

Emergent Phenomena in Minimal Models of Soft Matter

by

Zachary Kost-Smith

B.A., Colorado School of Mines, 2004

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Physics

2014

This thesis entitled:
Emergent Phenomena in Minimal Models of Soft Matter
written by Zachary Kost-Smith
has been approved for the Department of Physics

Matthew A. Glaser

Dr. Paul Beale

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Kost-Smith, Zachary (Ph.D., Physics)

Emergent Phenomena in Minimal Models of Soft Matter

Thesis directed by Dr. Matthew A. Glaser

Three models are examined to demonstrate novel and complex behavior stemming from simple specifications. First I explore the hard-core soft-shoulder (HCSS) system where spheres interact via isotropic, purely repulsive potentials involving a hard core repulsion of range σ and a finite step “shoulder” repulsion of range σ_s . This system’s diverse phase morphology is illustrated by the zero-temperature phase diagram calculated via a large scale simulated annealing procedure. I explore the prospect of directed design of self-assembled structures. The system’s incredibly complex behavior and numerous metastable states make it ideal as a test bed for examining the computational effectiveness of various advanced Monte Carlo techniques. I apply several existing methods as well as extend the virtual move Monte Carlo (VMMC) algorithm to models with purely repulsive interactions.

Second, I examine self-assembled bundles of achiral hard rods with distributed, short-range attraction. I show that in the majority of cases the equilibrium state of the bundle is chiral. I use umbrella sampling Monte Carlo and cell theory to compute the free energy as a function of a twist order parameter, and show that the formation of spontaneously chiral bundles is driven by maximization of orientational entropy through a process called *orientational escape*. I map out the phase diagram of bundles in terms of the density and bundle aspect ratio (L/D_b) finding transitions between untwisted, weakly twisted, and strongly twisted states.

Lastly I explore the phase behavior of tilted hard rods as a model of de Vries smectic behavior and the first order smectic C (SmC) to smectic A (SmA) phase transition. The free energy cost of azimuthal rotation of a molecule away from the local tilt direction is calculated via umbrella sampling. This calculation is used to map the hard rod system onto a lattice spin system which shows a cross-over from a continuous to first order phase transition as the tilt of the rods is increased.

This analysis offers a natural explanation of the first order SmA-SmC phase transition common to de Vries smectics.

Dedication

To my wife, family, and friends without whom this would not have been possible.

Acknowledgements

There have been many people that have helped me in this endeavor. First, my advisor Dr. Matthew Glaser whose guidance made this all possible. The LCMRC and Dr. Noel Clark for providing an environment and support network that allows this sort of research at this university. Dr. Paul Beale who helped with analytic theory and provided insight that I wouldn't have found on my own. Dr. Joel Eaves who introduced me to several advanced Monte Carlo techniques, guided me in my understanding of them, and challenged me to defend my positions when I claimed that I did. Dr. Per Rudquist and Leo Radzihovsky for helpful discussions regarding de Vries smectics along the doubtless forgotten colleagues that have helped me along my way.

In addition, thank you to my colleague Robert Blackwell whose long discussions on programming technique have been integral to the quality of the work presented here, and to my friend Dr. Patrick Alken who kindly provided me with a place to stay and human contact while finishing my dissertation away from home.

The funding of this research is almost entirely by NSF MRSEC Grant number DMR-0820579.

Contents

Chapter	
1	Introduction 1
1.1	A note on the pseudocode notation 3
2	Phase prediction in Hard Core/Soft Shoulder systems 5
2.1	Monte Carlo sampling and structure prediction 10
2.2	Finite temperature phase behavior 18
2.3	Zero temperature phase behavior 20
2.3.1	Simulated annealing and structure optimization 20
2.3.2	Structural enthalpy minimization as a convex hull problem 25
2.3.3	Phase Diagram 28
2.3.4	Interesting Phases and Features 31
2.4	Forward prediction self-assembly from geometric arguments 36
2.4.1	Directed design of colloidal phases 37
2.5	Conclusions 41
3	Advanced Monte Carlo sampling 43
3.1	Monte Carlo algorithms with collective mutation 44
3.2	Adding collective moves to the Metropolis algorithm 46
3.3	Swendsen-Wang style collective moves 48
3.3.1	Biased sampling 52

3.3.2	Virtual move Monte Carlo (VMMC)	58
3.3.3	Rejection and efficiency	64
3.3.4	VMMC with augmented potentials	66
3.4	Application of VMMC and VMMC _{aug}	67
3.5	Conclusion	75
4	Spontaneous twist in bundles of rods	77
4.1	The model and methodology	79
4.1.1	Umbrella sampling, multiple histogram, and flat histogram methods	82
4.2	The entropy of twist	85
4.3	Identifying the source of twist	87
4.4	Cell theory and orientational caging	89
4.5	Phase behavior	97
4.6	Conclusion	99
5	de Vries liquid crystals	102
5.1	The Implications of the Hollow Cone Model for de Vries SmA	105
5.2	Coarse Graining Transformation	107
5.3	Modeling	109
5.4	Experimental measurements	112
5.5	Discussion	115
6	Conclusion	118
	Bibliography	120
	Appendix	
A	Efficient sampling in Monte Carlo	126

A.1	Optimization of algorithm parameters	127
A.2	Numerical steepest descent	128
B	Odds and ends of soft matter computation	129
B.1	Treating excessive unit cell shear in small systems	129
C	VMMC implementation details	132
C.1	The search order	132
C.2	Comparison to the original VMMC algorithm	133
C.2.1	Reverse balance	133
C.3	The indicator function	134
C.4	Benefits of VMMC over GenGCA	135

Tables

Table

- 2.1 A selection of two dimensional zero temperature phases. Each phase is labeled by a nickname. The core contacts can be seen by touching circles, shoulder contacts are designated via a line drawn between the two circles. The computational unit cell is included in dashed red and black lines. 33
- 3.1 Measured decorrelation times for sample HCRS systems. Decorrelation times are measured from NVT fluid and crystalline micellar systems using 100% Metropolis sampling and 100% VM_{MC}_{aug} sampling. These are scaled to two different metrics of time, the wall time and the number of pair interactions. The pair interactions can be considered more accurate here as they are insensitive to implementation details and more accurately predict the theoretical limit of the algorithm's implementation. Metropolis wins across the board but the difference is particularly striking when dealing with crystalline systems where the cluster moves percolate to the system size and become unusable. 75

Figures

Figure

- 2.1 Illustration of (a) the penetrable sphere potential, and (b) the hard core repulsive step (HCRS) potential. Both include the penetrable sphere step and are thus susceptible to the soft shoulder clustering. Spheres in snapshots of HCRS systems are drawn via these disks with halos providing a pictorial representation of the system's structure and energy. The core areas of these particles cannot overlap due to volume exclusion, and each halo overlap between spheres represents a contribution of ϵ to the system energy. 7
- 2.2 Comparison of the system energy of the isotropic system versus the clustering system in one dimension. (a) Isotropic distribution of particles in the system produces a configuration with nearly double the potential energy than the clustered alternative. (b) The addition of cores to the potential produces the same effect, but the particles cannot collapse completely and occupy the same volume. This produces a competition between the tendency for collapse due to the shoulder potential and the excluded volume of the hard cores. 8
- 2.3 The real space and scaled space computational unit cell with periodic boundary conditions to remove the surface effects from the computation. Scaled space allows for simple calculation of periodic boundary conditions. 15

- 2.4 Behavior of the finite temperature system as calculated by Glaser et al.[27] (a) Approximate finite temperature phase diagram for the hard core repulsive step (HCRS) system with $\kappa = 5$. The phase boundaries mark changes in qualitative behavior between five different structural motifs, fluid micelles (M_1), crystal micelles (M_s), fluid stripes (L_1), crystal stripes (M_s), and crystal inverse micelles (IM_s). For each of these qualitative classes, there tend to be several distinct phases. (b) Example phases of the 3D HCRS system. This shows a qualitatively similar phase sequence of cluster/micellar, columnar, and lamellar often with crystal and liquid crystal versions of each. The higher density lattices are not shown in either diagram. 19
- 2.5 The number of phases that at least touch the minimum enthalpy curve in the HCRS system in two dimensions. This serves as a measure of the number of minimum phases that need to be considered to accurately measure the minimal enthalpy configuration set, which is independent of the method by which we generate the candidate configurations. This grows at least linearly with κ and thus increases the complexity scaling of the problem by a factor of κ 24
- 2.6 Enthalpy versus pressure p and the dual space representations of the stable-phase determination process (each enthalpy curve ($h(p) = u + pv$) is mapped to a point in the (u, v) plane); (a) $\kappa = 2$, (b) $\kappa = 5$. Because the HCRS system involves infinite forces, only the pv term in the enthalpy depends on the pressure. The enthalpy curves make it simple to see what pressure ranges correspond to what minimum enthalpy structures. The dual space representation gives us a better, less cluttered view of the simulated annealing results. You can easily identify how well the annealing is working, see how many unique configurations are being generated, identify metastable states and poorly optimized states, and see states that are stable only at single points in the pressure. 26

- 2.7 The sequence of stable phases at $\kappa = 5$ as determined by simulated annealing. This corresponds to the diagrams in Figure 2.6(b). The sequence shows a sequence of cluster phases, a sequence of stripe and ‘broken stripe’ phases, and a sequence of simple lattices with vacancies. Lines connecting spheres represent shoulder interactions. 29
- 2.8 The zero temperature phase diagram for hard sphere repulsive step potential. The low energy/density phases are determined and are identified, where possible, by their highest symmetry state, structures in Table 2.1. The pressure is plotted on a logarithmic scale for convenience only as the pressure differences tend to grow as the energy of the phase increases. Shaded regions denote regions of *entangled lattices*, though the three shown here are in no way comprehensive list. At the top of the diagram, *magic κ s* are listed which correspond to a new shoulder interactions and, typically, radical changes of the phase behavior. The complexity grows with the value of κ . We are unable to determine the nature of the phase behavior in highly complex regions of the diagram like the upper left portion of the plot. 32
- 2.9 Examples of the entangled lattice phenomenon for (a) the square lattice and (b) the triangle lattice, and (c) the honeycomb lattice. The lines represent shoulder-shoulder contacts. In entangled lattices, the shoulder-shoulder contacts define new sublattices which can be partially decoupled from each other. In the square lattice there are two oblique sublattices rotated by approximately 45° from the original. In the honeycomb lattice, there are two triangle sublattices. In the partially expanded triangle lattice, there are three triangle sublattices. In (b) and (c), the sublattices are decoupled from each other as they have no internal degrees of freedom. In (a), the sublattices can deform by becoming more or less oblique. This causes core-core contacts which causes the two lattices to be coupled. In the case of the case of the oblique lattices, this results in a deformed honeycomb like structure which is distinct from the other deformed honeycomb like structure with which it shares a phase boundary, see Figure 2.8. 35

- 2.10 The kappa values chosen needed to promote desired crystal structure and how they correspond to the geometry of the phase. In (a) the honeycomb lattice $\kappa = \sqrt{3}$ is chosen as the ratio of the next nearest to the nearest neighbor distances. In (b) the diamond lattice $\kappa = \sqrt{\frac{8}{3}}$ is chosen as the ratio of the next nearest to nearest neighbor distances. 38
- 2.11 The zero temperature phase sequence for the 3D HCRS $\kappa = \sqrt{\frac{8}{3}}$ system. This phase sequence shows a variety of novel phases including a random stacked diamond lattice in the pressure range of $p = [1.84, 4.34]$. This random stacking can to be removed by using a continuous generalization of the HCRS pair potential. The colors are there to help distinguish the third dimension. All spheres of a given color are in the same plane. 40
- 3.1 An explanation of the pairwise energy terms in the Wolff and GenGCA algorithms. The presentation of these two algorithms is just done for illustrative purposes. These same properties apply to any system that matches the criteria for these Swendsen-Wang style algorithm's restriction of a self-inverse operator that is a symmetry of the Hamiltonian. Note that when considering the special class of failed links where $j \notin \mathcal{C}$, $\epsilon_{ij}^{(\mu \rightarrow \nu)} = \epsilon_{i'j}^{(\nu \rightarrow \mu)}$ and $\epsilon_{ij}^{(\mu \rightarrow \nu)} = \epsilon_{ij}^{(\nu \rightarrow \mu)}$. This is distinct from the behavior of linked particles where $\epsilon_{ij}^{(\mu \rightarrow \nu)} = \epsilon_{ij}^{(\nu \rightarrow \mu)}$ and $\epsilon_{i'j}^{(\mu \rightarrow \nu)} = \epsilon_{i'j}^{(\nu \rightarrow \mu)}$ 56
- 3.2 An explanation of the pairwise energy terms in the VMMC and related algorithms. This is much like Figure 3.1 except that there are factors in the $i \in \mathcal{C}, j \notin \mathcal{C}$ product that do not cancel. This can be traced back to the fact that the r_3 is not equal to r_2 , which is due solely to the fact that the VMMC moves are not self-inverse. 59
- 3.3 A comparison of the results of Tavares et al. to our calculation using various VMMC procedures. Only the lowest temperature distribution is in the model's nematic phase, while the higher two temperatures are in the isotropic phase. 72

- 3.4 A comparison of the acceptance rates of the VMMC_{aug} algorithm using the $\lambda = 0$ augmented potential (isotropic linking potential) versus the normal VMMC procedure. In this histogram, the zero bin indicates moves that were rejected for whatever reason. The VMMC_{aug} method makes it more likely for larger clusters to be successfully moved, but at the cost of lowering the overall acceptance rate. This is seen by the fact that the acceptance rates of the VMMC_{aug} method actually drop below the acceptance rates of the standard VMMC procedure in the small cluster limit. This is particularly true of the lowest temperature data which is in the nematic phase. 73
- 4.1 A schematic of (a) the explicit bundle and (b) the ring systems. Colors denote the orientation of the tilt direction mapped to a color wheel. The parameters defining a bundle or ring are the aspect ratio of the rod, L/D , the aspect ratio of the bundle, L/D_b , and the density of the bundle (area density in the case of bundles; linear density, rings). 80
- 4.2 Example snapshots of the explicit bundle system as sampled via NVE Monte Carlo sampling at $\rho = 0.77$ with three different bundle aspect ratios. The system shows the tendency for twist for some system parameters but not others. Bundles with moderate or large bundle aspect ratios tend to display twisting behavior while bundles with small aspect ratios don't. 81

- 4.3 Entropy curves as calculated by umbrella sampling for three different bundle aspect ratios, $L/D_b = \{\frac{1}{2}, 1, \frac{3}{2}\}$, each at three different area densities, $\rho = \{0.7, 0.8, 0.9\}$. For small bundle aspect ratio ($L/D_b = \frac{1}{2}$), the entropy of twist is maximized at $\mathcal{T} = 0$. At larger bundle aspect ratio the entropy is maximized at a finite angle whose value is largely a function of only density, as demonstrated by comparing the $L/D_b = 1$ and $\frac{3}{2}$ bundles. The free energy barrier between left and right handed twist is quite large even for moderate bundle aspect ratios, on the order of $30 k_B T$, which indicates that the untwisted state is extremely unlikely in bundle systems. The low density, low \mathcal{T} behavior in the $L/D_b = \frac{3}{2}$ case is due to the system developing multiple twist domains that average to zero twist rather than sampling the straight configuration. 86
- 4.4 Examples of topological rubber glove behavior in the twisted bundle system. These transition states between the left and right handed configurations are induced during the umbrella sampling procedure by biasing the system against twist. The states have zero average twist but remain chiral. These types of configurations are preferred to the achiral transition configuration when twist is highly favored and occur when the density is small enough that there is enough room for the structure and domain boundary to form. 88
- 4.5 Demonstration of the increase of the available angular range as twist is introduced to the ring model. In the straight (untwisted) configuration at this rod aspect ratio, L/D , linear density, ρ , and bundle diameter, D_b/D , the probe rod has 7° of room in its orientational cage. When we increase the twist of the ring to tilt 17.2° , the orientational cage nearly doubles in size. The cage grows because the curvature of the ring allows for circumferential escape allowing the rods to partially miss each other, as is seen on the right. This increase in accessible angular range corresponds to an increase in orientational entropy making the twisted configuration the ground state. 90

