

²⁸Because this requires very low concentrations of excitations, this also corresponds to the low-temperature limit.

²⁹Although there are no *static* correlations in equilibrium, there are *dynamic* correlations arising from the kinetic constraint. We will see later that these dynamic correlations can lead to static correlations emerging out of equilibrium.




Figure 3.9: Inter-excitation correlations in East model glasses. Top: Excitation-excitation pair correlation functions for systems cooled to $T = 0$ at varying rates ν_c . On the right, the x -axis is scaled by $\langle\ell\rangle_{\nu_c}$, the average distance between excitations in the glass corresponding to cooling rate ν_c . Bottom: Distribution of distances between nearest-neighbor excitations, $P(\ell)$, for fast ($\nu_c = 10^{-3}$), moderate ($\nu_c = 10^{-6}$), and slow ($\nu_c = 10^{-9}$) cooling rates. On the right, scaling and shifting these distributions collapses these curves onto the same universal form. The dashed gray line, with slope -1 , is the equilibrium distribution.

(see equation 3.6) and therefore respond to changing temperature more quickly than longer domains comparable to or larger than the average length. As temperature is lowered and approaches T_g , larger domains will be first to fall out of equilibrium, followed later by shorter domains.³⁰ That is, different parts of the system will stop relaxing (their configurations fixed) at different temperatures.³¹

One consequence of this heterogeneity is that, even at temperatures well below T_g , small domains will continue to relax and adjust to the lower temperature by eliminating excitations. This explains why the curves in figure 3.6 plateau smoothly, rather than discontinuously at T_g . Even at $\beta > \beta_g$, excitations at the end of small domains may still be annihilated, so that concentration continues to decrease well after it deviates substantially from the equilibrium line.

Because nearby excitations continue to annihilate well below T_g , small inter-excitation distances, $\ell < \ell_{\text{eq}}(T_g)$, are depleted while large inter-excitation distances, $\ell > \ell_{\text{eq}}(T_g)$, remain. This suggests that $P(\ell)$ forms a peak at $\ell \neq 1$ (in marked contrast to the monotonic decay of the equilibrium distribution), and that very small values of ℓ become particularly improbable. This is consistent with the lack of small domains in the non-equilibrium configuration shown in figure 3.8.

These impressions are confirmed by the results in figure 3.9. The bottom pair of plots show $\ln P(\ell)$ versus ℓ for East lattices cooled to $T = 0$ at varying rates.³² For each cooling rate considered, a pronounced peak at $\ell_{\text{ne}} \neq 1$ forms, with probabilities at $\ell < \ell_{\text{ne}}$ substantially depleted. The location of this peak shifts to longer lengths for slower cooling rates, as larger domains are able to relax. The top two plots in figure 3.9, showing the corresponding excitation-excitation pair correlation functions, further supports our prior observation, from comparing the configurations in figure 3.8, that negative pair correlations between excitations accompany the glass transition in the East model.

Thus, the non-equilibrium structure that emerges upon cooling resembles nothing like an equilibrium structure. Consequently, changing the temperature of an East model glass with average inter-excitation distance ℓ_{avg} to the value T that satisfies $\ell_{\text{avg}} = \ell_{\text{eq}}(T)$ would *not* immediately re-establish equilibrium, since the *distribution* of distances is not the equilibrium distribution. Rather, the system will undergo an aging process in which key thermodynamic properties — such as the concentration of excitations or, equivalently, the energy — depend on time until the distribution of excitations becomes ideal-gas-like again.³³

³⁰It is for this reason that the glass transition is a broad crossover rather than a sharp transition.

³¹This is why the Keys, Garrahan, Chandler concept of fictive temperature as a spatially heterogeneous field is sensible.

³²At equilibrium, when ℓ is exponentially distributed, a logarithmic plot like this produces a straight line with slope $1/\ell_{\text{eq}}$. See the gray dashed line in the bottom right plot of figure 3.9.

³³In particular, the concentration of excitations will first *increase*, before falling back to its initial value.




Figure 3.10: Left and middle: Distribution of lengths $P(\ell)$ between kinks for various observation times. Right: Comparison of $P(\ell)$ for excitations (black) and for observed kinks with $t_{\text{obs}} = 1000$. All plots correspond to the non-equilibrium state point $(\nu_c, T) = (-10^{-9}, 0.20)$. Note that the distribution for $t_{\text{obs}} = 1$ (black/left) features many sharp peaks. These are finite size effects. At that temperature, for that observation time, the mean distance between observed kinks is roughly the same as the box size ($L = 1000$).

Excitations versus displacements

Although the emergence of correlations between excitations provide a distinct structural signature of the glass transition in the East model, observing the analogous transformation in an atomistic model, which we attempt to do in the next section, is nontrivial. Whereas excitations are inherent to the East model, they are emergent features of atomistic models of liquids at supercooled conditions. They correspond to configurations of particles that are locally unjammed and able to move, in intermittent fits, over lengthscales similar to the inter-particle spacing. Detecting these excitations, or “soft spots,” is challenging because a static order parameter that reliably distinguishes locally jammed versus unjammed regions of particles does not exist. Instead, we must rely on the dynamical manifestations of excitations: persistent particle motion over short timescales. That is, when we look at atomistic models, we can only look at the *displacements* of particles, made possible by the presence of underlying excitations, rather than the excitations themselves.

We examine how this difference, between excitations and their dynamical manifestations, changes the correlation functions we computed earlier. Specifically, we consider two ways of detecting excitations via dynamics: overlap and kinks. In the East model, a site i overlaps with its initial configuration at a later time t if $n_i(t) \neq n_i(0)$. With this idea in mind, we




Figure 3.11: Distribution of lengths $P(\ell)$ between sites that do not overlap with their initial state after some observation time, at the non-equilibrium state point $(\nu_c, T) = (-10^{-9}, 0.20)$. The observation time spans six orders of magnitude from the shortest t_{obs} (blue) to the longest (red), yet the distribution changes very little, in marked contrast with the same distributions calculated with kinks (rather than overlap) in figure 3.10.

can define an *overlap function*,

$$o_i(t) \equiv \delta[n_i(t) - n_i(0)], \quad (3.12)$$

where $\delta(x)$ is the Kronecker delta. The function $o_i(t)$ assumes a value of 1 if at time t site i overlaps with its initial configuration, and zero otherwise. Since only facilitated sites are able to change state, for short observation times sites that satisfy $n_i(t) \neq n_i(0)$ (or, equivalently, $o_i(t) = 0$) will be next to (or very close to) the excitations present in the system initially. That is, the set of sites satisfying the overlap condition will be strongly correlated with the initial locations of the excitations. Relatedly, a site has “kinked” if it has changed its state from its initial configuration *at any point* during the time interval of observation. That is, a site i has kinked in an observation time t if $p_i(t) = 0$, where $p_i(t)$ is the persistence function that we defined earlier in equation 3.7. Like overlap, kinks are highly correlated with the initial locations of excitations.

Our intuition might suggest that using kinks, or overlap, rather than excitations would not drastically change the distributions we have calculated, since kinks can only occur next to excitations. However, using kinks or overlap introduces a dependence on observation time, the effect of which is clearly illustrated for kinks in figure 3.10. For very short observation

times, the non-exponential distribution that is so clear for the bare excitations (black curve on the far right plot) is almost completely washed away. This result can be rationalized as follows: The effect of using sites where either $p_i(t) = 0$ or $o_i(t) = 0$ is to randomly subsample some fraction of the excitations, so that $P(\ell)$ for kinks, and non-overlapping sites, is a *convolution* of a uniform, ideal-gas-like distribution and the underlying non-exponential distribution of bare excitations.

For very short observation times and at very low temperatures, the subsampling dominates the resulting distribution of kinks. In figure 3.10, the distributions for $t_{\text{obs}} = 10$ (blue) and $t_{\text{obs}} = 100$ (red) are very exponential beyond ℓ_{ne} , with the post-peak curvature completely disappearing.³⁴ However, as the observation time increases, the distribution of kinks begins to resemble the distribution of excitations. With $t_{\text{obs}} = 1000$, the two are nearly identical (figure 3.10, right). This result suggests that kinks are a good proxy for the underlying excitations, provided the observation time is sufficiently long.

In contrast to the distribution of kinks, which changes noticeably as the observation time becomes long, the distribution of non-overlapping sites is relatively insensitive to t_{obs} . Figure 3.11 shows how $P(\ell)$ for sites satisfying the non-overlap condition, $o_i(t) = 0$, barely changes even as t_{obs} increases six orders of magnitude. Unlike $p_i(t)$, which remains zero once site i has flipped the first time, $o_i(t)$ alternates between 0 and 1 as site i changes its state. Consequently, the overlap condition always undercounts sites that have relaxed and even at long observation times only a small fraction of the underlying excitations is detected. We will return to this distinction between kinks and non-overlapping sites when we consider an atomistic model in the next section of this chapter.