- 4.6 A comparison of the umbrella sampling (solid) and cell theory calculations (dashed) of the entropy for a sampling of system parameters. The calculations have no free parameters other than the scaling factor α in (4.8) which is manifested as an additive constant, s_0 in (4.10). This parameter is fit to the umbrella sampling data revealing the accurate nature of the cell theory approximation. 92
- 4.7 The angular range of the orientational cage as a function of a 100 diameter ring's twist tilt, θ_c , for several different bundle aspect ratios, L/D_b . We see a clear transition from a twisted state to the straight configuration near $L/D_b = 0.8$ 94
- 4.8 The angular range at the edge of bundle of length 10 spherocylinders as a function of the bundle's twist tilt, θ_c , for several different bundle aspect ratios, L/D_b , and densities, ρ . We see a clear transition from a local tendency for twisted behavior to a local tendency for untwisted behavior as a function of the bundle aspect ratio, L/D_b . 95
- 4.9 The orientational cage of the rods in a bundle as a function of the bundle's twist tilt, θ_c , at several densities, ρ , and radial locations within the bundle. This is like Figure 4.8 except that the rods' cages also involve rods at larger radii than itself. The dotted lines indicate unphysical configurations where the cage rods would intersect other cage rods in order to achieve the desired θ_c . The transition as a function of bundle aspect ratio, L/D_b , is less clear here, but the expansion of the interior cage at small radius is still apparent. 96
- 4.10 Phase diagram for the ring system of density, ρ , versus bundle/ring aspect ratio, L/D_b for large bundle diameters, D_b/D . There is a clear transition between the twisted and untwisted states as a function of aspect ratio. In the close packed limit, the transition aspect ratio is one. There is also a weakly twisted phase that shows up in pockets near the twisted to untwisted phase transition. This phase easily switches it's handedness in finite temperature Monte Carlo simulations. The effect of the weakly twisted phase appears to lessen in the large D_b/D limit. 98

- 4.11 The entropy maximizing twist versus ring aspect ratio L/D_b for several densities, ρ , and versus density for several aspect ratios in the large D/D_b limit. There are scans of the phase diagram in Figure 4.10. Aside from the transition, the entropy maximizing twist doesn't vary much with bundle aspect ratio and the twisted state curves collapse upon each other in the large D/D_b limit. The transition is clearly discontinuous. 99
- 5.1 A snapshot from a Monte Carlo simulation of a single layer hard spherocylinder smectic with the hollow cone orientation distribution (color represents azimuthal orientation). This NPT simulation was performed with cone angle $\theta_A = 45^\circ$ and spherocylinder length $L/D = 5$ in the SmA phase near the SmA-SmC phase transition. Molecules in the hollow cone model have the orientation distribution in tilt, θ , and azimuthal angle, ϕ , of the form $f(\theta, \phi) = (2\pi)^{-1}g(\theta)$. The distribution is uniform in ϕ while $g(\theta)$ is narrowly peaked around the cone angle, θ_A , which we take as the delta function, $\delta(\theta - \theta_A)$. Hollow cone smectics show no global polar order in the SmA phase, but produce finite correlated regions. Macroscopic tilt, ψ , is the result of global bias of molecules to one side of the ϕ distribution. 103
- 5.2 Umbrella sampling calculations of the Gibbs free energy of azimuthal rotation within a smectic layer of a single spherocylinder relative to the polarization vector averaged among surrounding spherocylinders within radius $a = 5\sigma$ where σ is the spherocylinder diameter. These calculations are performed within the SmC phase near the SmA-SmC phase transition at different cone angles, θ_A . At larger cone angles a secondary minimum appears at $\phi = \pi$. Superimposed are fits of the γ -potential, see (5.1), where the secondary minimum is excluded from the fitting procedure if present by separately fitting G_{\max} as demonstrated in the 50° fit. As the cone angle is increased the width of the well is reduced, or equivalently, γ increases. 106

- 5.3 The mapping between the hard spherocylinder hollow cone system and the γ -potential XY model when using $a = 5\sigma$ as the coarse graining length scale where σ is the diameter of the spherocylinders. Spin coupling strength, J_{xy} , increases approximately linearly with pressure while the γ values are nearly independent of pressure but increases with larger cone angles, θ_A . While the specific values of J_{xy} and γ depend on a , these trends do not. 108
- 5.4 The $\gamma - T$ phase diagram of the γ -potential spin system showing tricritical behavior. The dashed and solid phase boundaries denote second and first order phase transitions, respectively. In the two dimensional system (■) and systems with weak inter-layer coupling, $J_z/J_{xy} < 0.1$, the phase transition crosses over from continuous to first order between $\gamma = 4$ and $\gamma = 5$. At stronger couplings, we see a sharp shift to larger values of γ as the cumulative potential begins to resemble the planar rotor. 111
- 5.5 The basic bookshelf cell geometry. Smectic layers are perpendicular to the cell substrate. Fields are applied along the cell normal, \hat{x} . The molecules lie on the surface of a cone and are tilted away from the layer normal, \hat{z} , by angle θ_A . The angle the polarization vector makes with the cell normal is ϕ , and ψ is the projection of the molecule onto the $y - z$ plane. 113
- 5.6 Materials W530 and W599 with phase sequences. These chiral de Vries materials show a double peak in current polarization current (or a sigmoidal polarization) which is characteristic of the de Vries SmA phase and of reminiscent of SmC_A phases. These materials both have special tails, perfloronated in W530 and silocarbane in W599, which are meant to support strong layering, a characteristic known to promote de Vries behavior.[63] 114

- 5.7 An example of the raw data from our polarization current measurements. The current consists of several background contributions (a linear resistive baseline, a capacitive jump, and an ‘ion peak’ whose origins are not entirely understood) which have to be removed before analysis. The inset shows the polarization current once the background has been removed. 114
- 5.8 Behavior of materials W530 and W599 near their respective SmA→SmC transitions at a series of temperatures. We see the characteristically large change in tilt and birefringence in the top two graphs. The bottom graphs show the polarization response which has a sigmoidal, or ‘S’ shaped, character. This is equivalent to having double peaked polarization current. 116
- 5.9 Comparison of the polarization density response of de Vries material W530 (symbols) when an electric field is applied across a liquid crystal cell with bookshelf alignment [73] and the model’s predictions for $\gamma = 6$ (curves). The model’s energy units at a given value of γ are set by equating the model’s SmA-SmC transition temperature, T_{AC} , to the transition temperature for W530, 39°C. The polarization density saturation and susceptibility is scaled to demonstrate simultaneous qualitative agreement with several polarization curves near the SmA-SmC transition. At lower temperatures we predict a discontinuous change in P which is not seen experimentally. The continuous experimental behavior may be due to quenched surface disorder in the cell. 117

Chapter 1

Introduction

Though the idea of emergence, that the collective behavior of many body systems does not relate simply to the individual behavior of its constituent parts, dates back for millennia, the concept of emergent phenomena has come to the forefront of condensed matter physics for at least four decades. This shift in thinking is attributed at least in part to the physicist Phillip Anderson who introduced the term *emergent phenomenon* to physics and has made the argument against reductionist reasoning in physics since the early 1970's.[2, 13] There is a long tradition in physics of describing physical systems by minimal models. These models give theorists and experimentalists alike intuition for new problems that they encounter. With the renaissance of emergent physics, these simple models are often treated via computational methods as analytical methods fail to reveal the complex behavior that emerges from systems described by simple interactions but many bodies.

Often this emergent behavior can be seen as a result of frustration, as is the case with the frustration of multiple length scales seen in hard-core soft-shoulder systems in Chapter 2 and the competition between orientational entropy and packing efficiency that maps to a defect condensation mechanism that drives de Vries phenomena presented in Chapter 5. Other times this emergent phenomenon is simply entropically driven and counterintuitive as in the spontaneous twist formation in bundles of rods in Chapter 4.

In Chapter 2, we explore the nature of the hard core soft shoulder system. The study of clustering in purely repulsive soft sphere systems, such as the penetrable sphere model, has led to

surprising results.[52] The fact that these systems cluster at all is counterintuitive by itself. The fundamental explanation of this phenomenon was given by Likos et al. in the framework of density functional theory as a clustering instability.[46] However, long before this discovery, Jagla had been characterizing and mapping the diverse phase behavior of a related system, the hard core soft shoulder (HCSS) systems.[35, 36] HCSS systems combine a soft potential that is susceptible to the clustering instability described by Likos, but includes an excluded volume interaction that directly interferes with the system's tendency to cluster. The complex phase behavior at zero temperature is in many ways just a glimpse of the interesting behavior that the system displays. Glaser et al. showed that this system exhibited liquid crystal phases, something that was previously unheard of for isotropic spheres.[27] This system with an isotropic, purely repulsive, simple pair potential, also displays the capability for directed design of self-assembling lattices that have previously only been thought to arise from highly directional binding such as the diamond lattice.

In Chapter 3, Monte Carlo techniques are developed for the general study of soft matter systems. The virtual move Monte Carlo (VMMC) cluster move by Whitlam and Geissler is examined and the fundamental connection between the VMMC and Swendsen-Wang like algorithms is laid out in exhaustive detail. New techniques are generalized from the VMMC algorithm designed to apply to the hard core repulsive step (HCRS) potential and other systems with hard core interactions.

In Chapter 4, we characterize the tendency for spontaneous chirality in self-assembled bundles of rods. This observation of a chiral ground state for the bundle system is keeping with a larger literature that finds chiral structures are often the ground state of even simple systems with no chiral components.[6, 7, 59, 79, 26] In particular, the work by Gibaud et al. and Yang et al. is based on a system that is remarkably similar to our self-assembled bundles, but their explanation is quite different from our findings. We isolate the source of the twisting phenomenon and identify it as an entropic effect. In particular, we show that one component of the orientational entropy increases as twist forms. This entropic term scales with the rod length and, for even moderately long rods, already makes the twisted state vastly more probable than any untwisted configuration.

In Chapter 5, we extend the understanding of the de Vries smectic A (SmA) liquid crystal phase by offering an explanation for the characteristic first order phase transition. While competing models exist, by and large, each of the properties of de Vries smectics have been explained within the framework of the hollow cone model. In this model the de Vries phase is considered to be a smectic phase with local smectic C (SmC) like structure with of a finite in-layer correlation length for the C director. This phase has global SmA like order but the tilted nature of the molecules in the model provide many predictions about the de Vries phase's properties. These include the anomalously small change in layer spacing and birefringence, the giant electroclinic effect and the anomalously large change in birefringence during electroclinic switching or over the SmA-SmC phase transition, and a set of relations between tilt, birefringence, and polarization.[11] In this work we demonstrate that the characteristic first order phase transition of de Vries materials is also a simple result of the hollow cone model.[41]

1.1 A note on the pseudocode notation

This thesis contains a significant amount of algorithmic analysis and development. I use pseudocode to keep the algorithms at a reasonable length and to focus on the important parts of the implementation rather than what I deem to be unessential to understanding the concepts entailed. For example I never discuss implementation of neighbor searching methods, an important part of any condensed matter computational code that allows the algorithm to run in linear time. This is not to say that these are unimportant implementation details, but rather that the reader will understand that these details can be dealt with in an efficient manner and that there are several sources that provide ample details on the subject such as Frenkel and Smit's book and Allen and Tildesley's book.[24, 1] In each algorithm I leaned toward making the algorithms easier to understand and more compact using mathematical notation whenever possible. These choices hopefully make the algorithms easier to reason about and, if one wished to implement, simple enough to translate into a lower level, real machine implementation.

I use free parameters throughout my functions and procedures which should universally be

considered parameters that can be tuned during the execution to optimize performance. These parameters are marked with with a “ \sim ” over them, for example \tilde{x} . Section A.1 discusses the importance of optimization in Monte Carlo sampling, but for the purposes of the algorithms discussed here there is no need to clutter the code with the explicit optimization of these parameters.

I use the words *function* and *procedure* a bit differently than their standard mathematical sense, but much more rigorously than the colloquial usage amongst programmers. Here a function is taken in the standard mathematical sense with the exception that there can be an implicit dependence on certain parameters, particularly optimization parameters, and that they may depend on a pseudo-random number source. Optimization typically happens at the highest level of the program and infrequently so I don't feel that this greatly hinders the meaning of the word function. The word function here is really meant to convey the meaning that this code takes parameters and returns answers and at no point is there modification of the parameters that are passed to the function. A particularly fastidious reader with a penchant for functional programming should consider them implicit arguments to every function that needs them.

As the primary technique in this thesis is Monte Carlo simulation, I use the function `random` to provide random numbers. This function takes an argument and returns a number between zero and that argument (not including the argument itself). To simplify the code, if that argument is clearly an integer, then the function returns an integer between zero and that number (again, not including the number itself). As an additional helper function, I use the `random_vector(x)` function to provide random vectors. These vectors are of dimension d , the dimensionality of the system, and each element is between $-x$ and x if x is a number. If x is a unit cell matrix, then it returns a random vector in the unit cell.

Chapter 2

Phase prediction in Hard Core/Soft Shoulder systems

Controlled self-assembly of colloidal systems has long been a goal of condensed matter physics and nanotechnology in general. The ability to create arbitrary designer lattices is necessary for many applications, but most of the structures that engineers are interested in are far from the minimum free energy for self-assembly in standard colloidal spheres. For a specific family of isotropic, repulsive pair potentials, so-called *hard core soft shoulder* (HCSS) potentials, competition between a ‘soft shoulder’ clustering instability and excluded volume interactions produces a rich variety of crystalline and liquid crystalline phases in both two and three dimensions. The rich phase behavior of HCSS systems dramatically expands the range of possible modes of colloidal self-assembly, and raises the interesting prospect of control of colloid self-organization by rational design of effective pair potentials.

Systems of particles with isotropic repulsive pairwise interactions typically display an isotropic fluid phase and a handful of crystalline phases (hexagonal lattice in two dimensions and FCC or BCC in three dimensions) owing to the tendency of mutually repelling objects to distribute themselves as uniformly as possible in space to minimize their overall repulsive energy. However, this behavior is not universal, as recognized as early as 1977 by Grewe and Klein.[32] For a particular family of isotropic repulsive pair interactions, a clustering instability can result in highly inhomogeneous thermodynamic states. More generally, for this class of pair potentials, the free energy of a homogeneous distribution of particles can be reduced through the formation of particle clusters, with a characteristic cluster size that depends on the range of the repulsive interaction potential.[46]

A mean field treatment of the isotropic fluid phase predicts that clustering will occur at sufficiently large density if the Fourier transform of the pair potential is negative at finite wavevectors.[46] This clustering behavior is known as ‘soft shoulder clustering.’ Soft shoulder clustering should be distinguished from clustering driven by competition between short range attraction and long range repulsion which can also produce clustering in colloidal self-assembly.[64] These two mechanisms lead to cluster sizes that scale with different parameters of the interaction potential. In systems with short range attraction and long range repulsion the characteristic cluster size is governed by the competition between surface and bulk free energies, that is, by the relative strength of the attractive and repulsive interactions. In soft shoulder systems, the characteristic cluster size is determined by the range of the soft interaction potential.