Glassy configurations are no more stable than their equilibrium counterparts

Another difference between equilibrium and non-equilibrium states lies in the decay of the mean persistence function defined by equations 3.7 and 3.8. If we consider two configurations, one in equilibrium and the other produced by cooling, with the same average inter-excitation distance ℓ_{avg} , their persistence functions $\bar{p}_{\text{eq}}(t)$ and $\bar{p}_{\text{ne}}(t)$ decay over similar timescales since the typical domain size is identical. However, some of the details of this decay differ. Because the non-equilibrium configuration lacks small domains, the initial decay of $\bar{p}_{\text{ne}}(t)$ will proceed slowly. In contrast, the equilibrium configuration contains many small domains that can relax almost immediately, and therefore $\bar{p}_{\text{eq}}(t) < \bar{p}_{\text{ne}}(t)$ for early times t .³⁵ Once t approaches the timescale required to relax the most probable length ℓ_{ne} in the non-equilibrium distribution, $\bar{p}_{\text{ne}}$ will decay quickly and overtake $\bar{p}_{\text{eq}}$. At longer times t , the equilibrium persistence function decays a bit more slowly, due to the greater abundance of very large domains in equilibrium.

³⁴For $\ell < \ell_{\text{ne}}$, the correlation hole dominates, since the convolution of a uniform distribution and zero is zero.

³⁵The rarity of domains smaller than ℓ_{ne} explains why Keys, Garrahan, and Chandler find no dip in the heat capacity upon heating. The depletion of probability $P(\ell)$ for $\ell < \ell_{\text{ne}}$ creates a minimum energy barrier that must be overcome before the system may relax appreciably.




Figure 3.12: Mean persistence functions for East model glasses prepared by cooling (solid lines) compared with their corresponding counterparts with the same average inter-excitation length (dashed lines). From left to right (maroon to navy), the cooling rate varies from 10^{-3} to 10^{-9} . Relative to their equilibrium counterparts, the non-equilibrium persistence functions decay more slowly for short times, but more quickly for longer times.

Figure 3.12 illustrates these differences for a wide range of cooling rates (10^{-3} to 10^{-9}) used to prepare the non-equilibrium state.

The results and discussion in this section establish that the emergence of static correlations between excitations is the distinguishing feature of East model glasses. Although it has been long understood that real-world glasses are liquids that have fallen out of equilibrium, the East model provides a microscopic understanding of this process, and the non-exponential distribution that emerges is its microscopic signature.

3.3 The Wahnström model: An atomistic model of glass-forming liquids

The previous sections introduced the glass transition and a simple model — the East model — for understanding it. In this section, we introduce the Wahnström model, an atomistic model of glass-forming liquids [39]. In reference [15], Keys et al. establish a quantitative connection between atomistic models and the one-dimensional East model. In this chapter, we exploit this connection to make predictions about the Wahnström model when it is subjected to the

kinds of cooling protocols considered for the East model. Although atomistic models are much more complex and computationally expensive to study in comparison to kinetically constrained models, they are much more realistic and allow us to bridge the gap between simple but highly abstract models like the East model and real-world glass formers. In the appendix to this chapter, we also briefly consider the two-dimensional atomistic model from reference [3]. However, our main focus (for reasons discussed in the appendix) is the three-dimensional Wahnström model.

Table 3.3 lists relevant parameters for the two models that we consider. Both models are binary mixtures of Lennard-Jones particles, so that the pairwise potential $u_{ij}(r)$ between particles of type i and j is given by

$$u_{ij}(r) = \begin{cases} 4\epsilon_{ij} [(\sigma_{ij}/r_{ij})^{12} - (\sigma_{ij}/r_{ij})^6] + \text{constant}, & r < r_c, \\ 0, & r > r_c, \end{cases}$$

with the constant chosen to ensure continuity at the cutoff. Note that the choice of cutoff for the two-dimensional model (see table 3.3) turns $u_{ij}(r)$ into the purely repulsive potential of Weeks, Chandler, and Andersen (WCA) [40].

Also included in table 3.3 are the East model parameters computed in reference [15]. With these East parameters, we can predict the equilibrium relaxation time at any temperature $T < T_o$.³⁶ With $\tau(T)$ known quantitatively, we can use equation 3.2 to predict the glass transition temperature T_g ³⁷ and then, by evaluating the $\ell_{\text{eq}}(T)$ at T_g , predict the non-equilibrium lengthscale that will emerge. We can also do the reverse, calculating the cooling rate required to achieve a particular non-equilibrium length ℓ_{ne} . For detailed discussions of these types of calculations, consult references [17, 16]. Interestingly, of all the models considered listed in Table II of reference [15], the Wahnström model requires the least cooling time (i.e., the fastest rate) to achieve a particular ℓ_{ne} . It is for this reason that we chose the Wahnström model over other possibilities.

3.3.1 Avoiding crystallization

Typical liquids undergo a phase transition to a crystalline state at modestly cold temperatures. The presence of this transition at temperatures comparable to T_o greatly complicates the study of supercooled liquids. Some effort to avoid or inhibit crystallization is therefore necessary in any study of the glass transition.³⁸

The choice of model is crucial in this respect. Binary mixtures of particles with different sizes tend to resist crystallization relative to their monodisperse counterparts, as the

³⁶In addition to the four parameters listed in table 3.3, we also need the *fractal dimensionality* d_f , which in accordance with reference [15] we take to be $d_f = 2.6$.

³⁷In doing so, we assume that equation 3.2 holds as a strict equality.

³⁸In reference [15], all the models studied were found to crystallize at sufficiently low temperatures after sufficiently long equilibration times. To prevent this from spoiling their results, Keys et al. discarded all trajectories in which the system was found to have crystallized, retaining those that remained in the metastable liquid basin.

Model	ρ	f_A	σ_{AA}	σ_{AB}	σ_{BB}	$r_{ij}^{(c)}/\sigma_{ij}$	m_A	m_B	J_0	γ	T_o	τ_o
Wahnström	1.296	0.5	1	11/12	5/6	2.5	2	1	9.4	0.36	0.88	1.8
WCA2D	0.75	0.32	1	11/10	7/5	$2^{1/6}$	1	1	21.8	0.57	2.13	60

Table 3.2: Parameters of the models considered in this work. The parameters listed between the double column bars (columns 2 through 9) fully define the models. Definitions of these parameters are as follows: ρ is the number density N/V ; f_A is the fraction of particles that are type A ; σ_{ij} is the Lennard-Jones distance parameter for ij pairs; and m_i is the mass of type- i particles. For both models, for all pairs, the Lennard-Jones interaction parameter has the value $\epsilon = 1$. The four parameters past the double column lines — J_0, γ, T_o, τ_o — are the East model parameters calculated in reference [15]. Note that the final one, τ_o , the relaxation time at the onset temperature, is misprinted in reference [15] for the Wahnström model. The number listed there is 100 times the correct number printed here.

presence of multiple microscopic lengthscales frustrates the ability of particles to order themselves over more than a few particle diameters. For this reason, binary mixtures dominate theoretical studies of supercooled liquids and the glass transition. The Kob-Andersen model is an especially popular choice in the literature, owing to its set of non-additive force-field parameters that favor interactions between A and B particles and allow them to come closer together than like pairs.³⁹ These non-additive parameters further stabilize the mixed state, ensuring that separation into two pure phases, which could then crystallize easily, does not occur. Other binary mixture models, like the Weeks-Chandler-Andersen (WCA) purely repulsive model, or the Wahnström model considered in this study, which have normal additive force-field parameters, are also widely and successfully used to study supercooled liquids.

In principle, even models designed specifically to inhibit crystallization, like the Kob-Andersen model, cannot entirely avoid the problem. At some (possibly very low) temperature, the liquid state becomes metastable.⁴⁰ At that point, the timescale for crystallization, τ_{xtl} , places an upper bound on the timescales over which we can study the liquid state. If, for example, we wish to form a glass, we must limit ourselves to cooling times much shorter than τ_{xtl} . Studying a supercooled liquid in equilibrium⁴¹ similarly requires limiting ourselves to temperatures where τ_{rel} , the liquid relaxation time, satisfies $\tau_{\text{rel}} \ll \tau_{\text{xtl}}$. Thus, crystallization places bounds on the conditions under which we may study the supercooled liquid, both in and out of equilibrium.

One consequence of the timescale constraint imposed by crystallization is that the stability of glasses that may be formed by cooling is bounded. If ν_c^{min} is the slowest possible cooling rate that can be employed without the system crystallizing, and T_g^{min} the glass transition

³⁹Specifically, for the Kob-Andersen model, σ_{AB} and ϵ_{AB} are smaller and larger, respectively, than their counterparts for like pairs.

⁴⁰And at an even lower temperature, the liquid state will become unstable and the system will immediately begin to crystallize.

⁴¹Or quasi-equilibrium, since the liquid state may no longer be the global free energy minimum.




Figure 3.13: Potential energy per particle $u \equiv U/N$ for the Wahnström model at $\beta = 1.72$. The system is initially out of equilibrium at $t = 0$, the configuration taken from equilibrium at a higher temperature. The time scale presented here is roughly three orders of magnitude larger than the measured relaxation time of the metastable liquid.

temperature corresponding to this rate, then $\ell_{\text{ne}}^{\text{max}} = \ell_{\text{eq}}(T_g^{\text{min}})$ is the largest non-equilibrium lengthscale that can be embedded into the material. Models that readily crystallize have small values of $\ell_{\text{ne}}^{\text{max}}$ and are considered poor glass formers, while those that resist crystallization — like the Kob-Andersen model, and other models like it with non-standard force-field parameters — have large $\ell_{\text{ne}}^{\text{max}}$ and are considered good glass formers.

3.3.2 Crystallization of the Wahnström model at modestly supercooled conditions

Unfortunately, the Wahnström model appears to be a relatively poor glass former, displaying clear signs of crystallization at temperatures not far below the onset temperature. Figure 3.13 shows the potential energy per particle as a function of equilibration time at $T = 0.588$. The persistent downward drift of the energy is a clear indication that the system has failed to reach global equilibrium, even though the time scale depicted is 3-4 orders of magnitude larger than the measured relaxation time. As we will see shortly, this is a clear sign that the system is undergoing a transition to the ordered state. The temperature is well below T_o but not significantly so, and is well within the range of temperatures studied in reference

[15].⁴²

Since the system in figure 3.13 is initially out of equilibrium, one might be tempted to attribute the drift in the potential energy to an aging process in which the supercooled liquid slowly approaches equilibrium. Although the timescale of the drift is much longer than the measured relaxation time τ , there could be a separation of timescales between τ and t_{age} . That is, the time t_{age} over which the system exhibits significant aging dynamics⁴³ could be significantly longer than the time required for the self-correlation function to decay. Even after τ time has elapsed, F_s will, by definition, not have fully decayed, and a large fraction of particles will have yet to move significantly. For all particles to move, and for F_s to decay more fully, typically requires waiting at least an order of magnitude longer than τ . And even once $F_s \approx 0$, much of the system will retain memory of the initial out-of-equilibrium condition, as many particles will not have moved enough to significantly rearrange their local configurations. It is therefore reasonable to expect that $t_{\text{age}} \gg \tau$.⁴⁴

Nevertheless, on the time scales pictured in figure 3.13, which are much longer than any reasonable estimate of t_{age} for the supercooled liquid, we would expect the system to have reached equilibrium. That the energy continues to drift slowly downward suggests that crystallization dynamics are at work. In reference [28], Pedersen et al. show that the Wahnström model forms a crystal with a 2:1 ratio of small to large particles, i.e., AB_2 , forming the well-known MgZn_2 structure. This structure consists of a tetrahedral network of so-called Frank-Kasper (FK) bonds between neighboring A particles [29]. These bonds are characterized by the formation of a hexagonal ring of B particles about the vector connecting the A particles, as depicted in figure 3.14.

To detect the formation of these FK bonds, we begin by computing the set of nearest neighbors for every A particle in the system via the solid-angle-based procedure described in reference [21]. An A particle forms an FK bond with one of its nearest neighbor A particles if there is an overlapping set of six B particles that are nearest neighbor to both. Figure 3.15 shows a system equilibrating at $\beta = 1.7$, showing only A particles that form exactly four FK bonds. These particles display clear ordering, with the formation of a tetrahedral network of FK bonds particularly clear in rendering (B).

The system depicted in figure 3.15 is clearly beyond the nucleation phase and well into the growth phase, judging by the presence of multiple crystalline clusters of differing orientations. As this system continues to equilibrate, these clusters will continue to grow until they meet each other at grain boundaries, at which point the system will undergo coarsening as the crystallites compete to dominate the global ordering.

Because the systems we consider contain an equal mixture of particles while the crystal

⁴²The liquid also appears to become thermodynamically unstable at around $\beta \approx 1.75$ (data not shown).

⁴³In the present context, we define t_{age} as the time required for the system to approach equilibrium, as judged by some reason thermodynamic quantity such as the potential energy.

⁴⁴Indeed, the remarkable agreement in figure 3.18 between the simulation results and the predictions of equation 3.2 is likely due to this separation of timescales. In section 3.2.2, we argued that our inability to predict the out-of-equilibrium slope for the East model (see figure 3.5) was due to the *lack* of this separation of timescales, as the East model would age significantly in time τ .




Figure 3.14: Illustration of Frank-Kasper (FK) bonds between neighboring A particles (gray). Each FK bond is characterized by a hexagonal ring of B particles (green) that belong to the set of nearest neighbors of both A particles. The six B particles belonging to the FK bond between particles 1 and 2 are shown connected in red. In the crystal phase, each A particle forms four FK bonds with neighboring A particles, as particle 1 in the illustration does. As a result, the crystal phase is characterized by a tetrahedral network of FK bonds, similar to the structure of ice Ih.

structure has a 2:1 ratio of large to small particles, formation of the crystal must be accompanied by phase separation in which “unused” A particles are expelled from the crystallizing mixture. Crystallization should thus be accompanied by large fluctuations in the local composition. Figure 3.16(b) confirms this, showing that the distribution of the number of B particles in $5 \times 5 \times 5$ sub-volumes broadens considerably as the system crystallizes.

In crystallization studies, it is common to employ bond-orientational order parameters to distinguish the liquid and crystal phases, similar to the ones originally discussed in reference [35]. For each particle i , we compute a bond-orientational vector

$$\psi_{\ell m}^i = \frac{1}{|\text{nn}(i)|} \sum_{j \in \text{nn}(i)} Y_{\ell m}(\hat{\mathbf{r}}_{ij}).$$

Here, $\text{nn}(i)$ denotes the set of particles that are nearest neighbors of particle i , as determined by the solid-angle-based method in reference [21], $|\text{nn}(i)|$ is the size of this set, $Y_{\ell m}$ is the ℓm complex spherical harmonic, and $\hat{\mathbf{r}}_{ij}$ is the unit vector pointing from i to j . Since m takes on integer values ranging from $-\ell$ to $+\ell$, $\psi_{\ell m}^i$ is a $(2\ell + 1)$ -component vector. We choose $\ell = 6$ since order parameters based on ψ_{6m}^i are known to be particularly effective at distinguishing amorphous and ordered phases. Figure 3.16(a) shows the distribution of ψ_6^i , the magnitude




Figure 3.15: *A* particles participating in exactly four Frank-Kasper bonds. Both images (A and B) were rendered in VMD, with image (A) using the van der Waals representation and image (B) using the dynamic bonds representation with cutoff 1.6. Regular ordering of the *A* particles is evidence in both images, but the tetrahedral network of FK bonds is particularly evidence in image (B).

of this 13-component vector:

$$(\psi_6^i)^2 = \sum_{m=-\ell}^{\ell} \psi_{6m}^i \psi_{6m}^{i*}.$$

Figure 3.16(a) shows that crystallization of the Wahnström mixture is accompanied by abrupt increases and decreases in ψ_6^i for *B* and *A* particles, respectively.