The prototypical system that displays soft shoulder clustering behavior is the penetrable sphere model with pair potential shown in Figure 2.1(a). At densities high enough that the mean inter-particle separation is smaller than the step diameter σ_s , the potential energy of a uniform distribution of penetrable spheres can be reduced through the formation of clusters of characteristic size $\lambda \sim \sigma_s$.

To first understand the nature of the clustering behavior we consider the energy of a collection of penetrable spheres. An isotropically distributed system of penetrable spheres at sufficiently high density can reduce its energy through the formation of inhomogeneous structures with characteristic size on the order of the shoulder diameter, as illustrated in Figure 2.2(a). Consider a linear chain of penetrable spheres at a low enough density that neighboring particles are non-interacting. If this chain is then homogeneously compressed to a higher density, the interaction energy increases to $N\epsilon$. The dimerization of the chain will reduce the energy by a factor of two. Hence clustering is energetically favorable in the penetrable sphere model. The clustering behavior can be seen as a direct result of the ‘flatness’ of the penetrable sphere potential; once a pair of particles are within the interaction range, there’s no further energetic penalty to core collapse, when two particles become completely coincident. In this case the ground state of the system consists of fully collapsed groups of particles with no overlap between the groups.

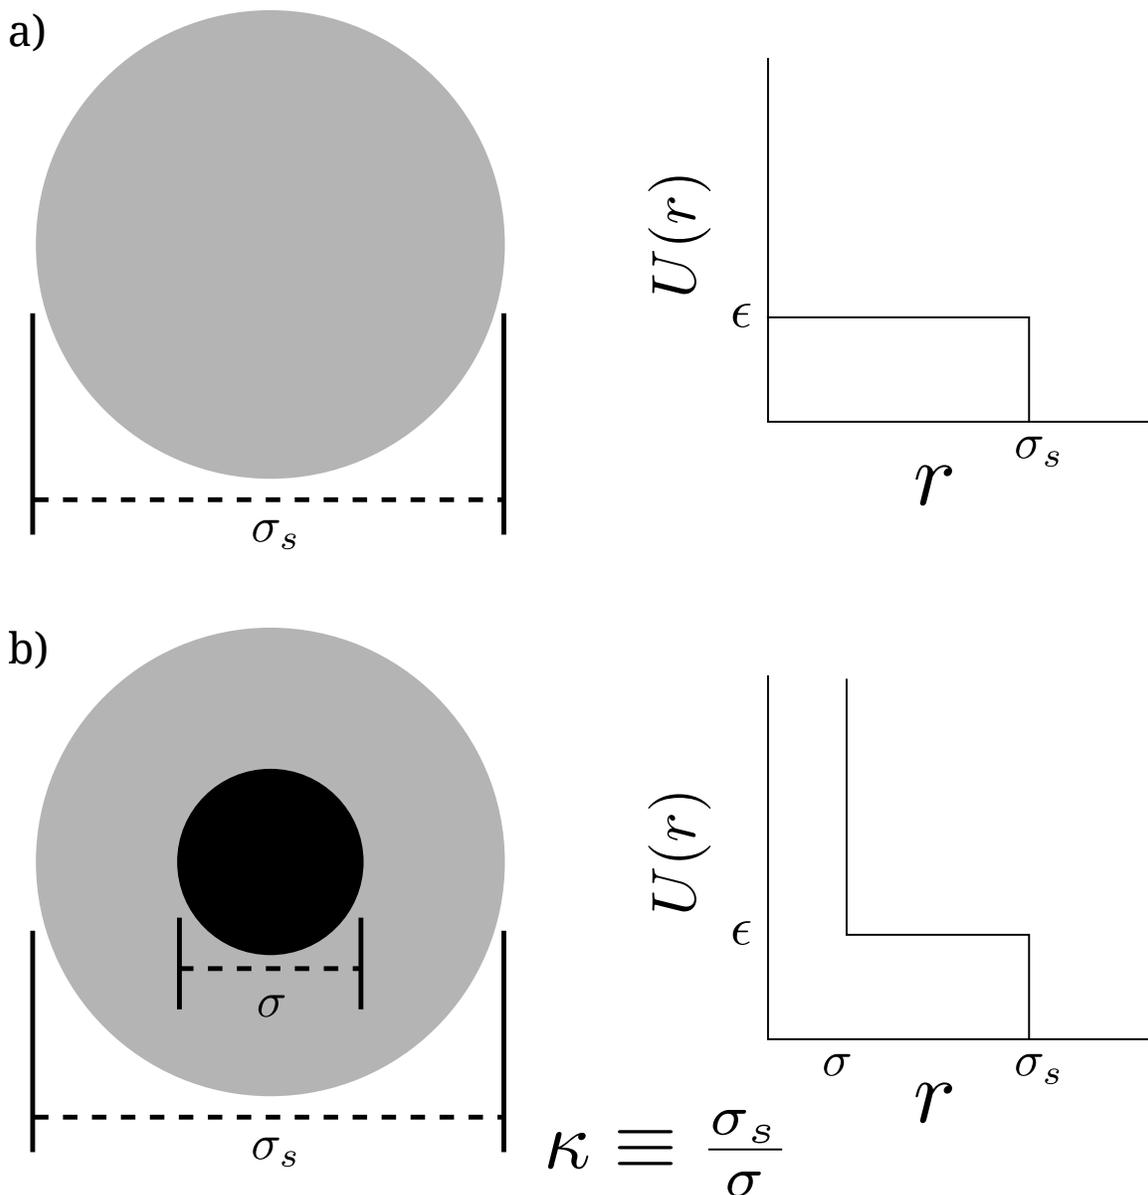


Figure 2.1: Illustration of (a) the penetrable sphere potential, and (b) the hard core repulsive step (HCRS) potential. Both include the penetrable sphere step and are thus susceptible to soft shoulder clustering. Spheres in snapshots of HCRS systems are drawn via these disks with halos providing a pictorial representation of the system's structure and energy. The core areas of these particles cannot overlap due to volume exclusion, and each halo overlap between spheres represents a contribution of ϵ to the system energy.

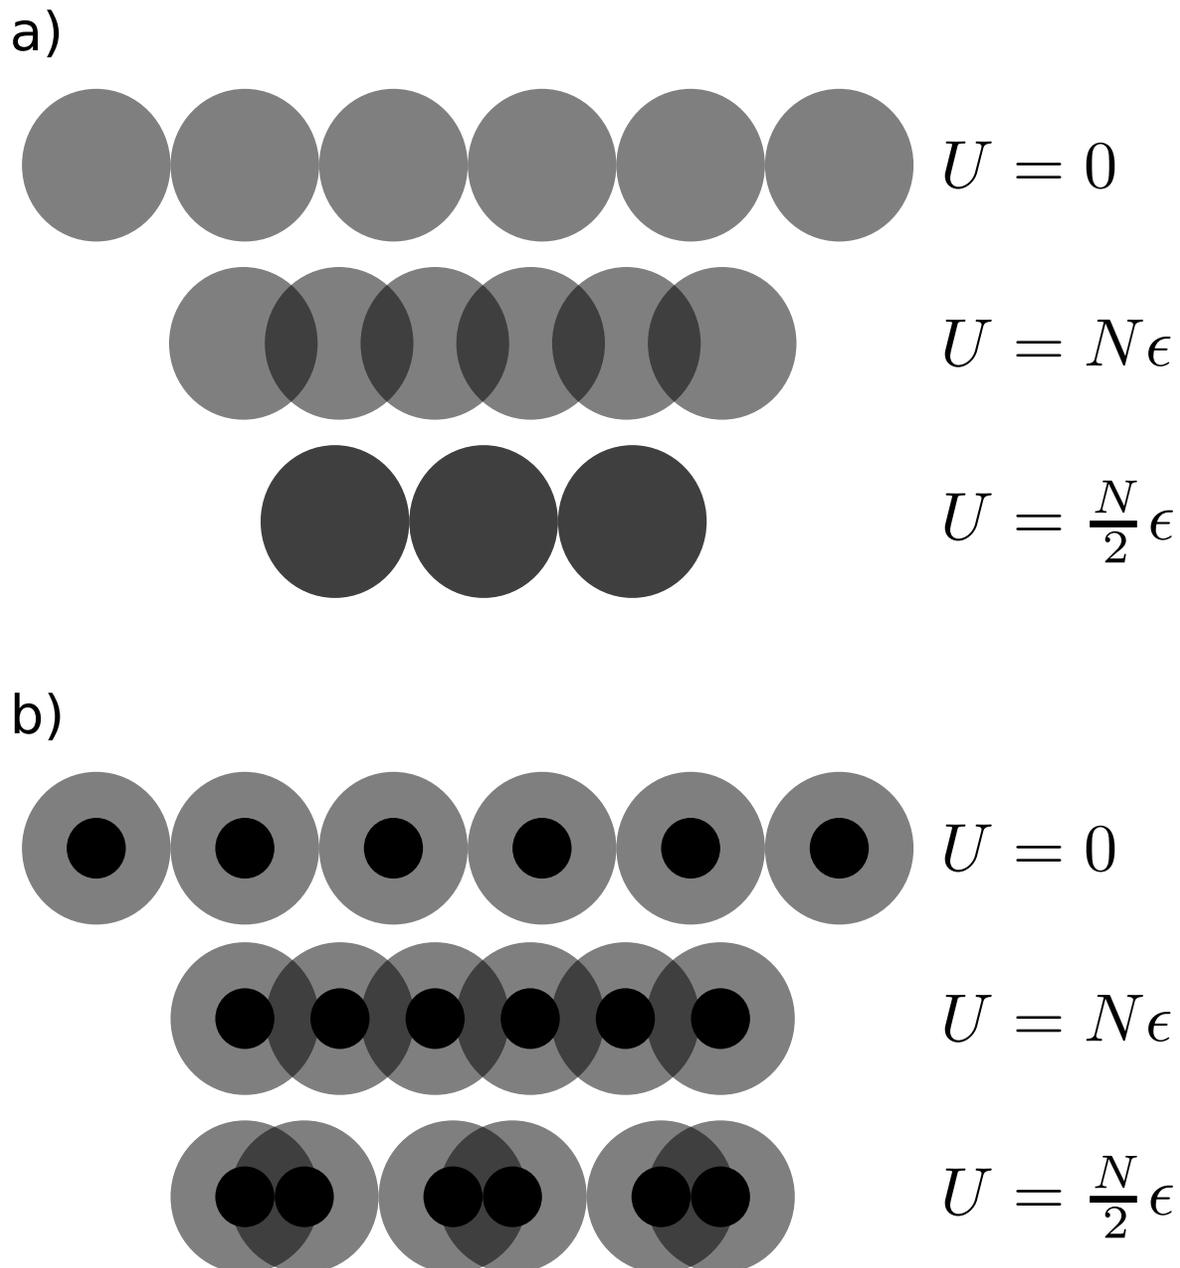


Figure 2.2: Comparison of the system energy of the isotropic system versus the clustering system in one dimension. (a) Isotropic distribution of particles in the system produces a configuration with nearly double the potential energy than the clustered alternative. (b) The addition of cores to the potential produces the same effect, but the particles cannot collapse completely and occupy the same volume. This produces a competition between the tendency for collapse due to the shoulder potential and the excluded volume of the hard cores.

While it's obvious how our argument works in the penetrable sphere model, it's not obvious what constitutes sufficiently flat in other systems. Many potentials other than the penetrable sphere potential exhibit similar soft sphere clustering behavior. These potentials include linear and non-linear ramp, generalized exponential for $\gamma > 2$, and Fermi distribution potentials.[46, 35, 36, 52] As an example, the generalized exponential model ($v(r) = \exp[-(r/\sigma_s)^\gamma]$) smoothly approaches the form of the penetrable sphere potential as γ becomes large. We see that clustering only occurs with values of γ greater than 2. The dimensionality of the system also comes to bear; linear ramp potentials show clustering in two and three dimensions, but not in one. This is where the Likos criterion is useful. It provides a simple test for determining whether a potential can induce clustering behavior. Likos et al. considered the conditions under which an isotropic fluid is unstable towards the formation of a density wave.[46] Their mean field treatment showed an instability can develop if the Fourier transform of the pair potential is oscillatory. We've investigated a number of different models via NPT Monte Carlo and have found the Likos criterion to be an extremely accurate indicator of the tendency to cluster.

Mladek et al. mapped out the phase diagram of the generalized exponential model in three dimensions, using Monte Carlo and density functional theory.[52] They found excellent agreement between theory and simulation, but the phase behavior was dominated by simple close packed structures with multiple occupancy, triangle lattices in 2D, BCC and FCC in 3D. The behavior changes dramatically when excluded volume interactions are introduced in addition to soft shoulder potentials, see Figure 2.1(b) and Figure 2.2(b). The inclusion of the hard core does not disrupt the clustering instability and the Likos criterion still applies, but the hard core does preclude the possibility for complete core collapse. The competition between the soft shoulder potential clustering and the excluded volume constraints introduces a novel type of frustration that leads to incredibly rich phase behavior. While the hard sphere repulsive step potential was first studied as a model of crystalline polymorphism for cesium and cerium, the complex phase behavior wasn't appreciated until the zero temperature phase diagram was studied by Jagla.[80, 35, 36] He looked at the hard sphere repulsive step and related HCSS systems reporting a diverse variety of minimal

enthalpy crystal structures. More recently, this class of systems has investigated using Monte Carlo sampling, genetic algorithms, lattice models, and density functional theory revealing the system’s incredibly diverse phase behavior including liquid crystal phases, quasicrystals, and interesting crystal lattices.[49, 8, 27, 57, 19]

In this chapter, I describe Monte Carlo studies of the phase behavior of HCSS systems, in both 2D and 3D, focusing on a technique to systematically determine the zero temperature phase behavior. Because the clustering is enthalpically driven, the phase behavior becomes increasingly complex with decreasing temperature. We find a variety of crystal structures including basic lattices, open crystal lattices such as the honeycomb and diamond lattices, stripe, cluster, and void structures, and snub square tilings hinting at the parameters necessary for entropy stabilized quasicrystalline phases at finite temperatures. Due to the extraordinarily rich phase behavior and the very large number of metastable states, merely mapping the zero temperature phase behavior is a formidable task. We then apply the diverse phase behavior of these HCSS systems to design a method of directed self-assembly where a target lattice is chosen and the pair potential is tuned to promote its self-assembly. This technique is applied to the self-assembly of diamond lattices.

2.1 Monte Carlo sampling and structure prediction

The primary technique I use in studying the systems in this chapter and thesis as a whole fall under the technique of Monte Carlo sampling. As this is the basis of most of the results shown here, I feel that it is useful to have a brief overview of the basic theory and methods of Monte Carlo sampling. A reader that feels comfortable with basic Metropolis importance sampling may wish to skip this section, but note that this also serves as the background for Chapter 3.