3.3.3 Rapid cooling prevents crystallization in the Wahnström liquid

In the previous section, we established that the Wahnström fluid begins to crystallize at modestly supercooled temperatures. Specifically, the liquid state becomes metastable near $\beta \approx 1.6$ and thermodynamically unstable at $\beta \approx 1.8$. Unfortunately, this temperature coincides with the conditions where the hierarchical dynamics of the (metastable) liquid become important. In reduced units, this temperature range is roughly $1.6 < \tilde{\beta} < 2.3$, which is where the curvature in the $\ln \tau$ versus T starts to become nontrivial. Nevertheless, we can




Figure 3.16: (a) Distribution of ψ_6^i for several temperatures. The liquid phase corresponds to the center peak, which shifts to modestly higher values as temperature is lowered. The crystal phase corresponds to the the low- ($\psi_6^i \approx 0.2$) and high-value ($\psi_6^i \approx 0.6$) peaks, the former for *A* particles and the latter for *B* particles. (b) Composition fluctuations for *A* (left) and *B* (right) particles as temperature is lowered. The system is divided into approximately $5 \times 5 \times 5 \sigma^3$ subvolumes, then the number of each type of particle in each subvolume is counted. The distribution for *A* particles changes little with temperature, while the distribution for *B* particles broadens considerably as the system undergoes crystallization (black, right).

study what happens to the liquid as it glassifies provided that we cool at a sufficiently rapid pace.⁴⁵

For the present study, we consider two cooling rates, which we denote “fast” and “slow.” (More details about the protocols are provided in the next section.) Figure 3.17 plots the number of Frank-Kasper bonds as a function of the temperature, which starts at T_o and ends at $T = 0$. Although the number of FK bonds increases noticeably at intermediate temperatures — where the liquid becomes metastable — it quickly levels off as the temperature is lowered even further and the material is arrested. Even for the slow protocol, which ends with roughly twice as many FK bonds as the fast protocol, the number of FK bonds remains very small compared to the system size. For this reason, we are confident that our results for the out-of-equilibrium liquid, which will be presented in subsequent sections, are not spoiled by the crystallization process.

⁴⁵Since we are limited to fast cooling rates by finite computational resources, in practice this the constraints imposed by crystallization is not as consequential as it seems.




Figure 3.17: Number of Frank-Kasper bonds N_{FK} versus temperature during the cooling protocols. Note that $N_{\text{FK}}/2 = N_A$ if all A particles participate in four FK bonds — that is, there are two FK bonds per A particle (FK/ A). Even for the slower cooling rate (blue), there are fewer than 0.02 FK/ A , or less than 1% of the value for a fully crystalized system.

3.4 Forming a glass from a Wahnström fluid

In this section, we make predictions for our chosen atomistic model when it is cooled from a very high temperature to $T = 0$. Each prediction is related to one the features of the East model that we explored in the section 3.2. Our goal in studying an atomistic model, which is considerably more complex and computationally expensive to simulate than the East model, is to strengthen the correspondence between the two. In particular, we intend to show that the out-of-equilibrium signatures of East model glasses have analogues in atomistic models of glass formers. In doing so, we hope to bridge gap between the abstract lattice model and real-world glasses.

In what follows, we consider cooling protocols that start from the onset temperature T_o and cool to $T = 0$ in a finite time interval at a constant rate. We identify a cooling protocol by the absolute value⁴⁶ of the rate of change of temperature. Unless otherwise noted, in this

⁴⁶We adopt this convention even though the rate of change of temperature in a cooling protocol is negative because it will always be clear from context or notation whether we are referring to heating (positive rate) or cooling (negative rate).

work the rate may always be assumed to be in units of T_o/τ_{LJ} , where

$$\tau_{LJ} = \left(\frac{m\sigma^2}{\epsilon} \right)^{1/2}$$

is the natural Lennard-Jones time unit. For example, a cooling rate of $\nu_c = 0.1$ corresponds to a protocol in which the temperature decreases linearly from T_o to zero in a time $10\tau_{LJ}$.

In this section, we present results for two cooling protocols, one with rate 10^{-4} (“fast”) and the other with rate 10^{-5} (“slow”). Studying cooling rates that differ by an order of magnitude allows us to explore the effect of the rate on the properties of the glasses that form, and to probe the validity of some of the predictions we will make in the next several sections. As we mentioned in section 3.3.3 and demonstrated in figure 3.17, these protocols are rapid enough to arrest crystallization. The remainder of this chapter will be spent examining a series of East-model-based predictions and their manifestations (or lack thereof) in the Wahnström fluid. Before diving into these predictions and results, we take a brief detour to review the concept of averages, both *ensemble* and *iso-configurational*, since they are so crucial to the calculations we perform later in this chapter.

3.4.1 Calculating non-equilibrium averages

We often refer to averages in and out of equilibrium with some degree of ambiguity. Here, we clarify the meaning of an average and offer some comments about how these averages are computed in practice. To begin, consider a quantity x that is a function of the configuration $\mathcal{C}$ of the system, $x = x(\mathcal{C})$. Of interest is the *ensemble average* of x ,

$$\langle x \rangle \equiv \sum_{\mathcal{C}} w(\mathcal{C}) x(\mathcal{C}), \quad (3.13)$$

where $w(\mathcal{C})$ is a weight function that determines the contribution of configuration $\mathcal{C}$ to the average. Choosing an ensemble means choosing the function w . The micro-canonical ensemble weights all allowed configurations equally, while the canonical ensemble weights configurations by the Boltzmann factor $\exp(-E(\mathcal{C})/T)$, where $E(\mathcal{C})$ is the energy of configuration $\mathcal{C}$ and T is the temperature. Non-equilibrium ensembles, corresponding to systems driven out of equilibrium by a protocol of time-varying external parameters, are also possible, though an analytical expression for w may be known only in special cases.

In practice, the ensemble average of x is computed by a simple average of x over configurations drawn from the corresponding distribution w . If we sample M configurations, then we estimate $\langle x \rangle$ with

$$\bar{x} \equiv \frac{1}{M} \sum_{i=1}^M x(\mathcal{C}_i). \quad (3.14)$$

Provided the set of configurations $\{C_i\}$ is representative of the ensemble,⁴⁷ $\bar{x}$ will be a good estimate of the true ensemble average. For equilibrium ensembles, this representative subset is typically generated by periodically sampling configurations from a sufficiently long time series following a generous period of equilibration.

Non-equilibrium ensembles require specification of an *initial condition*, i.e., the ensemble that prevails at $t = 0$, as well as the time dependence of all external parameters for $t > 0$.⁴⁸ This time dependence means that time averaging is not an option out of equilibrium. Instead, we generate representative subsets by independently preparing a large number of systems according to the protocol, drawing initial configurations from the known ensemble at $t = 0$ and then subjecting those configurations to the protocol of time-varying external parameters. The configurations remaining at the end of the protocol comprise a representative subset of the non-equilibrium ensemble.

Because the cooling protocols we consider begin with the system in equilibrium at a high initial temperature T_i , the starting configurations are drawn from the canonical ensemble at temperature T_i .⁴⁹ All parameters besides the temperature are fixed and T decreases linearly with time, so the rate ν_c is sufficient to fully specify a cooling protocol.

So far we have implicitly limited ourselves to cooling protocols that take the temperature all the way down to zero. However, we could imagine cooling to an intermediate temperature T such that $T_i > T > 0$. The rate ν_c then corresponds to a family of protocols and associated non-equilibrium ensembles, differentiated by the final temperature T to which the system is cooled. Note that the ensemble corresponding to ν_c and T could also be produced by *pausing* the full cooling protocol when it reaches temperature T and sampling the configurations at that point, which is the approach we take in practice. If $x = x(C)$ is our quantity of interest, then we denote the average over this non-equilibrium ensemble by $\langle x(T) \rangle_{\nu_c}$. For a fixed rate ν_c , this quantity is a function of temperature indicating how the average of x changes as the system is cooled to lower temperatures at rate ν_c . As ν_c changes, this function also changes, as in figures 3.6 and 3.18.

For any $\nu_c > 0$, there will be a glass transition temperature, T_g , satisfying equation 3.2, below which this function will deviate substantially from the equilibrium result. In the limit of infinitely slow cooling, or $\nu_c \rightarrow 0$, we expect $\langle x(T) \rangle_{\nu_c}$ to approach the equilibrium ensemble average. In the work that follows, we will typically drop the subscript ν_c for ease of notation,

⁴⁷The term *representative* implies that the set of configurations $\{C_i\}$ contains a sufficiently large subset of statistically independent configurations drawn from the relevant ensemble. Note that this is not necessarily equivalent to making M large, since the set $\{C_i\}$ may contain correlations, as it does when the configurations are sampled from a time series at time intervals on the order of, or less than, the relevant correlation time.

⁴⁸We assume, without loss of generality, that all protocols begin at $t = 0$.

⁴⁹This drawing of initial configurations is accomplished as follows: We place particles on a regular lattice until the desired density is achieved, then randomly select particles to be type B until the total fraction of particles of type B matches the specification of the model. With this initial configuration in hand, we draw initial velocities from the Maxwell-Boltzmann distribution at temperature T_i and equilibrate for at least 10^3 times the structural relaxation time at T_i . Configurations produced in this way are statistically independent and representative of the canonical ensemble at temperature T_i .

but it should be clear from context whether the average is equilibrium or non-equilibrium, and for the latter, what the rate should be.

3.4.2 Iso-configurational averaging

In the previous section, we restricted our attention to static observables that depend only on a single configuration. However, in much of the work that follows, we will be concerned with *dynamic* observables, which depend on the configuration at multiple points in a trajectory, such as

$$x = x(\mathcal{C}(t_1), \mathcal{C}(t_2), \dots, \mathcal{C}(t_N)),$$

where $t_1 < t_2 < \dots < t_N$. Typically we will focus on observables that depend only on the endpoints of the trajectory, so that the above simplifies to

$$x = x(\mathcal{C}(0), \mathcal{C}(t)),$$

where t is the total length of the trajectory, which we often call the *observation time*.

The so-called *iso-configurational ensemble* (see reference [41]) associated with an initial configuration $\mathcal{C}(0)$, a temperature T , and an observation time t is generated by drawing initial velocities from the Maxwell-Boltzmann distribution at temperature T and then integrating Newton's equations of motion forward to time t to produce a new configuration $\mathcal{C}(t)$. The average of a dynamical observable x over this new ensemble is termed the *iso-configurational average*, and is itself a *static* observable that depends only on the initial configuration $\mathcal{C}(0)$. For convenience, we denote this observable by $x_{\text{iso}}(\mathcal{C})$.

We will often write $\langle x(\mathcal{C}(0), \mathcal{C}(t)) \rangle$ when we mean to refer to the ensemble average of x_{iso} , so that the following expressions are equivalent:

$$\langle x(\mathcal{C}(0), \mathcal{C}(t)) \rangle \equiv \langle x_{\text{iso}} \rangle.$$

That is, when dealing with a dynamical observable, the angle brackets denote an ensemble average as well as an iso-configurational average.⁵⁰ In practice, this average is computed via an extension of equation 3.14:

$$\bar{x} = \frac{1}{M} \sum_{i=1}^M \frac{1}{N_{\text{iso}}} \sum_{j=1}^{N_{\text{iso}}} x(\mathcal{C}_i, \mathcal{C}_{ij}(t)).$$

In this expression, N_{iso} is the number of trajectories run per initial configuration $\mathcal{C}_i$ to sample its iso-configurational ensemble, and $\mathcal{C}_{ij}(t)$ is the configuration produced at time t in the j th trajectory that starts with configuration $\mathcal{C}_i$.

For $\bar{x}$ to produce a good estimate of $\langle x(t) \rangle$, it is important that M is large. By contrast, it is *not* particularly important that N_{iso} be large, provided M is large. The iso-configurational average is essentially just an average over the initial velocities of the trajectories, and for a

⁵⁰More accurately, it denotes averages over two ensembles, one of which is the iso-configurational ensemble.




Figure 3.18: Logarithm of the relaxation time versus inverse temperature for the Wahnström model. The dashed curve shows the parabolic law. The solid red and blue curves show non-equilibrium relaxation times for cooling protocols with $\nu_c = 10^{-4}$ (“fast”) and $\nu_c = 10^{-5}$ (“slow”), respectively. The relaxation time τ satisfies the relation $F_s(q_0, \tau) = 1/e$, where $q_0 \approx 7.7$ coincides with the first peak in the static structure factor and $F_s(k, t)$ is the self-correlation function. The thin dashed lines in red and blue show the prediction of equation 3.2 for the fast and slow cooling rates, respectively, assuming unit proportionality constant and an abrupt transition at T_g (see discussion in main text).

sufficiently large value of M , this velocity averaging is satisfied even for $N_{\text{iso}} = 1$. In practice, however, the value of M will be limited (by computational resources), and a larger value of N_{iso} will sometimes be necessary.

3.4.3 Prediction 1: Crossover from parabolic to Arrhenius

A defining feature of the glass transition, both in the laboratory and in the East model, is the crossover from super-Arrhenius to Arrhenius behavior in the relaxation time for temperatures below T_g (see figure 3.1). As explained in section 3.2, this crossover is a consequence of the structural arrest of the liquid, as the free energetic barrier to relaxation plateaus with inverse temperature.

To test this prediction, we compute the relaxation time $\tau(T)$ for the Wahnström model. We define the relaxation time as the time τ that satisfies $F_s(q_0, \tau) = 1/e$, where F_s is the self-correlation function and $q_0 = 2\pi/\sigma_0$, where $\sigma_0 \approx \sigma$ is a principal structural length in the

system. The exact choice of σ_0 is unimportant, provided that it is roughly of the order of the inter-particle spacing in the system. In practice, we choose q_0 to coincide with the first peak in the static structure factor, which for the model under consideration is $q_0 \approx 7.7$. The self-correlation function, F_s , is defined as

$$F_s(k, t) \equiv \langle \exp [i\mathbf{k} \cdot (\mathbf{r}_j(t) - \mathbf{r}_j(0))] \rangle. \quad (3.15)$$

The angle brackets denote an average over the appropriate ensemble, which in this case is the non-equilibrium ensemble associated with the cooling protocol. We thus compute two relaxation time curves, one for each cooling rate (10^{-4} and 10^{-5}) under consideration.

Figure 3.18 shows the results of our computations. We show $\ln \tau(\beta; \nu_c)$ for the two values of ν_c (fast and slow), juxtaposed against the prediction of the parabolic law based on the parameters in table 3.3. The results in figure 3.18 confirm our expectations. There is a clear crossover to Arrhenius behavior and $\ln(\tau)$ becomes linear in β at temperatures well below T_g . Furthermore, slower cooling pushes this crossover to lower temperatures.

In figure 3.18 we also compare the computed data (red/blue points) to out-of-equilibrium predictions (dashed red/blue lines). These predictions are based on equation 3.2, reproduced here,

$$\frac{d\tau}{dT} = \frac{1}{\nu_c}, \quad (3.16)$$

assuming that the relation holds as a strict equality with unit proportionality constant. For a given cooling rate ν_c , equation 3.16 defines a glass transition temperature T_g where the supercooled liquid falls out of equilibrium. As we discussed in the section introducing the East model, this process is not a sharp transition, but rather occurs gradually around T_g . However, as a rough approximation we can assume a sharp transition, as we did in producing figure 3.1. In this case, the relaxation time will obey the relation

$$\ln(\tau/\tau_0) = E(T_g) \left(\frac{1}{T} - \frac{1}{T_o} \right), \quad (3.17)$$

for $T < T_g$, where $E(T)$ is the equilibrium energy barrier to relaxation at temperature T . At this point, it is convenient to work in units of J , where J^2 is the curvature of the equilibrium relaxation time curve and related to the East model parameters of table 3.3 by

$$J^2 = \gamma J_\sigma^2 / d_f.$$

Letting $\tilde{\beta} \equiv J(1/T - 1/T_o)$ and $\tilde{E} \equiv E/J$, equation 3.17 becomes

$$\ln(\tau/\tau_0) = \tilde{E}(\tilde{\beta}_g) \tilde{\beta},$$

where $\tilde{\beta}_g = J(1/T_g - 1/T_o)$ and $\tilde{E}(\tilde{\beta}) = \tilde{\beta}$. Thus, once we have computed T_g from equation 3.17 for a particular cooling rate, it is trivial to predict the out-of-equilibrium behavior of $\ln(\tau/\tau_0)$ under the assumptions we outlined above.⁵¹ These predictions are shown in figure

⁵¹In reference [16], Limmer goes through a similar calculation for supercooled water.

3.18 for the two cooling rates considered. Considering the roughness of the approximation we have employed (i.e., assuming unit proportionality constant in equation 3.17 and an abrupt transition at T_g), the agreement with the data is remarkably good.

3.4.4 Prediction 2: Concentration of soft spots plateaus

Recall that the structural arrest of an East model glass is due to the inability of excitations to de-excite their nearest neighbors over the finite timescale imposed by the non-equilibrium protocol. As a result, the concentration of excitations cannot adjust to the changing temperature below T_g and $c(T)$ plateaus to a constant value that decreases with slower cooling. This result was shown for the East model in figure 3.6, for cooling rates spanning six orders of magnitude.

Establishing a similar result for the atomistic model requires identifying the localized soft spots that are analogous to the discrete excitations of the East model. Unfortunately, a structural order parameter that can reliably identify these soft spots does not yet exist, so that $c(T)$, and the excitations themselves, cannot be observed directly. Instead, we must rely on the dynamical manifestations of soft spots, in the form of enduring particle displacements. In an atomistic model, we may think of an excitation as a local configuration of unjammed particles, where small displacements in particle positions are possible, and which are higher in energy than jammed configurations where such motion is difficult or impossible. Because particles in these soft spots are much more likely to move, changes in the inherent structure are reasonable, if approximate, indicators of underlying excitations.