In the field of condensed matter physics,¹ Monte Carlo refers to a method to measure ensemble averages of the equilibrium distribution of a system by randomly mutating the system

¹ I draw a strong distinction between the definition of “Monte Carlo” used in condensed matter physics and the definition used in computer science. While these definitions are undoubtedly commensurate on some deep level, in practice (and in particular, in this thesis), the definitions and results from computer science poorly describe the type of analysis physicists use and thus the work is of limited use.

state, μ , into another feasible system state, ν , according to a transition matrix $k_{\mu \rightarrow \nu}$. This mutation process is independent of previous states of the system and thus defines a discrete time Markov process, or Markov chain. Such a chain has a special property that if the chain is ergodic,² meaning in a finite system you can reach any state from any other state given a finite number of steps, the limiting distribution of the occupancy of the states is independent of the initial system state.

This Markov chain can be used to generate samples corresponding to the equilibrium distribution of the system. This is done by executing the chain, performing a random walk through the system's configuration space where the transition rates are biased by the ratio of the Boltzmann factors of the initial and final state.[24]

The correctness of such a scheme can be motivated by looking at the master equation for each microstate in the system, see (2.1). The master equation simply encodes the fact that the rate of change of the probability that a system will occupy a given microstate, σ_i , is equivalent to the sum of the rates of transition into that state minus the rates of transitions out of that state. In (2.2) we have used the fact that in equilibrium the probability of any microstate is constant. This states that the total transition rate into a state must equal the sum of the transition rate out of the state. In Monte Carlo sampling we usually take it one step further and impose a *detailed balance* condition, used in (2.3), where the equilibrium transition rate from states i to j is equal to the equilibrium rate of the reverse transition from j to i . When we have imposed detailed balance, we only deal with two states: the current state, $\mu \equiv \sigma_i$, and the final state, $\nu \equiv \sigma_j$.

$$\dot{P}(\sigma_i) = \sum_j P(\sigma_j)k_{j \rightarrow i} - P(\sigma_i) \sum_j k_{i \rightarrow j} \quad (2.1)$$

$$0 = \sum_j (P(\sigma_j)k_{j \rightarrow i} - P(\sigma_i)k_{i \rightarrow j}) \quad (2.2)$$

$$P(\sigma_i)k_{i \rightarrow j} = P(\sigma_j)k_{j \rightarrow i} \quad (2.3)$$

$$\frac{k_{\mu \rightarrow \nu}}{k_{\nu \rightarrow \mu}} = \frac{P(\nu)}{P(\mu)} = e^{-(k_B T)^{-1}[E(\nu) - E(\mu)]} \quad (2.4)$$

Equation (2.3) is the core of Monte Carlo sampling in condensed matter physics. All of the algorithms used in this thesis derive from this basic equation. Equation (2.3) does not fully

² This is a sufficient condition not a necessary one, but we rely on ergodicity later in the theory.

determine the $k_{\mu \rightarrow \nu}$ values. The probability of a given microstate is given by Boltzmann statistics, $P(\mu) = \mathcal{Z}^{-1} e^{-\beta E(\mu)}$, thus any definition of the $k_{\mu \rightarrow \nu}$ values that satisfies (2.4) is sufficient to satisfy the master equation and thus guarantee sampling of the system's equilibrium distribution.

To see how this is applied to a sampling scheme, we examine the Metropolis algorithm. The Metropolis algorithm is the first and still most common method of performing Monte Carlo sampling in many condensed matter systems due to its simplicity and intuitive nature and it is the method used in this chapter.[51] The basic Metropolis sampling method is quite simple as shown in Algorithm 1.

Algorithm 1 Perform Metropolis sampling starting with initial configuration \vec{r} , reciprocal temperature β , and step size δ . The procedure's output is done through the call to `emit`, which will send the generated configuration to be consumed by another program. The function `PBC` is defined in Algorithm 2. The functions `random` and `random_vector` are as described in Section 1.1.

```

1: procedure METROPOLIS( $\{\vec{r}\}, M$ )
2:   loop ▷ Loop until desired statistical accuracy is reached.
3:      $i \leftarrow \text{random}(\text{size}(\{\vec{r}\}))$  ▷ Randomly select a single particle in your system.
4:      $\Delta\vec{r} \leftarrow \delta \times \text{random\_vector}(1.0)$  ▷ Select a small random displacement.
5:      $\{\vec{r}'\} \leftarrow \{\vec{r}\}_{\vec{r}_i \rightarrow \vec{r}_i + \Delta\vec{r}}$  ▷ Displace single particle by  $\Delta r$ .
6:      $\Delta E \leftarrow E(\{\vec{r}'\}) - E(\{\vec{r}\})$  ▷ Compute the energy difference.
7:     if  $\text{random}(1.0) < e^{-\beta\Delta E}$  then
8:        $\{\vec{r}_i\} \leftarrow \text{PBC}(\{\vec{r}_i'\})$  ▷ We accept the move, update the system state.
9:     end if
10:    emit( $\{\vec{r}'\}$ ) ▷ Emit configuration data for ensemble sampling.
11:  end loop
12: end procedure

```

$$P(\mu) \times \alpha_{\mu \rightarrow \nu} \times \text{acc}_{\mu \rightarrow \nu}(\mu, \nu) = P(\nu) \times \alpha_{\nu \rightarrow \mu} \times \text{acc}_{\nu \rightarrow \mu}(\nu, \mu) \quad (2.5)$$

$$P(\mu) \times \text{acc}_{\mu \rightarrow \nu}(\mu, \nu) = P(\nu) \times \text{acc}_{\nu \rightarrow \mu}(\nu, \mu) \quad (2.6)$$

The Metropolis sampling method further constrains the detailed balance condition from (2.3) by imposing that the transition rates $k_{\mu \rightarrow \nu}$ are composed of two parts, a *move proposal rate*, $\alpha_{\mu \rightarrow \nu}$, which is symmetric and independent of the system state, and an *acceptance criterion*, $\text{acc}_{\mu \rightarrow \nu}$, which in general depends on the initial and final states of the system. The Metropolis detailed balance condition can be written as (2.5). In the algorithm, the move proposal rate, $\alpha_{\mu \rightarrow \nu}$, is

simply defined by line 4 of Algorithm 1. Transitions are proposed between systems that differ by a small random displacement made to a single particle in the system. The stipulation that move proposal is symmetric and independent of the system state ensures that $\alpha_{\mu \rightarrow \nu} = \alpha_{\nu \rightarrow \mu}$ allowing the Metropolis detailed balance condition to reduce to (2.6). The acceptance criterion, $\text{acc}_{\mu \rightarrow \nu}$ is defined on line 7 of Algorithm 1. As seen in the listing, if the system energy decreases, then the move is accepted; if it increases, the move is accepted according to the ratio of the Boltzmann factors. This satisfies Equation 2.4 as when the system energy is increased by a transition, $k_{\nu \rightarrow \mu} = 1$ and $k_{\mu \rightarrow \nu} = e^{-[E(\sigma_j) - E(\sigma_i)]}$, and when the system energy is decreased by the transition, $k_{\mu \rightarrow \nu} = 1$ and $k_{\nu \rightarrow \mu}^{-1} = e^{-[E(\sigma_j) - E(\sigma_i)]}$. Together, the move proposal rate and the acceptance criterion define a transition rate that satisfies the master equation in its restrictive, equilibrium, detailed balance form and thus will sample from the equilibrium distribution of the NVT ensemble.

It should be noted that nearly all of the transition rates $k_{\mu \rightarrow \nu}$ are zero in this scheme. We exclude these vanishing transition rates from our summations in (2.1 - 2.3) without any change to the analysis. Further, due to the complete basis set defined by the move proposal method, it is trivial to show that ergodicity is retained; any system state can be reached from any other system state via a finite number of transitions within the set of non-zero $k_{\mu \rightarrow \nu}$ values. However, the local nature of the Metropolis transitions hints at one of the biggest barriers of Monte Carlo sampling, the difficulty of efficiently sampling statistically independent system states. We will see in Chapter 3 that ergodicity considerations like this don't ensure efficient sampling and can, in effect, cause a loss of ergodicity for finite sampling.

Due to the guarantee of a unique limiting distribution of the Markov chain, the correctness of this method doesn't depend on initial state. This means that, given enough mutation steps, the state of the system will eventually reach a region of phase space that has a significant Boltzmann weight. Even with this guaranteed limiting behavior, the initial samples that are produced from the Markov chain can be a poor representation of the limiting distribution. Including them in any analysis would introduce systematic error to the desired ensemble averages. Thus, it is standard to perform an initial series of mutations and discard the samples before actually computing ensemble

averages.

This process is known as equilibration and it tends to greatly improve the accuracy of the calculation for a given amount of effort. It is often difficult to predict exactly how many iterations should be sacrificed to the process of equilibration. Often we pick a given fixed number of iterations and confirm that they were enough after the fact by looking at the autocorrelation times and relaxation of various fluctuating properties of the system like energy in constant temperature Monte Carlo. There are several ways to approximate the autocorrelation time; one particularly good method of this is the blocking analysis by Flyvbjerg et al., though it tends to be overly conservative.[22]

Most interesting systems have several modes of relaxation and each can have vastly different time scales for a given Monte Carlo transition matrix, $k_{\mu \rightarrow \nu}$. This means that while there are some general tools for determining if a system has equilibrated or not, they provide necessary but not sufficient conditions. In reality each system needs to be considered on an individual basis with intimate knowledge of the relaxation modes to truly predict proper equilibration. To keep this discussion general, I use a function `EQUILIBRATED_P` without definition as a place holder for a function that should determine if a sequence of samples represents a proper equilibration for the purposes of the given system and calculation.

In the systems we are concerned with it is almost always necessary to employ periodic boundary conditions³ in order to eliminate surface effects. These boundary conditions are a generalized form of toroidal boundary conditions. We define a unit cell by d unique lattice vectors, $\{\vec{a}\}$. These lattice vectors define a unit cell matrix whose columns equal the lattice vectors, $M \equiv [\vec{a}_1; \dots; \vec{a}_d]$. The unit cell matrix has an inverse and the two matrices, M and M^{-1} define a reversible transformation of points between the ‘real space’ coordinates (r -space) and the ‘scaled space’ coordinates (s -space) defined in (2.7) and depicted in Figure 2.3. The orthogonal nature of s -space simplifies many aspects of the computation including reducing the boundary condition to standard toroidal

³ The notable exception is the computations in Chapter 4 where the effect is intimately related to the surface effects.

boundaries.

$$\{\vec{s}\} = M^{-1}\{\vec{r}\} \quad (2.7)$$

$$\{\vec{r}\} = M\{\vec{s}\} \quad (2.8)$$

Algorithm 2 A function for computing the position of within the fundamental unit cell given a point in the system. To simplify the algorithm listings, it is convenient to allow this function to take single vectors or sets of vectors.

```

1: procedure PBC( $\{\vec{r}\}, M$ )
2:    $\{\vec{s}\} \leftarrow M^{-1}\{\vec{r}\}$ 
3:   for all  $\vec{s}_k \in \{\vec{s}\}$  do
4:     for  $i = 1..d$  do
5:        $\vec{s}_k[i] \leftarrow s_k[i] - \text{round}(s_k[i])$ 
6:     end for
7:   end for
8:    $\{\vec{r}\} \leftarrow M\{\vec{s}\}$ 
9: end procedure

```

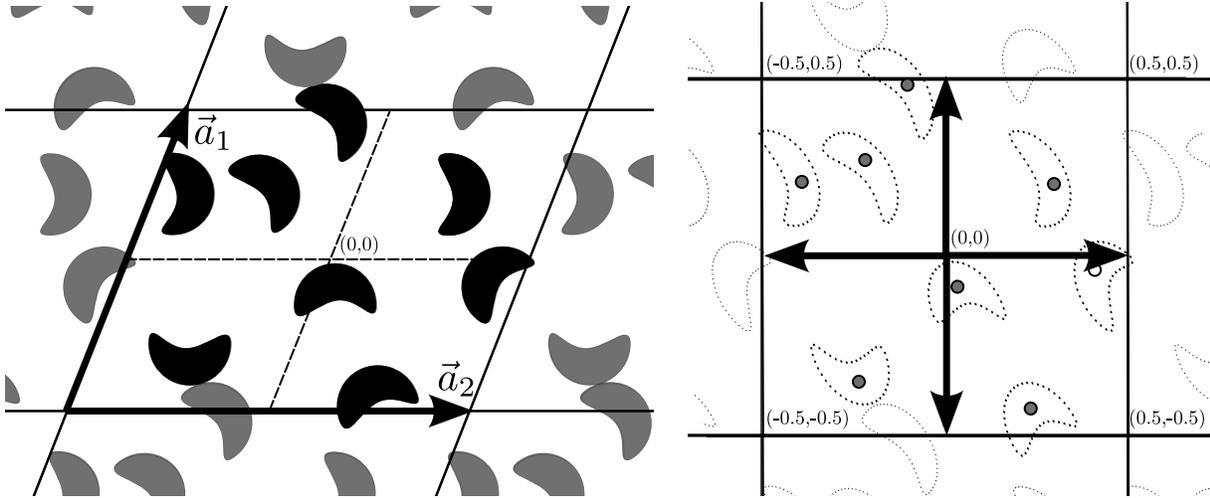


Figure 2.3: The real space and scaled space computational unit cell with periodic boundary conditions to remove the surface effects from the computation. Scaled space allows for simple calculation of periodic boundary conditions.

Metropolis Monte Carlo lends it self very naturally to NVT sampling. With a minor modification the ensemble sampling can be done against the NPT ensemble. For this to work, we only

need introduce new transitions in the Markov chain that allow the system to change its volume and include that volume dependent component of the free energy in the Boltzmann factors. Volume changes can be done in several ways, depending on your system. For systems that should express only isotropic strain, simple isotropic scaling of the system is often enough to sample any configuration that the system may need to explore. Due to the small unit cells and crystal structure, the systems in this chapter require a more general form for unit cell mutation, one that includes shears as well as volume changes.

To accommodate this modification, I factor the process of performing a state mutation, or a *move*, into separate functions with a specified interface. Each move should take a system state⁴ and optionally return a new system state and any associated information regarding the change such as energy changes.⁵ The move functions are randomly selected based on a discrete probability distribution, $\{P_{\text{move}}\}$. The volume change involves a small change to M parametrized by small parameter Δ . This change can be done in several different ways, such as a simple random mutation of a unit cell vector or one of an explicit shear or volume change move. Any way that it is performed, it involves a volume change against external pressure p and a potential energy change which are included in the acceptance criterion.

Just like volume changing moves modify the ensemble from the Canonical ensemble (NVT) to the isobaric-isothermal ensemble (NPT), other ensembles sampling can be achieved by adding moves, and hence introducing new transitions, $k_{\mu \rightarrow \nu}$, in the Markov chain. The Grand Canonical ensemble (μ VT) can be achieved by adding insertion/removal moves and even the so-called Gibbs ensemble by adding both insertion/removal and volume changing moves.[24] In principle, micro-canonical sampling can be achieved via NVT Metropolis in the limit of zero temperature, but this doesn't work in practice for systems with continuously varying pair potentials. A better approach is to look to the rejection-free methods from Chapter 3 or to use molecular dynamics which is

⁴ For most cases, this is a set of coordinates and a unit cell matrix, but in other cases, such as the Potts and Ising models of Section 3.3, this might not be the most appropriate representation.

⁵ In principle the final state of the system is all that is necessary to find the new ensemble averages. With sufficient annotation, the final state can also provide those values in an efficient manner consistent with Appendix A. I return extra information to clearly show where this can be made efficient when this is not trivial.