For atomistic models, we therefore study the *rate of particle displacements*, denoted by $r(T; a, t_{\text{obs}})$, which we expect to be related in a straightforward way to the (unobservable) concentration of excitations $c(T)$. We define it as

$$r(T; a, t_{\text{obs}}) \equiv \langle \Theta (|\bar{\mathbf{r}}_i(t_{\text{obs}}) - \bar{\mathbf{r}}_i(0)| - a) \rangle, \quad (3.18)$$

where the angle brackets denotes both iso-configurational and ensemble averaging.⁵² The observation time t_{obs} will typically be set to a value not much larger than the mean instanton time Δt_a , the typical time required for a particle to complete a displacement of length a .

Figure 3.19 shows the rate of displacements $r(T; a)$ for the Wahnström model for the choice of displacement length $a = 0.6\sigma$. Interestingly, the high-temperature deviation from the equilibrium prediction is qualitatively similar to that for the East model kink rate. As expected, the slower cooling protocol follows the equilibrium line to lower temperatures than the faster protocol. Furthermore, the rate does not plateau, just as the East model kink rate does not plateau. However, in contrast to the latter, the rate of displacements in figure 3.19 does not clearly approach a straight line even for temperatures well below T_g . This also contrasts with the relaxation times shown in figure 3.18, which cover the same range of temperatures but show a clear crossover to Arrhenius behavior below T_g .

⁵²We will be interested in how r behaves in and out of equilibrium, so that equilibrium and non-equilibrium ensembles will both be considered.




Figure 3.19: Concentration of 0.6σ displacements ($t_{\text{obs}} = 10$) versus inverse temperature for $T < T_o$. The dashed black line shows the expected equilibrium scaling, $c \propto \exp[-J_a(\beta - \beta_0)]$ with $a = 0.6\sigma$. The slope of this line is $J_a = 7.67$, which we compute from the relation $J_a = J_\sigma(1 + \gamma \ln(a/\sigma))$, using the parameters listed in table 3.3. For ease of comparison, we shift this line vertically so that it coincides with the slow cooling data (blue points/lines) at $T = 0.682$. The thin red and blue lines are the East model predictions for $r(T)$ for the fast and slow cooling rates, respectively. The predicted slope, $(1/2)J_a$, overestimates the actual slope implied by the data. See figure 3.7 and the associated discussion.

3.4.5 Prediction 3: Distribution of soft spots becomes correlated

In section 3.2, we showed that the distribution of excitations in the East model is ideal-gas-like, owing to the lack of energetic coupling between excitations in the Hamiltonian (see equations 3.3 and 3.11). We also saw that cooling an East lattice out of equilibrium introduces nontrivial correlations between excitations, a physical consequence of the East model’s kinetic constraint and the hierarchical dynamics that emerges from it. Figure 3.9 demonstrated how, upon cooling an East lattice down to zero temperature, a so-called “correlation hole” — a lack of excitations near other excitations — develops, manifested in $g(r) < 1$ for $r \lesssim \ell_{\text{ne}}$ and $P(\ell)$ peaking near ℓ_{ne} , where ℓ_{ne} is a non-equilibrium, protocol dependent lengthscale.

The emergence of this non-equilibrium correlation length is the key structural signature of East model glasses, and our expectation is that it also emerges in more realistic atomistic models undergoing a glass transition. In this section, we search for these non-equilibrium




Figure 3.20: Illustration of the directional condition. For each particle that displaces at least the minimum distance a , we compute the nearest neighbor distance and radial distribution function for other displacing particles that lie in its wake. The size of the wake is determined by the angular cutoff θ , which, to be useful, should be no larger than 45° .

correlations in the Wahnström model after it has been cooled well below T_g . As in the last section, we must resort to enduring particle displacements, the dynamical manifestations of excitations, rather than the excitations themselves to probe this question. To start, we compute nearest neighbor distribution functions $P(\ell)$ for the East model, as in figure 3.9, but instead of using excitations (sites with $n_i = 0$) we instead use kinks (sites with $p_i(t) = 0$).

To perform the corresponding atomistic calculations, we limit our attention to particles that, in the observation time t_{obs} , displace by at least a distance a .⁵³ That is, we focus on the subset of atoms for which

$$\Theta(|\bar{\mathbf{r}}_i(t_{\text{obs}}) - \bar{\mathbf{r}}_i(0)| - a) = 1, \quad (3.19)$$

where Θ is the Heaviside function, and the overbars over the position vectors denote inherent structure coordinates. For each particle i satisfying this condition, we compute the distance to all other particles j satisfying this condition, *plus* the additional condition

$$\Theta(\cos^{-1} [(-\hat{\mathbf{u}}_i) \cdot d\mathbf{r}_{ij}] - \theta) = 1,$$

⁵³In reference [15], the authors consider displacements ranging between 0.2 and 2.

where $-\hat{\mathbf{u}}_i$ is the direction of facilitation of particle i and points in the opposite direction as its displacement vector,

$$-\mathbf{u}_i \equiv \bar{\mathbf{r}}_i(t_{\text{obs}}) - \bar{\mathbf{r}}_i(0),$$

$\mathbf{dr}_{ij}$ is the vector pointing from particle i to particle j , and θ is an angular cutoff. The effect of this additional condition is to limit our attention to particles that lie in a cone of angle θ centered about particle i 's facilitation vector, which we take to be the direction opposite its displacement.⁵⁴ This directional condition is analogous to the directional constraint in the East model, which is a crucial component of the model.⁵⁵ Figure 3.20 illustrates the angular condition.

The preceding procedure generates, for each displacing particle, a set of distances to other displacing particles lying in its wake. To compute the atomistic version of $P(\ell)$, we take, for each particle, the minimum of these distances, then *cube* and bin the resulting value to yield the distribution $P(\ell^3)$, which is the $d = 3$ analogue of $P(\ell)$.⁵⁶ Figure 3.21 shows the results of this computation for $a = 0.6\sigma$ over a wide range of temperatures extending to well below T_g , for both protocols. In all cases studied, the distribution of $P(\ell^3)$ is remarkably exponential beyond approximately $\ell = 5\sigma$, indicating that the distribution of soft spots is ideal-gas-like and exhibits none of the inter-excitation correlations that our experience with the East model would suggest.⁵⁷

With the set of distances between displacing particles we can also compute pair correlation function, which we denote $\mu(r; a, \theta)$, to directly probe the question of whether any excitation-excitation correlations emerge out of equilibrium. Figure 3.22 shows the results of this calculation, which are consistent with our results for $P(\ell^3)$. As in figure 3.21, correlations disappear after a few particle diameters.

Figures 3.21 and 3.22 are generated for observation times that are relatively short, on the order of 1–10 mean instanton times. This time is long enough to allow displacing particles to complete their sojourns. However, as we saw earlier in the context of the East model, the choice of observation time, as well as the choice between overlap and kinks, can make a large difference in the observed distribution of distances. Figure 3.23 shows $P(\ell^3)$ vs ℓ^3 between particles that have kinked, i.e., that have displaced at any point during the time interval of observation, for much longer observation times than the ones considered in figures 3.21 and

⁵⁴That is, displacing particles facilitate in the wake of their motion. This choice is physically motivated, since a particle leaves a small void in its wake. Keys et al. also demonstrate quantitatively that motion begets motion at a later time in the opposite direction [15].

⁵⁵Without the directional constraint in the East model, excitations could diffuse. Consequently, the dynamics would not be hierarchical and the energy barrier would not grow logarithmically with length.

⁵⁶In section 3.2, we derived analytically the form of $P(\ell)$ for a one-dimensional ideal gas, showing that $P(\ell)$ is exponential in ℓ . It can be shown that the generalization of this to arbitrary dimensions is $P(\ell^d) \propto \exp(-\ell^d)$.

⁵⁷The large correlation at small values of ℓ corresponds to intra-excitation correlations between particles that displace together as part of a larger collective motion.




Figure 3.21: Distribution of ℓ^3 for the Wahnström model cooled at two different rates. Results for the fast protocol are on the left, for the slow protocol on the right. Part (a) shows the raw distributions for several temperatures. Part (b) removes the contributions from $\ell < 5$ and rescales by the mean value of ℓ^3 , so that exponential (ideal-gas-like) distributions collapse onto the line $y = -x$. For both protocols, at all conditions studied (including many temperatures well below T_g), this distribution is highly exponential beyond $\ell \approx 5\sigma$.