Algorithm 3 Perform Metropolis sampling starting with initial configuration \vec{r} , unit cell matrix M , reciprocal temperature β , and step size δ .

```

1: procedure METROPOLIS( $\{\vec{r}\}, M$ )
2:   loop                                     ▷ Loop until desired statistical accuracy is reached.
3:
4:     MOVE  $\leftarrow$  {
5:         DISPL with probability  $\tilde{P}_{\text{displ}}$            ▷ Select move.
6:         VOLUME_CHANGE with probability  $\tilde{P}_{\text{vol}}$      }
7:
8:      $(\{\vec{r}'\}, M') \leftarrow \text{move}(\{\vec{r}\}, M)$            ▷ Perform move.
9:      $\Delta E \leftarrow E(\{\vec{r}'\}, M') - E(\{\vec{r}\}, M)$      ▷ Compute the energy difference.
10:    if  $\text{random}() < e^{-\beta\Delta E}$  then
11:       $(\{\vec{r}'\}, M) \leftarrow \text{PBC}(\{\vec{r}'\}, M')$        ▷ We accept the move, update the system state.
12:    end if
13:     $\text{emit}((\{\vec{r}'\}, M))$                                ▷ Emit configuration data for ensemble sampling.
14:  end loop
15: end procedure
16:
17: function DISPL( $\{\vec{r}\}, M$ )
18:    $i \leftarrow \text{random}(\text{size}(\{\vec{r}_i\}))$            ▷ Randomly select a single particle in your system.
19:    $\Delta\vec{r} \leftarrow \delta \times \text{random\_vector}(1.0)$      ▷ Select a small random displacement.
20:    $\{\vec{r}'\} \leftarrow \{\vec{r}\}_{\vec{r}_i \rightarrow \vec{r}_i + \Delta\vec{r}}$      ▷ Displace single particle by  $\Delta r$ .
21:   return  $(\{\vec{r}'\}, M)$ 
22: end function
23:
24: function VOLUME_CHANGE( $\{\vec{r}\}, M$ )
25:    $M' \leftarrow M + \tilde{\Delta}$                                ▷ Apply a small change to the unit cell.
26:    $\{\vec{r}'\} \leftarrow M' \times M^{-1} \times \{\vec{r}\}$        ▷ Compute new positions.
27:   return  $(\{\vec{r}'\}, M')$ 
28: end function

```

naturally microcanonical.

2.2 Finite temperature phase behavior

While HCSS systems show incredibly complex phase behavior, their limiting behavior is quite simple. For $\kappa \equiv \frac{\sigma_s}{\sigma} \approx 1$ or for $k_B T/\epsilon \gg 1$ the model reduces to the hard sphere system which has been studied extensively. For $\kappa \gg 1$, the HCSS system approaches a simple bounded potential system which, while subject to the Likos clustering instability, the particles within a cluster completely collapse onto one another resulting in a lattice of clusters which exhibits simple close packed lattices.[52] The phase behavior in between these limits is much more difficult to determine.

Glaser et al. used Metropolis NPT Monte Carlo to determine the approximate finite temperature phase diagram for the hard core repulsive step system shown in Figure 2.4(a).[27] With increasing κ , the hard core repulsive step system exhibits increasingly complex phase behavior dominated by highly inhomogeneous states reminiscent of lyotropic liquid crystals or block copolymers. Figure 2.4(b) illustrates a diverse subset of the 2D phase behavior showing snapshots from the Monte Carlo Markov chain and slices through the three dimensional pair distribution functions, $g(\mathbf{r})$, for the $\kappa = 5$ system at four pressures along the $t = 0.5$ isotherm.

With decreasing pressure, a variety of phases are observed, including a hexagonal lattice of voids ('inverse micelles'), periodic arrays of stripes ('lamellae') with either crystalline or liquid-like ordering within the stripes, and hexagonal lattices of clusters ('micelles') with either crystalline or liquid-like intra-cluster ordering. All of the phases shown in Figure 2.4(a) exhibit inhomogeneity on a length scale comparable to the Likos instability length scale for the penetrable sphere system. The behavior of three-dimensional HCRS system is qualitatively similar to the two-dimensional case showing a similar progression of clusters, columnar, lamellar, and fully core connected lattices, see Figure 2.4(b). The cluster, columnar, and lamellar phases show crystal and liquid crystal versions.[27] Again, the 3D phases display inhomogeneity on a length scale comparable to the Likos instability length scale for the 3D penetrable sphere system.

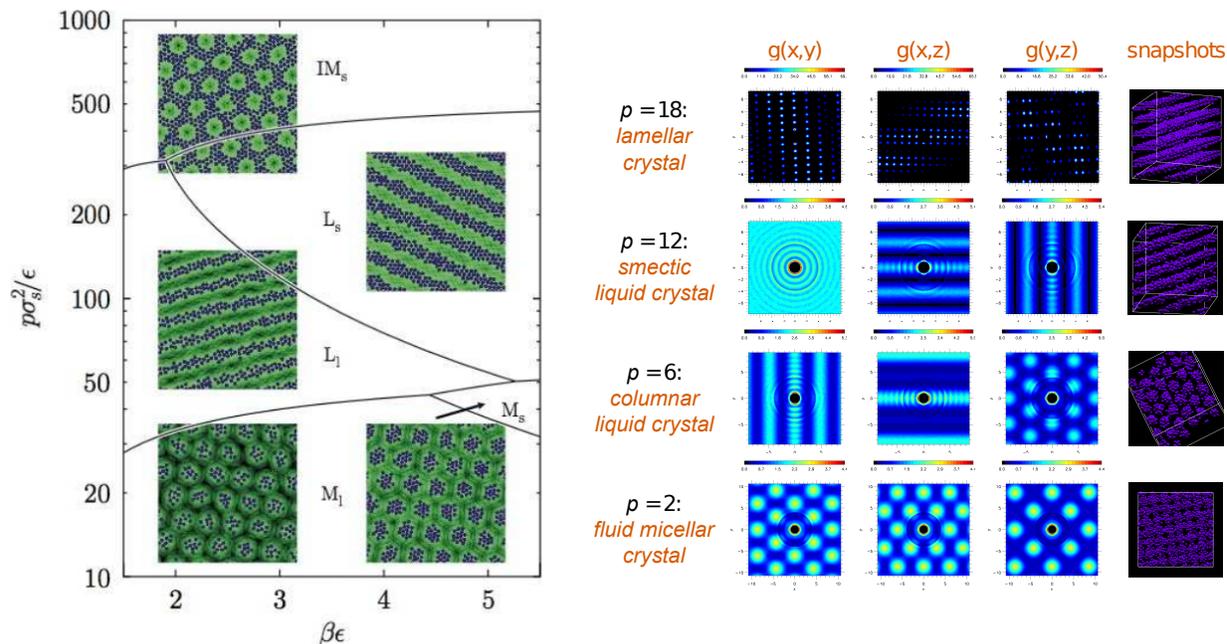


Figure 2.4: Behavior of the finite temperature system as calculated by Glaser et al.[27] (a) Approximate finite temperature phase diagram for the hard core repulsive step (HCRS) system with $\kappa = 5$. The phase boundaries mark changes in qualitative behavior between five different structural motifs, fluid micelles (M_l), crystal micelles (M_s), fluid stripes (L_l), crystal stripes (M_s), and crystal inverse micelles (IM_s). For each of these qualitative classes, there tend to be several distinct phases. (b) Example phases of the 3D HCRS system. This shows a qualitatively similar phase sequence of cluster/micellar, columnar, and lamellar often with crystal and liquid crystal versions of each. The higher density lattices are not shown in either diagram.

2.3 Zero temperature phase behavior

Some of the first steps in understanding the behavior in Figure 2.4 lie in the low temperature behavior where entropic effects do not play a role. While we can only approximate the phase behavior at finite temperature, we can essentially solve for the zero temperature phase behavior exactly. Work by Jagla and others shows that even in this simple case, diverse phase behavior is still demonstrated.[35] One goal we set out to accomplish is to map this behavior in a systematic and automated way. Methods that require minimal human intervention are necessary for this problem as the extreme diversity of phases precludes manual enumeration of ‘likely candidates’ and the sheer number of possible configurations makes manual inspection and characterization a costly, if not impossible, task.

Due to the primary role that enthalpy plays in this process, as temperature decreases, enthalpy becomes increasingly important relative to entropy, and the phase behavior becomes increasingly complex. The enthalpy of a structure is trivial to calculate. The principal difficulty in determining the zero temperature behavior is in enumerating all reasonable ground-state crystal structures for the HCRS system. We find that the complexity and diversity of crystal structures increases rapidly with increasing κ , making the number of possible configurations very large. Work by Arkus et al. has demonstrated the combinatorial growth of possible isolated cluster packings as a function of cluster size N . [3] Each rigid cluster they generate has $dN - d!$ contacts. While the work by Arkus et al. deals with spheres with one length scale, the HCRS system has two, making $2^{dN-d!}$ times as many packings to consider. This direct enumeration is beyond the possibility of computation at this time for even modest cluster sizes and another approach must be used.

2.3.1 Simulated annealing and structure optimization

Since it is beyond our capability to enumerate all possible configurations, some other technique must be utilized to find minimal enthalpy states. While the pure scale of the problem is daunting, this is at its core a global optimization problem, one of the most researched fields of

numerical analysis. For this problem, we use simulated annealing to generate likely configurations with a high probability. Simulated annealing is a probabilistic method for finding the global minimum for a given function. The process performs a Metropolis style random walk over the function domain while slowly reducing the system temperature which reduces the thermal energy available to climb barriers in the function.[40] While simulated annealing isn't the only technique suitable for this task (for example the use of evolutionary algorithms by Pauschenwein and Kahl have proven very useful for this problem), it is by far the most conceptually clear and obvious approach.[57]

The process of simulated annealing is based on the physical technique of annealing that is commonly used in materials manufacturing. It involves elevating a system to a higher temperature where the thermal energy is greater than the barriers between basins of attraction of various free energy minima making substantial reconfigurations possible. This allows the system to explore a greater area of the configuration space than would be otherwise possible. The temperature of the system is slowly reduced, eventually reintroducing the barriers and uncovering the structure of the energy landscape. Granville et al. showed that when cooling is sufficiently slow ($T \sim (\ln t)^{-1}$ where t is the iteration), the probability of finding the global minimum approaches 1.0.[31]

Although simulated annealing is a quite general technique and applies to any optimization problem, we are using it in the most literal sense of its namesake. We are starting with high temperature systems and slowly reducing the temperature allowing the structure to anneal and relax into a low energy configuration via NPT Metropolis Monte Carlo. As the temperature is decreased, the system naturally falls into low enthalpy configurations which defines the stable phase at zero temperature. However, because the cooling schedule of Granville et al. is prohibitively slow, and due the stochastic nature of the process itself, the final minimum is not guaranteed to be the global minimum.

While the simulated annealing process samples many basins of attraction of local minima during the anneal and is much better at finding low enthalpy states than simple quenches, in systems plagued by a multitude of metastable states, such as HCSS systems, the process often results in metastable local minima in the configurational enthalpy landscape. Many runs of the

simulated annealing procedure starting with randomized initial states are required to reach an acceptable probability of finding the global minimum. After many annealing runs, one is left with a set of configurations in local minima, some of which may be the global minimum. The task then becomes selecting the minimum enthalpy configuration within that set of local minima. At any given pressure and set of candidate phases, it is simple to determine the most stable state by comparing enthalpies.

Algorithm 4 Simulated annealing procedure using Metropolis Monte Carlo for system pressure p starting with temperature T subject to an exponential decay by multiplicative factor α .

```

1: procedure ANNEAL( $\{\vec{r}\}, p, T, \alpha, \delta$ )
2:   stream  $\leftarrow$  start(metropolis( $\{\vec{r}\}, p, T, \delta$ ))            $\triangleright$  Start sampling and save stream.
3:   samples  $\leftarrow$  []                                          $\triangleright$  Initialize samples with an empty list.
4:   repeat
5:     samples  $\leftarrow$  append([consume(stream)], samples)  $\triangleright$  Consume sample and prepend.
6:   until EQUILIBRATED_P(samples)
7:   if  $T < T_{\text{final}}$  then                                        $\triangleright$  Exit at low temperature.
8:     first(samples)                                            $\triangleright$  Return the latest state in the sampling.
9:   else
10:    anneal(first(samples),  $p, \alpha \times T, \delta$ )              $\triangleright$  Reduce  $T$  and recurse.
11:  end if
12: end procedure

```

The HCSS system’s minimum enthalpy configuration is a function of the reduced system pressure, $p \equiv \frac{P\sigma_c^2}{\epsilon}$, and the ratio of the shoulder potential diameter to the hard core diameter, $\kappa \equiv \frac{\sigma_s}{\sigma_c}$. For every value of κ , we perform a set of simulated annealing runs at different pressures and initial configurations. Care must be taken to accommodate every possible lattice by using initial conditions that range in the number of spheres from one to some number greater than $N_{\text{max}}(\kappa)$, the number of spheres that lie within interaction range of a central sphere in a close packed configuration. This ultimately produces a large set of candidate structures for each (p, κ) point. While each simulated annealing run does not carry a prohibitive computational cost, taking anywhere from a matter of minutes to over half an hour, in aggregate the cost can be substantial.

As $N_{\text{max}}(\kappa)$ scale with κ^d , on order $O(\kappa^d)$ minimizations, where d is the dimensionality of the system, must be used to accurately determine the stable phase at any given (p, κ) point. Importantly, the computational expense of each of these simulations scales as N^2 , giving the time scaling

of finding the minimal enthalpy configuration as $O(\kappa^{3d})$. However, we are interested in computing the entire phase sequence which will require several of these minimization runs at different pressures. As I will show in Section 2.3.2, minimizations at one pressure be used at other pressures, but the minimization procedure is only likely to find the minimal energy configuration within that configuration's stable pressure range. This means that if there are m different configurations in the phase sequence on a certain κ line, there needs to be at least m stable phase calculations at different (p, κ) points.

For $\kappa \approx 1$ the phase behavior is simple, there are only four stable phases, and thus a relatively sparse grid in pressure can be used. If we chose our pressures perfectly and the simulated annealing did a good job finding the minimal enthalpy structures, then only four (p, κ) phase determination runs would suffice. However, at larger values of κ there are a great many more stable phases on any given κ line, and thus many more (p, κ) points must be used. While the number of stable phase as a function of κ is difficult to estimate *a priori*, based on our results we can empirically see that the number of phases that are among the minimum enthalpy structures scales at least linearly with κ , as shown in Figure 2.5. This corresponds to an increase in the computational expense, giving a total computational cost to compute a κ grid line that scales as $O(\kappa^{3d+1})$.

In addition, spacing between κ grid lines needs to be fairly tight to capture the steeply varying phase boundaries. The end result is on the order of tens of simulation runs at each (p, κ) grid point, hundreds to potentially thousands of simulated annealing runs at each κ grid line, and overall tens to hundreds of thousands of simulated annealing runs in total to cover the range of $\kappa \in [1, 5]$. Our computational effort of over 56,000 simulated annealing runs ultimately proved only sufficient to reliably uncover the phase behavior for systems with κ less than 2.5.