Figure 3.22: Radial distribution function $\mu(r; a, \theta)$ for particles that displace at least a distance a in an observation time $t_{\text{obs}} = 10$. Part (a) compares, for two different choices of the angular cutoff, the results for each cooling protocol at $T = 0.374$, which is well below T_g in both cases. The results from the slower protocol (blue) are offset vertically to make the comparison easier. Estimated 1σ error bars are shown. Part (b) compares the radial distribution function across a wide range of temperatures for $\theta = 45^\circ$.




Figure 3.23: Distribution of ℓ^3 between particles that have “kinked” during an observation time t_{obs} . Kinks count all particles that displace at least a at any point in the time interval $(0, t_{\text{obs}})$, even if they return to the initial position before the end of the time interval. The results shown are for $T = 0.374$ for the slower cooling protocol ($\nu_c = 10^{-5}$), for two different choices of the observation time. The inset shows the same distributions with the intra-excitation correlations removed, and the rest of the distribution scaled by $\langle \ell^3 \rangle$.

3.22. The particles included in the computation of 3.23 all satisfy the condition

$$\prod_{t'=0}^t [1 - \Theta(|\bar{\mathbf{r}}_i(t') - \bar{\mathbf{r}}_i(0)| - a)] = 0,$$

where $\prod$ indicates a product over each timepoint in the interval $(0, t)$,⁵⁸ which ensures that a displacement is counted even if a particle later returns to its original basin. This is a much more liberal condition than equation 3.19.

Nevertheless, choosing a much longer observation time and computing the distribution for kinks does not change the qualitative results. The nearest-neighbor distribution functions in figure 3.23 remain remarkably exponential beyond $\approx 5\sigma$. Indeed, the effect of a much longer observation time appears to be far less dramatic for particle kinks than for East model kinks

⁵⁸Ideally this would be a product over each infinitesimal timestep. In practice, we divide the time interval $(0, t)$ into a finite number of smaller time intervals $(0, \delta t)$, $(\delta t, 2\delta t)$, and so on, where $\delta t = 1\tau_{LJ}$.

(see figure 3.10). That the choice of kinks versus overlap is less impactful in the atomistic model likely has to do with the probability of a model’s constituent parts — particles in the atomistic model, spins in the East model — returning to its original state within a reasonable amount of time. Because the East model is a binary model — each spin n_i assumes only two possible values — every spin that flips will *necessarily* return to its original state upon its subsequent flip. In contrast, a particle that displaces need not, and typically will not, return to its original basin after its next displacement. Consequently, the difference between “non-overlapping at time t ” and “kinked by time t ” is less important in the atomistic model than the East model.

3.4.6 Conclusions and discussion

The results presented for the Wahnström model in this section stand in stark contrast to our prior expectations. The emergence of inter-excitation correlations out of equilibrium is a defining signature of the East model, a natural and intuitive consequence of the microscopic mechanism by which it falls out of equilibrium. The analogous picture for an atomistic system, in which discrete excitations are too far apart to interact with and de-excite each other, offers a compelling explanation for the vitrification process observed in the real world. That these excitations appear to be distributed like an ideal gas, with many excitations close together, raises numerous questions.

Superficially, it may appear that these results rule out the dynamical facilitation picture of supercooled liquid dynamics and the glass transition. The East model makes a clear prediction, and that prediction appears to fail. However, there are several other possible, and more compelling, reasons why our results contradict our prior expectations. First, the one-dimensional East model, for all the rich complexity of its dynamics, is a highly abstract model, and mapping between the $d = 1$ East model and our much more complicated atomistic model is far from a trivial task. It may be the case that our conceptual mapping between the two is incorrect, such that we are looking for correlations “in the wrong place.” In particular, our concept of how excitations excite and de-excite may be too simplistic. Another possibility is that displacements are a poor proxy for the underlying excitations. In the triangular plaquette model, for example, every spin can always flip (there is no kinetic constraint), but underlying excitations cannot be eliminated except through complicated series of spin flips [8]. At very low temperatures, the underlying excitations will be unable to relax, but trivial spin flips can readily occur. The displacements we observe in our atomistic system may be of this variety and thus have little relation to the actual, underlying distribution of excitations. If that is the case, then confirming the emergence of inter-excitation correlations in glasses will likely require the ability to detect excitations directly, which is a daunting challenge that has yet to be met.




Figure 3.24: Average number of particles that have displaced by at least σ in an observation time $t = 75\tau_{LJ} = 5 \langle \Delta t \rangle_\sigma$, where $\langle \Delta t \rangle_\sigma$, versus reduced temperature $J_\sigma(\beta - \beta_o)$, for the WCA2D-68:32 model cooled at rate $\nu_c = 5.13 \times 10^{-7}$. The gray dashed line shows the equilibrium scaling, shifted in the y -direction so that it coincides with the high-temperature point shown.

3.A Forming a glass from a two-dimensional fluid

For many reasons, two-dimensional models would seem to be ideal candidates for studying the glass transition in realistic systems. Their computational demands are far fewer than for three-dimensional models: the number of degrees of freedom per particle is lower, each particle has fewer neighbors, so that there are fewer force calculations, and for a given box size, there are fewer total numbers of particles. A two-dimensional system is also ideal for visualization. However, two-dimensional systems are problematic in ways that do not hold for their three-dimensional counterparts.

In reference [7], Flenner and Szamel claim that glassy dynamics is very different in two dimensions than it is in three dimensions. In particular, the dynamics, as observed in the decay of the self-correlation function, displays significant system size effects that imply very large correlation lengths. Furthermore, the “caging” effect that is so emblematic of the onset of glassy dynamics disappears in very large system sizes. While the case presented by Flenner and Szamel could have benefited from a rigorous corresponding states analysis, the results they do present suggest that glassy dynamics in two dimensions may well be different in key ways from glassy dynamics in three dimensions.




Figure 3.25: Nearest-neighbor distribution function $P(\ell^2)$ for atoms that have displaced at least σ in a very long observation time $t = 15000\tau_{LJ} = 1000 \langle \Delta t \rangle_\sigma$. On the right, we remove the intra-excitation correlations for $\ell < 10\sigma$, then scale and shift the resulting distributions. This collapses the curves onto the exponential, ideal-gas distribution (black dashed line).

These concerns notwithstanding, we attempted to form a glass with the two-dimensional model of reference [3] (also studied in reference [15]), a two-dimensional binary mixture model with purely repulsive interactions, consisting of 32% small (type-*A*) and 68% large (type-*B*) particles. See table 3.3 for the full set of parameters. Figure 3.24 shows the number of one-sigma displacements in an observation time t ,

$$\langle \Theta(|\bar{\mathbf{r}}_i(t) - \bar{\mathbf{r}}_i(0)| - 1) \rangle,$$

versus inverse temperature. At high temperatures, this quantity follows the expected equilibrium scaling, but, as expected, it deviates substantially at lower temperatures, indicating that the system has fallen out of equilibrium. At these lower temperatures, we compute the nearest-neighbor distribution function for 0.6σ displacements, shown in figure 3.25. Like the Wahnström model, the two-dimensional model here exhibits no correlations between displacing particles beyond short range.

Bibliography

- [1] A. Amadei, A. B. M. Linssen, and H. J. C. Berendsen. Essential dynamics of proteins. *Proteins: Structure, Function, and Bioinformatics*, 17 (4): 412–425, 1993.
- [2] G. Bussi, D. Donadio, and M. Parrinello. Canonical sampling through velocity rescaling. *Journal of Chemical Physics*, 126 (1): 014101, 2007.
- [3] R. Candelier, A. Widmer-Cooper, J. K. Kummerfeld, O. Dauchot, G. Biroli, P. Harrowell, and D. R. Reichman. Spatiotemporal hierarchy of relaxation events, dynamical heterogeneities, and structural reorganization in a supercooled liquid. *Physical Review Letters*, 105 (13): 135702, 2010.
- [4] D. Chandler. Hydrophobicity: Two faces of water. *Nature*, 417 (6888): 491, 2002.
- [5] D. Chandler. Interfaces and the driving force of hydrophobic assembly. *Nature*, 437 (7059): 640–647, 2005.
- [6] J. A. Dunkle, L. Wang, M. B. Feldman, A. Pulk, V. B. Chen, G. J. Kapral, J. Noeske, J. S. Richardson, S. C. Blanchard, and J. H. D. Cate. Structures of the bacterial ribosome in classical and hybrid states of tRNA binding. *Science*, 332 (6032): 981–984, 2011.
- [7] E. Flenner and G. Szamel. Fundamental differences between glassy dynamics in two and three dimensions. *Nature Communications*, 6: 7392, 2015.
- [8] J. P. Garrahan. Glassiness through the emergence of effective dynamical constraints in interacting systems. *Journal of Physics: Condensed Matter*, 14 (7): 1571–1579, 2002.
- [9] D. M. Huang and D. Chandler. Temperature and length scale dependence of hydrophobic effects and their possible implications for protein folding. *Proceedings of the National Academy of Sciences*, 97 (15): 8324–8327, 2000.
- [10] D. M. Huang and D. Chandler. The hydrophobic effect and the influence of solute-solvent attractions. *Journal of Physical Chemistry B*, 106 (8): 2047–2053, 2002.
- [11] G. Hummer, J. C. Rasaiah, and J. P. Noworyta. Water conduction through the hydrophobic channel of a carbon nanotube. *Nature*, 414 (6860): 188–190, 2001.