With this series of candidate states, we eliminate identical states and find the candidate that has the lowest enthalpy and identify that as the global minimum state while the others are identified as metastable states. While there is no guarantee that this set of candidate states necessarily contains the minimum state, we can have some assurance that the phase sequence is complete by continuing to add simulated annealing runs until there is a clear steady state in the list of predicted

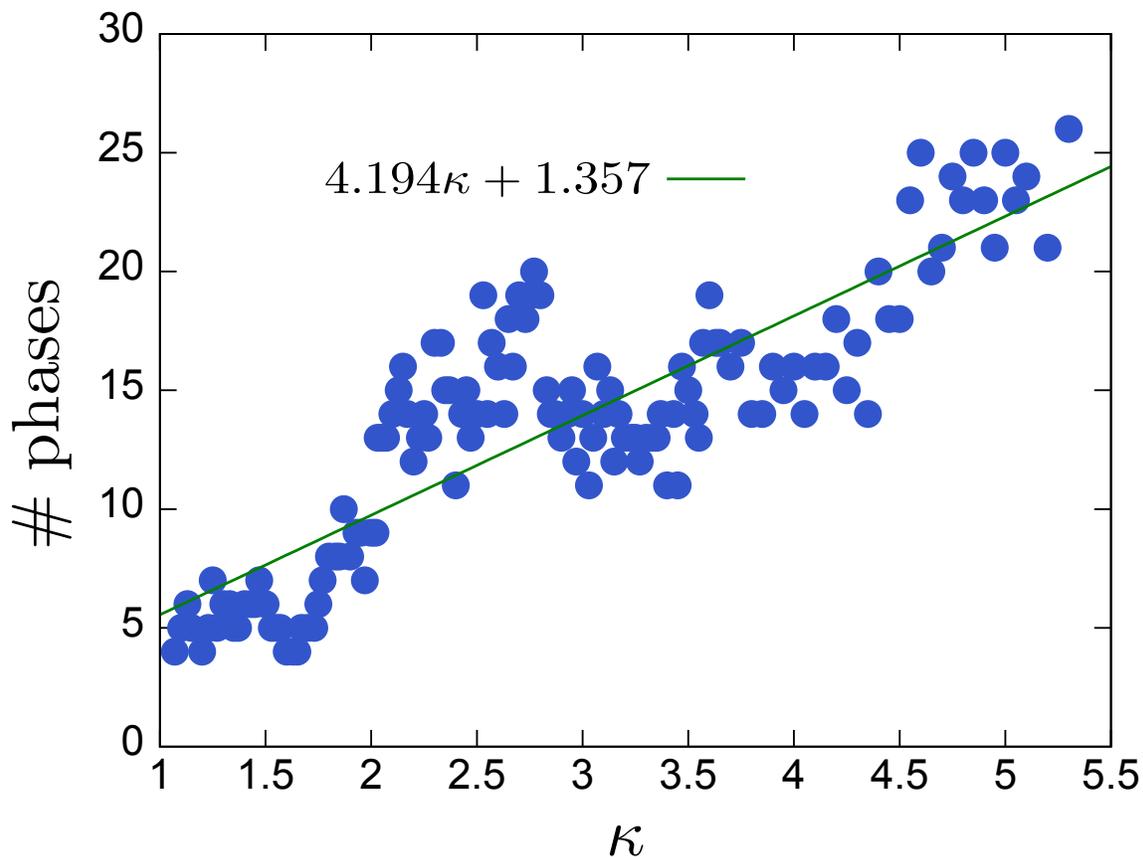


Figure 2.5: The number of phases that at least touch the minimum enthalpy curve in the HCRS system in two dimensions. This serves as a measure of the number of minimum phases that need to be considered to accurately measure the minimal enthalpy configuration set, which is independent of the method by which we generate the candidate configurations. This grows at least linearly with κ and thus increases the complexity scaling of the problem by a factor of κ .

phases. With this procedure it is straightforward to produce the zero temperature phase diagram with no a priori knowledge of what phases might or might not appear. Further, this procedure can be fully automated.

2.3.2 Structural enthalpy minimization as a convex hull problem

Due to the simple linear relationship between pressure and enthalpy in the HCRS system, a single phase annealed at a particular pressure can be extrapolated to the entire pressure range.⁶

All that remains is to find the set of line segments that border the lower-right portion of the Cartesian plane as shown in Figure 2.6. Each intersection between the configurations' enthalpy curves are detected and reported as phase transitions.

The process of determining the stable phases is as simple as determining the set of line segments that define the lowest enthalpy as a function of pressure. This requires finding the intersections between the enthalpy curves, and scales as $O(m \times n)$ for m stable phases and n total candidate phases. A transformation of the problem into a dual space defined by the intercept and slope of the lines (u, v) casts this problem into a classic computational geometry task, the determination of the convex hull of a set of points. Under this map, each enthalpy curve is reduced to a point in the (u, v) space. The convex hull connects the candidates that comprise the structures that minimize enthalpy where each edge in the hull defines an intersection of two enthalpy curves. While computational effort isn't really an issue for the number of candidate structures we are concerned with, extensive work has been done to find optimal and correct algorithms to calculate the convex hull of a set of points including algorithms such as Jarvis March (also known as Gift Wrap), Graham Scan, and Quick Hull which are readily available in computational geometry packages.[14, 37]

The dual space representation is also useful as a tool to view the simulated annealing procedure as a whole as shown on the right side of Figure 2.6. It is easy to gauge how well the procedure

⁶ In the case of other potentials, the volume and energy of the configuration are both functions of pressure, which means that the minimal enthalpy changes as the pressure changes in a non-trivial way. This can still be dealt with by performing stepwise minimizations while scanning in pressure.

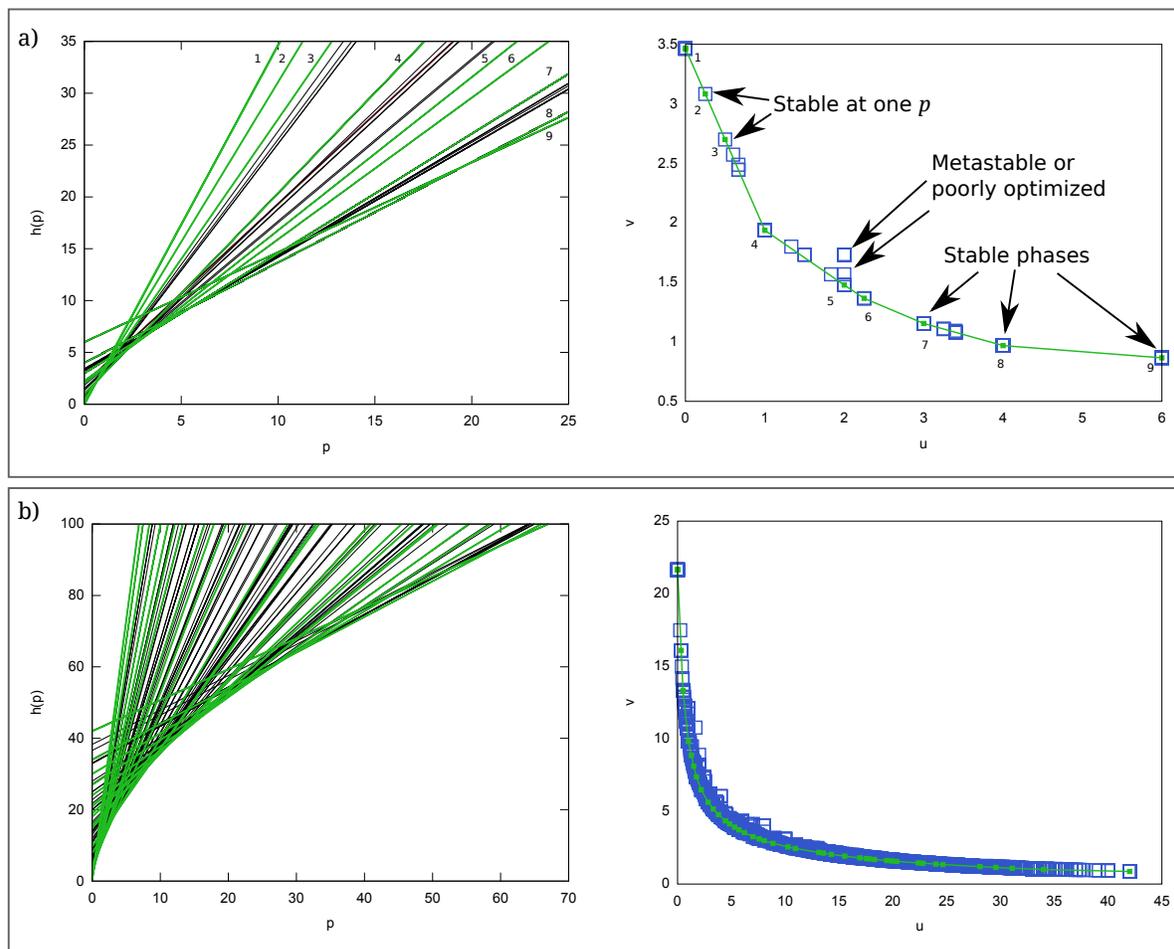


Figure 2.6: Enthalpy versus pressure p and the dual space representations of the stable-phase determination process (each enthalpy curve ($h(p) = u + pv$) is mapped to a point in the (u, v) plane); (a) $\kappa = 2$, (b) $\kappa = 5$. Because the HCRS system involves infinite forces, only the pv term in the enthalpy depends on the pressure. The enthalpy curves make it simple to see what pressure ranges correspond to what minimum enthalpy structures. The dual space representation gives us a better, less cluttered view of the simulated annealing results. You can easily identify how well the annealing is working, see how many unique configurations are being generated, identify metastable states and poorly optimized states, and see states that are stable only at single points in the pressure.

is working by comparing the total set of generated candidate phases with the hull defining the true stable phases. If the two trace out similar curves, as is the case in both the $\kappa = 2$ and $\kappa = 5$ results in Figure 2.6, then the simulated annealing process is working well. We also can get an idea of how many unique structures are being found by the procedure. In the $\kappa = 2$ case, there are relatively few, while in the $\kappa = 5$ case there are many. The number of unique candidate phases is proportional to the number of metastable phases. The near continuum of metastable phases along the equilibrium hull hints at the difficulty of mapping out the zero-temperature phase diagram at large κ . One common occurrence in the phase sequence is to find phases that are stable at only one point in the diagram, such as the phases marked ‘2’ and ‘3’ in the in the $\kappa = 2$ diagrams. These configurations are only stable at precisely $p \approx 0.655$. In this case, these phases are dimerizations or partial dimerizations of the structure where some, but perhaps not all, of the spheres have collapsed to core-core interactions with one of their neighbors. Phases of this type are not stable when $\kappa \leq 2$. In the enthalpy curve representation this corresponds to three or more enthalpy curves intersecting at one point, while in the dual space it represents three or more co-linear phase points. At these points the system phase has more than two equilibrium phases, only two of which have some finite stable pressure range. This is precisely the phenomenon of Jagla’s pentagon tiling which is stable at only one pressure.[36] It turns out that this is an extremely common phenomenon in the HCRS system and we choose to exclude these structures from our phase diagram.

Looking at the 2D $\kappa = 5$ HCRS system we find a zero-temperature phase sequence that includes more than 40 distinct crystal phases between the low-density (shoulder-packed) hexagonal lattice with lattice constant $a = \sigma_s$ and close-packed hexagonal phase with $a = \sigma$, see Figure 2.7. The enthalpy as a function of pressure and the dual space diagram of zero-temperature crystal structures are shown in Figure 2.6. The phase morphologies in the $T = 0$ phase sequence qualitatively resemble those at $T = 0.5$ (e.g., cluster and stripe morphologies), but there are some interesting exceptions, most notably the absence of stable ‘inverse’ phases at $T = 0$. It should be noted that the absence of these inverse phases does not seem to be a failure of the simulated annealing procedure. The process generates several candidate inverse phases, but none prove to be

the minimum enthalpy phase. We also observe morphologies that are not present at $T = 0.5$, for example ‘broken stripe’ phases, which blur the distinction between cluster and stripe configurations, and vacancy lattices, which can be seen as crystal structures of missing particles.

2.3.3 Phase Diagram

Our investigation of the zero temperature phase behavior of HCSS systems has indeed displayed the complex nature of these systems. Figure 2.3.3 shows the zero temperature behavior of the HCRS system and the phases are cataloged in Table 2.1. When comparing the $\kappa = 5$, zero temperature phase behavior to the finite temperature behavior, we note that the sequence is considerably more complicated at zero temperature since the clustering phenomenon is enthalpically driven. These zero temperature configurations act as a starting point to understanding finite temperature behavior by defining structural motifs that recur at finite temperature. Indeed, the same qualitative structures that appear in the finite temperature phase diagram, clusters, stripes, and open and close-packed lattices, are present in the zero temperature diagram.

While the behavior is complex and hard to predict, some simple patterns emerge. The phase sequence from low density close packed to high density close packed is characterized by core collapse and percolation of core-core interactions in particular dimensions. In the two dimensional system, this corresponds to a low density system with clusters, followed by stripe phases where one dimension has collapse to the hard core separation, and finally a collapse to hard core separation in the second dimension producing a series of lattices where core-core interactions have percolated throughout the system. In three dimensions the sequence is to collapse on one dimension into columnar phases, followed by a collapse to layered phases, finally followed by a collapse of the system to a set of fully percolated structures. Each of these classes of systems have several different, often exotic, structures and as κ gets larger the distinction between these groups begins to blur, but the general progression holds.

As stated before, the HCRS system should reduce to the penetrable sphere model at large κ , recovering the phase behavior described by Mladek et al.[52] This seems to be contradicted by

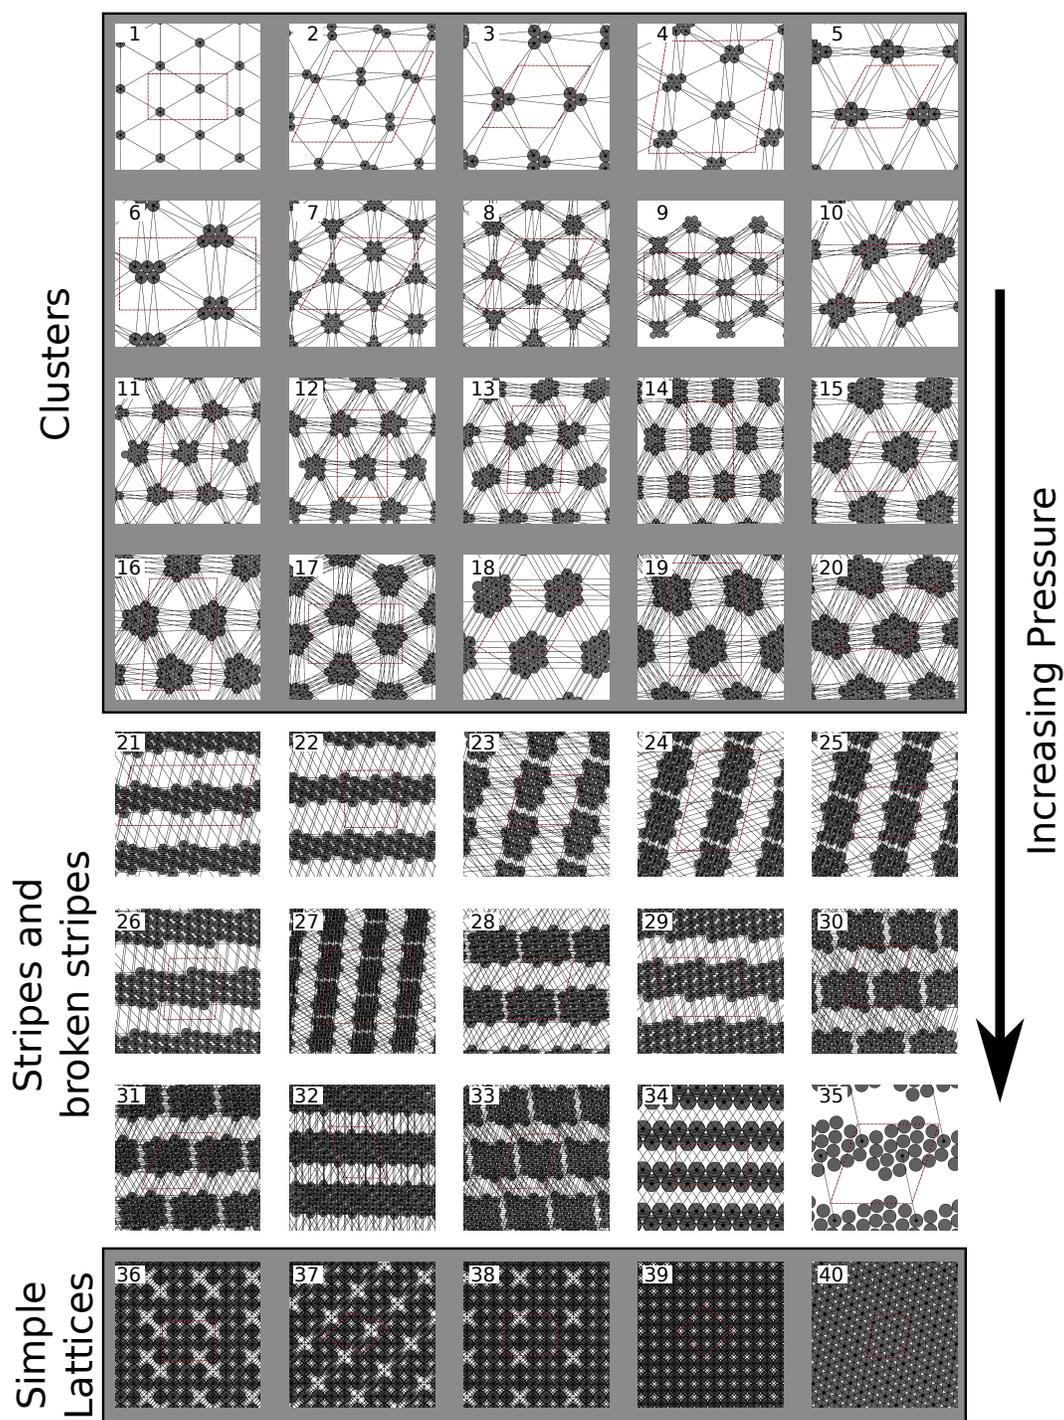


Figure 2.7: The sequence of stable phases at $\kappa = 5$ as determined by simulated annealing. This corresponds to the diagrams in Figure 2.6(b). The sequence shows a sequence of cluster phases, a sequence of stripe and ‘broken stripe’ phases, and a sequence of simple lattices with vacancies. Lines connecting spheres represent shoulder interactions.

the explosion of complexity at increasing κ . The phase diagram in Figure 2.3.3 and phase sequence at $\kappa = 5$ demonstrates how this happens. In the large κ limit, the clustering phases indeed show up and behave similarly to the Mladek et al. results of multiply occupied clusters in simple close packed cluster lattices. The multitude of cluster phases are due to the occupancy number of each cluster, which Mladek et al. addressed, but also the imperfect core collapse and the many ways that these imperfectly collapsed clusters can pack. We see this behavior in the $\kappa = 5$ phase sequence. As an example, compare phase 4 and 5. Their clusters have the same occupancy but differ in structure based on the way the clusters pack. This packing between clusters is coupled to the cluster's internal packing, which is coupled to the occupancy of each cluster. This gives rise to configurations like 7 and 8 which have two types of clusters, each with different occupancy numbers. In the large κ limit, perfect core collapse is recovered and the different packings all become degenerate, restoring the behavior from Mladek et al. Looking at the phase diagram that is only successfully mapped for small κ belies the nature of the large κ limit. Cluster structures don't even appear until $\kappa = 2$ which is already quite large for characterizing the full phase diagram. The non-cluster phases don't play a role in this phase progression in the large κ limit as they require a core-core percolation in at least one dimension. The enthalpy of such a configuration diverges as κ becomes large.

The assumption that the simple micellar clusters dominate provides some predictive power even at moderate κ . The overall enthalpy of the equilibrium phases approaches the functional form of $h(p) \sim \sqrt{p}$ as κ gets large. This is illustrated in the $\kappa = 5$ enthalpy plots in Figure 2.6. This dependence can be understood by considering the $\kappa \rightarrow \infty$ limit of the HCSS model (the penetrable sphere model). With increasing pressure, the 2D penetrable sphere model exhibits a sequence of hexagonal crystal phases with increasing integer site occupancy n . [52] The multiple-occupancy ($n > 1$) cluster phases of the penetrable sphere model are analogous to the finite- κ crystalline micelle phases of the HCRS model.

At $T = 0$, the volume per particle v and potential energy per particle u for the penetrable sphere cluster phases have a simple dependence on n , namely $v_n = \sqrt{3}/(2n)$ and $u_n = (n-1)/2$, and transitions between adjacent cluster phases occur at pressures $p_{n,n+1} = -(u_{n+1} - u_n)/(v_{n+1} - v_n) =$

$n(n+1)/\sqrt{3}$ and enthalpies $h_{n,n+1} = u_n + p_{n,n+1}v_n = n$. Thus, $h_{n,n+1} \sim \sqrt{p_{n,n+1}}$ for large n (large p). Comparing this behavior of the penetrable sphere model to the $\kappa = 5$ HCSS model shows the general $h \sim \sqrt{p}$ trend, but the enthalpy of the HCSS model is significantly larger than that of the penetrable sphere model owing to excluded volume effects, which lead to an increase in the pV contribution to the free energy. A simple first-order estimation of excluded volume effects for the HCSS model can be obtained by increasing the diameter of a n -particle penetrable sphere cluster by an amount sufficient to accommodate a close-packed lattice of n cores, yielding a volume per particle $v'_n = v_n[1 + (\sqrt{2\sqrt{3}n/\pi} - 1)/\kappa]^2$ and the potential energy per particle is assumed to be unchanged. This simple correction to the penetrable sphere free energy leads to very good agreement with the HCSS model for low pressures, but overestimates the enthalpy at higher pressures, where the HCSS model displays more complex morphologies with collapse in multiple dimensions such as stripe and fully percolated structures.

2.3.4 Interesting Phases and Features

The phase diagram in Figure 2.8 contains a lot of information. First there seems to be a cascading effect where at certain values of κ a new phase boundary will appear, often bifurcating several times with increasing κ , and spawn several new phases. This cascade is simple to understand, it occurs when the shoulder interaction reaches another coordination shell in the triangle lattice. When this happens the system must split between at least two phases: one where the shoulders overlap and one where the lattice expands to accommodate no overlaps. This produces a phase boundary, and typically produces several as the lattices can partially expand, keeping the lattice constants in some directions and expanding the lattice in others.

Although most apparent in the hexagonal lattice, this behavior is not limited to coordination shells in the hexagonal lattice. Indeed, any time increasing the shoulder diameter introduces new interactions in a stable phase, it is common to see cusp-like changes in the phase boundaries. Due to

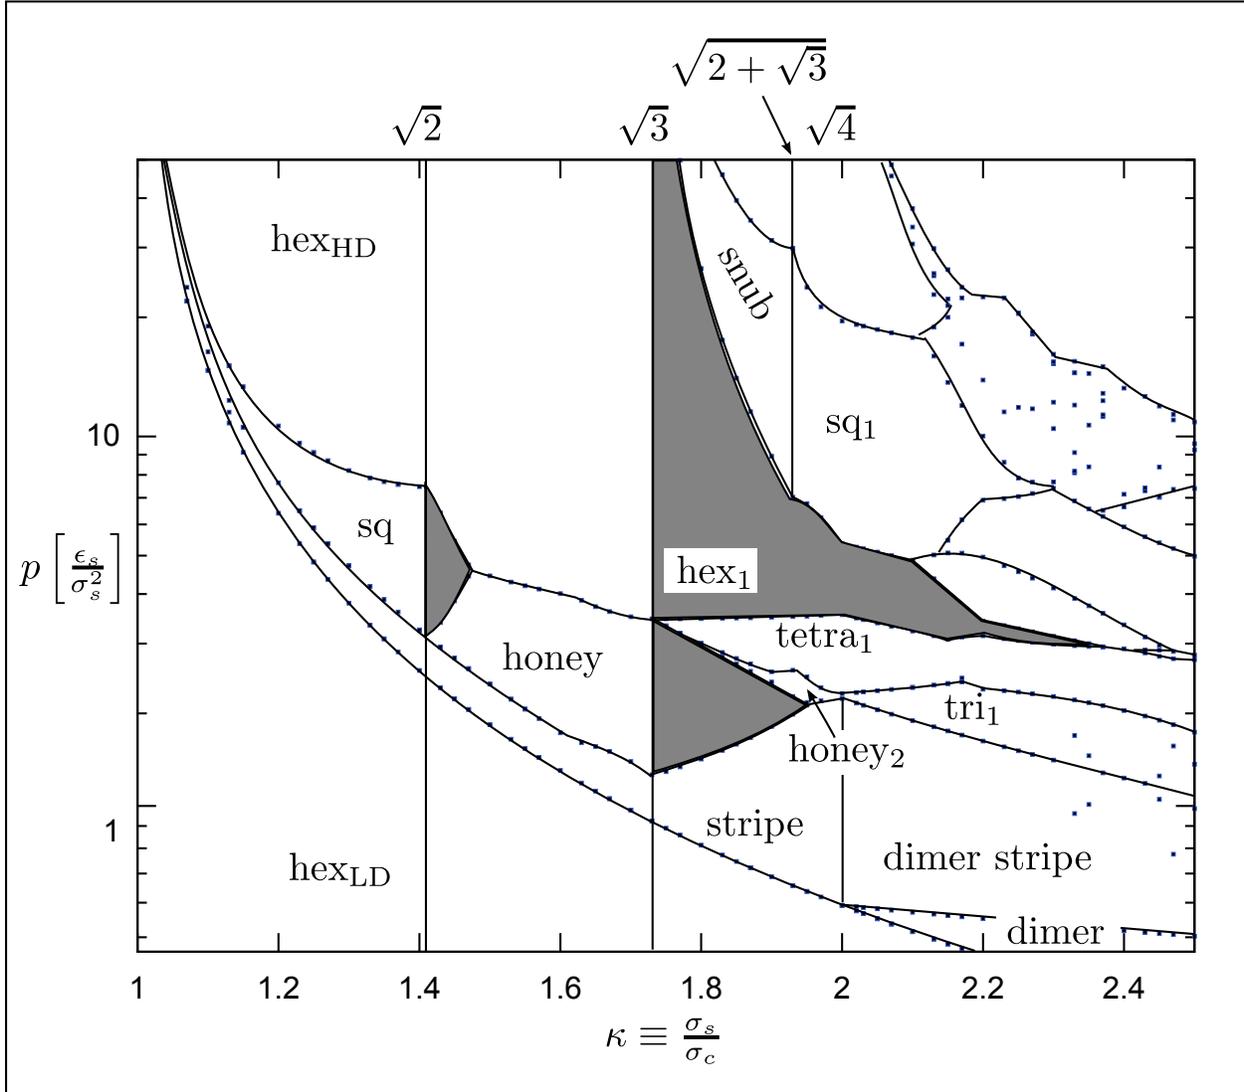


Figure 2.8: The zero temperature phase diagram for hard sphere repulsive step potential. The low energy/density phases are determined and are identified, where possible, by their highest symmetry state, structures in Table 2.1. The pressure is plotted on a logarithmic scale for convenience only as the pressure differences tend to grow as the energy of the phase increases. Shaded regions denote regions of *entangled lattices*, though the three shown here are in no way comprehensive list. At the top of the diagram, *magic* κ s are listed which correspond to a new shoulder interactions and, typically, radical changes of the phase behavior. The complexity grows with the value of κ . We are unable to determine the nature of the phase behavior in highly complex regions of the diagram like the upper left portion of the plot.

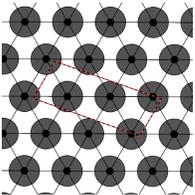
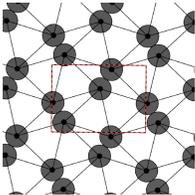
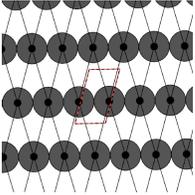
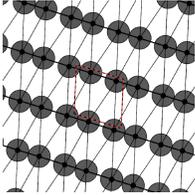
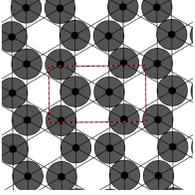
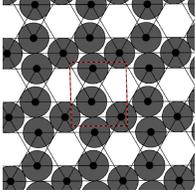
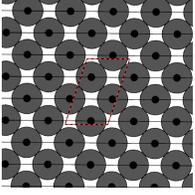
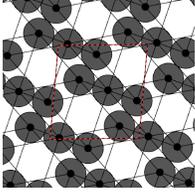
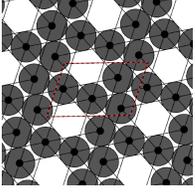
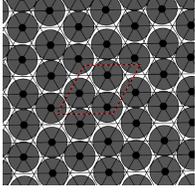
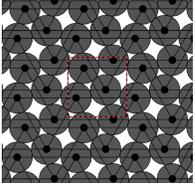
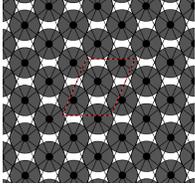
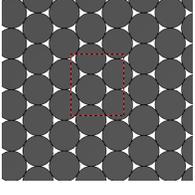
Name	Snapshot		Name	Snapshot
hex _{1d}			dimer	
stripe			dimer stripe	
honey			honey ₂	
sq			tri ₁	
tetra			hex ₁	
snub			sq ₁	
hex _{hd}				

Table 2.1: A selection of two dimensional zero temperature phases. Each phase is labeled by a nickname. The core contacts can be seen by touching circles, shoulder contacts are designated via a line drawn between the two circles. The computational unit cell is included in dashed red and black lines.

the repulsive nature of the pair interaction, the new shoulder contact always increases the enthalpy of the phase. All things being equal, this will decrease the stable range of pressures for a phase. This can be complicated if two adjacent phases in the phase sequence happen to find a new shoulder contact at exactly the same κ .