- [12] J. Jäckle and S. Z. Eisinger. A hierarchically constrained kinetic Ising model. *Physik B – Condensed Matter*, 84 (1): 115–124, 1991.
- [13] S. Katira, K. K. Mandadapu, S. Vaikuntanathan, B. Smit, and D. Chandler. Pre-transition effects mediate forces of assembly between transmembrane proteins. *eLife*, 5: e13150, 2016.
- [14] W. Kauzmann. Some factors in the interpretation of protein denaturation. *Advances in Protein Chemistry*, 14: 1–63, 1959.
- [15] A. S. Keys, L. O. Hedges, J. P. Garrahan, S. C. Glotzer, and D. Chandler. Excitations are localized and relaxation is hierarchical in glass-forming liquids. *Physical Review X*, 1 (2): 021013, 2011.
- [16] D. T. Limmer. The length and time scales of water’s glass transitions. *Journal of Chemical Physics*, 140 (21): 214509, 2014.
- [17] D. T. Limmer and D. Chandler. Theory of amorphous ices. *Proceedings of the National Academy of Sciences*, 111 (26): 9413–9418, 2014.
- [18] D. T. Limmer, A. P. Willard, P. Madden, and D. Chandler. Hydration of metal surfaces can be dynamically heterogeneous and hydrophobic. *Proceedings of the National Academy of Sciences*, 110 (11): 4200–4205, 2013.
- [19] K. Lum, D. Chandler, and J. D. Weeks. Hydrophobicity at small and large length scales. *Journal of Physical Chemistry B*, 103 (22): 4570–4577, 1999.
- [20] Lutz Maibaum and David Chandler. A coarse-grained model of water confined in a hydrophobic tube. *Journal of Physical Chemistry B*, 107 (5): 1189–1193, 2003.
- [21] J. A. van Meel, L. Filion, C. Valeriani, and D. Frenkel. A parameter-free, solid-angle based, nearest-neighbor algorithm. *Journal of Chemical Physics*, 136 (23): 234107, 2012.
- [22] T. F. Miller, E. Vanden-Eijnden, and D. Chandler. Solvent coarse-graining and the string method applied to the hydrophobic collapse of a hydrated chain. *Proceedings of the National Academy of Sciences*, 104 (37): 14559–14564, 2007.
- [23] J. Noeske, M. R. Wasserman, D. S. Terry, R. B. Altman, S. C. Blanchard, and J. H. D. Cate. High-resolution structure of the Escherichia coli ribosome. *Nature Structural and Molecular Biology*, 22 (4): 336–341, 2015.
- [24] A. J. Patel, P. Varilly, and D. Chandler. Fluctuations of water near extended hydrophobic and hydrophilic surfaces. *Journal of Physical Chemistry B*, 114 (4): 1632–1637, 2010.

- [25] A. J. Patel, P. Varilly, D. Chandler, and S. Garde. Quantifying density fluctuations in volumes of all shapes and sizes using indirect umbrella sampling. *Journal of Statistical Physics*, 145 (2): 265–275, 2011.
- [26] A. J. Patel, P. Varilly, S. N. Jamadagni, H. Acharya, S. Garde, and D. Chandler. Extended surfaces modulate hydrophobic interactions of neighboring solutes. *Proceedings of the National Academy of Sciences*, 108 (43): 17678–17683, 2011.
- [27] A. J. Patel, P. Varilly, S. N. Jamadagni, M. F. Hagan, D. Chandler, and S. Garde. Sitting at the edge: How biomolecules use hydrophobicity to tune their interactions and function. *Journal of Physical Chemistry B*, 116 (8): 2498–2503, 2012.
- [28] U. R. Pedersen, N. P. Bailey, J. C. Dyre, and T. B. Schrøder. Crystallization of the Wahnström binary Lennard-Jones liquid. eprint: [arXiv:0706.0813](https://arxiv.org/abs/0706.0813).
- [29] U. R. Pedersen, T. B. Schrøder, J. C. Dyre, and P. Harrowell. Geometry of slow structural fluctuations in a supercooled binary alloy. *Physical Review Letters*, 104 (10): 105701, 2010.
- [30] S. Pronk, S. Páll, R. Schulz, P. Larsson, P. Bjelkmar, R. Apostolov, M. R. Shirts, J. C. Smith, P. M. Kasson, D. van der Spoel, B. Hess, and E. Lindahl. GROMACS 4.5: A high-throughput and highly parallel open source molecular simulation toolkit. *Bioinformatics*, 29 (7): 845–854, 2013.
- [31] B. Rotenberg, A. J. Patel, and D. Chandler. Molecular explanation for why talc surfaces can be both hydrophilic and hydrophobic. *Journal of the American Chemical Society*, 133 (50): 20521–20527, 2011.
- [32] B. S. Schuwirth, M. A. Borovinskaya, C. W. Hau, W. Zhang, A. Vila-Sanjurjo, J. M. Holton, and J. H. D. Cate. Structures of the bacterial ribosome at 3.5 Å resolution. *Science*, 310 (5749): 827–834, 2005.
- [33] M. R. Shirts and J. D. Chodera. Statistically optimal analysis of samples from multiple equilibrium states. *Journal of Chemical Physics*, 129 (12): 124105, 2008.
- [34] P. Sollich and M. R. Evans. Glassy dynamics in the asymmetrically constrained kinetic Ising chain. *Physical Review E*, 68 (3): 031504, 2003.
- [35] P. J. Steinhardt, D. R. Nelson, and M. Ronchetti. Bond-orientational order in liquids and glasses. *Physical Review B*, 28 (2): 784–805, 1983.
- [36] S. Vaikuntanathan and P. L. Geissler. Putting water on a lattice: The importance of long wavelength density fluctuations in theories of hydrophobic and interfacial phenomena. *Physical Review Letters*, 112 (2): 020603, 2014.

- [37] S. Vaikuntanathan, G. Rotskoff, A. Hudson, and P. L. Geissler. Necessity of capillary modes in a minimal model of nanoscale hydrophobic solvation. *Proceedings of the National Academy of Sciences*, 113 (16): E2224–E2230, 2016.
- [38] P. Varilly, A. J. Patel, and D. Chandler. An improved coarse-grained model of solvation and the hydrophobic effect. *Journal of Chemical Physics*, 134 (7): 074109, 2011.
- [39] G. Wahnström. Molecular-dynamics study of a supercooled two-component Lennard-Jones system. *Physical Review A*, 44 (6): 3752, 1991.
- [40] J. D. Weeks, D. Chandler, and H. C. Andersen. Role of repulsive forces in determining the equilibrium structure of simple liquids. *Journal of Chemical Physics*, 54 (12): 5237–5247, 1971.
- [41] A. Widmer-Cooper, P. Harrowell, and H. Fynewever. How reproducible are dynamic heterogeneities in a supercooled liquid? *Physical Review Letters*, 93 (13): 135701, 2004.
- [42] A. P. Willard and D. Chandler. The molecular structure of the interface between water and a hydrophobic substrate is liquid-vapor like. *Journal of Chemical Physics*, 141 (18): 18C519, 2014.
- [43] A. P. Willard and D. Chandler. The role of solvent fluctuations in hydrophobic assembly. *Journal of Physical Chemistry B*, 112 (19): 6187–6192, 2008.
- [44] P. R. ten Wolde and D. Chandler. Drying-induced hydrophobic polymer collapse. *Proceedings of the National Academy of Sciences*, 99 (10): 6539–6543, 2002.
- [45] P. R. ten Wolde, S. X. Sun, and D. Chandler. Model of a fluid at small and large length scales and the hydrophobic effect. *Physical Review E*, 65 (1): 011201, 2001.