As Jagla noted, there are a several non-close packed lattices that self-assemble in the 2D HCRS system including honeycomb, oblique, and square lattices, as well as some extremely exotic phases.[36] The range of pressure for these phases tends to be maximized around a certain value for κ , a value that is commensurate for its structure. For the oblique lattice and increasing κ , we see that the range of stable pressure values increases until a perfect square lattice at $\kappa = \sqrt{2}$. We call these κ values that maximize the symmetry, and often the phase's stable pressure range, the structure's *magic* κ at κ^* . Phases with κ less than κ^* show deformed versions of the structure. Phases with κ greater than κ^* show *entangled lattices*, the shaded regions of the square and honeycomb regions in Figure 2.8. Entangled lattices are two or more sublattices with lattice constant σ_s that are only coupled to each other via core-core interactions and, in the case of triangle lattices, can be effectively decoupled from one another as demonstrated in Figure 2.9. These sublattices are generally of the same lattice group as each other but potentially different from the lattice at κ^* . For the ideal HCRS system, these phases can have a finite entropy even at zero temperature, and thus technically have no stable configuration. In the case of triangle sublattices, these systems have undetermined degrees of freedom, however the structure is still macroscopically rigid. Further, if an infinitesimal gradient were added to the shoulder potential this effect goes away and, if the gradient is repulsive, the perfect symmetry of the phase at κ^* is restored. This entangled lattice phase behavior differs from the results found by Jagla, who suggested that the phase region abruptly stops at that κ^* value and the surrounding phases subsume that pressure range. This is provably false by comparing against the enthalpy of the structures in the work presented here.

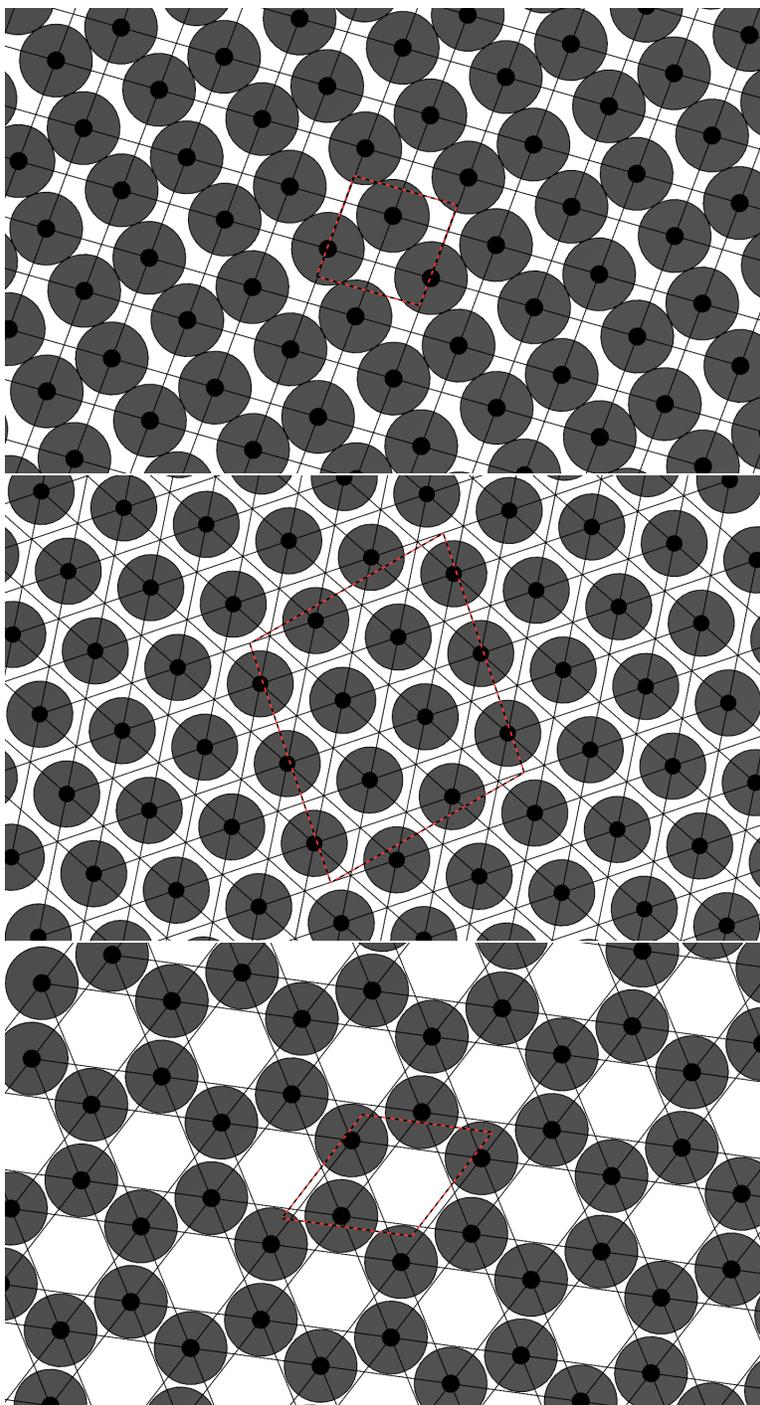


Figure 2.9: Examples of the entangled lattice phenomenon for (a) the square lattice and (b) the triangle lattice, and (c) the honeycomb lattice. The lines represent shoulder-shoulder contacts. In entangled lattices, the shoulder-shoulder contacts define new sublattices which can be partially decoupled from each other. In the square lattice there are two oblique sublattices rotated by approximately 45° from the original. In the honeycomb lattice, there are two triangle sublattices. In the partially expanded triangle lattice, there are three triangle sublattices. In (b) and (c), the sublattices are decoupled from each other as they have no internal degrees of freedom. In (a), the sublattices can deform by becoming more or less oblique. This causes core-core contacts which causes the two lattices to be coupled. In the case of the case of the oblique lattices, this results in a deformed honeycomb like structure which is distinct from the other deformed honeycomb like structure with which it shares a phase boundary, see Figure 2.8.

2.4 Forward prediction self-assembly from geometric arguments

The creation and characterization of appropriate systems where the position of colloids can be controlled to a high degree of precision would enable great advances in the field of engineering nanoscale devices including fields such as photonics, optics, electronics, and solar cells. One technique used for this purpose is direct manual assembly of lattices via robot-aided sedimentation which uses robotic manipulation and sacrificial sites to build desired colloidal structures. While this works, the cost and effort of building a lattice by hand is often too high for many applications. While manual assembly processes don't scale well to larger systems or production, self-assembly uses natural equilibrium states to organize the system in parallel on a massive scale.

Forward prediction of such structures, the ability to declare a target structure and derive the system parameters necessary for its self-assembly, would be the ideal situation. This problem has many technical hurdles. *Versatility* and *predictability* are two of these necessary (but not sufficient) hurdles.⁷ First, one must find systems that show the *versatility* to self-assemble into interesting crystal lattices rather than the standard close packed lattices. Second, one must develop a procedure by which you can systematically *predict* the phase behavior based on careful characterization or, preferably, by tuning system parameters in such a way that you expect the system to prefer certain structural behavior without fully characterizing the phase diagram.

Several directionally binding colloidal systems have been proposed for the purpose of self-assembly of designer phases.[28] Janus particles and patchy colloids, which utilize different chemical environments grafted onto their surface according to a designed pattern, are attractive candidates for self-assembly.[10, 75, 38] In a similar vein, lock and key colloid style schemes have been proposed as a versatile self-assembly system combining directional interaction sites with species-specific binding.[74] Both systems allow for versatility and some level of, typically forward, predictive power.

However, directional interactions are not necessary in all cases. As we have seen, systems

⁷ There are a multitude of important factors that need to be considered. One important one that is not address here is the robustness of the self-assembly mechanism. In order for the self-assembly mechanism to work in realistic scenarios, it should not require extreme manufacturing tolerance or be strongly sensitive to the random variations.

with isotropic interactions, including the HCSS systems, also have the capability to form non-close packed lattices.[35, 36, 57, 60] For simplicity, and potentially due to the difficulty of manufacturing the colloids, an isotropic interaction like those found in colloids with uniformly treated surfaces is much preferred. Work by Rechtsman et al. demonstrated a mechanism for forward design of phases involving an iterative process where you attempt to stabilize a structure by randomly mutating the pair interaction.[60] This shows the ability for isotropic interactions with attraction to form many phases that they target. This work, while promising, produces complex pair potentials that are likely difficult to manufacture in real world systems.

The HCSS systems, on the other hand, are simple in nature and more likely to be realizable in a real world system. Since their diverse phase behavior extends to the zero temperature limit, they make an attractive target for self-assembly. While this diversity of phase behavior caused technical hurdles when computing the zero temperature behavior, it also opens new possibilities. This kind of diversity allows us to create designer phases as we can clearly satisfy both the versatility and predictability criteria for targeted self-assembly. In fact, due to the simplicity of the HCRS potential, it is a simple matter to tune the system potential to a particular target phase or lattice without even computing the zero temperature phase behavior. In the HCRS system, we are able to mandate two characteristic length scales, the core diameter and the repulsive step diameter. If we pick the ratio of these two lengths to be a ratio present in the target lattice, we can promote the self-assembly of that target.

We demonstrate this kind of directed design of phases by targeting the diamond lattice. The existence and stability of honeycomb lattices in two dimensions suggests that diamond lattices are likely a stable phase in three dimensions. For our ratios of length scales in the lattice we picked two coordination shells.

2.4.1 Directed design of colloidal phases

When we examine structures in the 2D phase diagram it is clear why particular structures are promoted at certain values of κ . For example, in the oblique lattice case, we see that the

parallelogram unit cell is stabilized by the core interactions along the unit cell lattice vectors, (\vec{a}_1, \vec{a}_2) , and the shoulder interaction across its diagonal, $\vec{a}_2 - \vec{a}_1$. We can see that a perfect square lattice is promoted by tuning our κ to the square lattice's magic κ , $\kappa^* = \sqrt{2}$, the ratio of the diagonal of the square to its side. While this is strictly observation after the fact, with the proper foresight, we could have self-assembled several interesting structures using nothing more than this insight. The honeycomb lattice, for instance, has a ratio of the nearest neighbor to next nearest neighbor of $\kappa = \sqrt{3}$, see Figure 2.10(a). Looking at the phase diagram, we indeed see that this is the honeycomb lattice's $\kappa^* = \sqrt{3}$ and this is the location in the phase diagram where it is most stable.

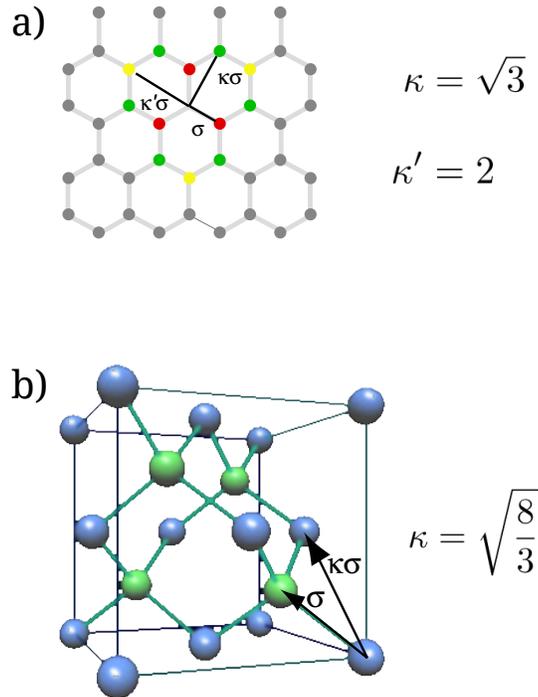


Figure 2.10: The kappa values chosen needed to promote desired crystal structure and how they correspond to the geometry of the phase. In (a) the honeycomb lattice $\kappa = \sqrt{3}$ is chosen as the ratio of the next nearest to the nearest neighbor distances. In (b) the diamond lattice $\kappa = \sqrt{\frac{8}{3}}$ is chosen as the ratio of the next nearest to nearest neighbor distances.

This is not limited to zero temperature phases. Entropy stabilized phases often exhibit a

corresponding zero temperature tiling. If we can predict the location of that tiling, we can have some predictive power regarding the location of the finite temperature phase. In the case of entropically stabilized twelve-fold quasicrystals made from random packings of squares and triangles, the zero temperature tiling is the snub square. This tiling has nearest to next nearest neighbor distance ratio of $\kappa = \sqrt{2}$. However, since $\sqrt{2}$ is the κ^* of the square tiling, and we know that there is no snub square in that zero temperature phase sequence, we must try a different κ . We used the next-next nearest to nearest neighbor ratio, $\kappa = \sqrt{2 + \sqrt{3}}$, and we see the snub square tiling and deformations of that tiling in nearby values of κ .

Any κ value that corresponds to a ratio of a distance between spheres in the structure divided by the nearest neighbor distance in the system can, in principle, be used. This provides some leeway when choosing κ . There is a trade off between coordination number and proximity of other coordination shells. The higher the coordination number the more likely the phase due to the higher energy of core collapse at that distance. Since coordination numbers tend to grow at larger distances, one might expect that a large κ value is optimal. This is a bad idea, however, as coordination shells become more closely spaced at larger distances leading to tight restrictions on the tolerances allowed for σ_s to avoid coincidental promotion of competing structures. In addition it is almost certainly more difficult to produce an experimental system with large interaction lengths that can pass arbitrarily through other colloids in the system. It's not clear where the proper balance is for each target system. There is some trial and error in this process, but the search space for this is greatly reduced from the general problem. We typically only need to examine one or two κ values before a suitable tiling is found.

We performed targeted assembly for the diamond lattice in the HCRS system. We picked $\kappa = \sqrt{\frac{8}{3}}$ which corresponds to the next nearest to nearest neighbor ratio Figure 2.10(b). The zero temperature phase determination procedure was applied and the diamond lattice was found. The resulting phase sequence is detailed in Figure 2.11. The results from the HCRS system actually show only a random stacking of the diamond lattice as each stacking of the lattice is degenerate. If we use a generalized exponential model ($\gamma = 64$), a very close approximation to the HCRS system,

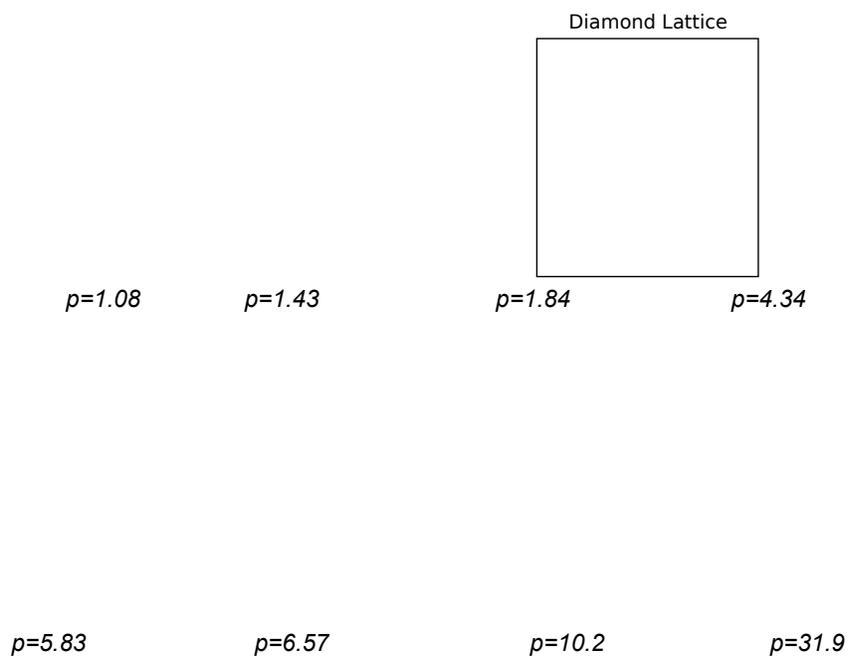


Figure 2.11: The zero temperature phase sequence for the 3D HCRS $\kappa = \sqrt{\frac{8}{3}}$ system. This phase sequence shows a variety of novel phases including a random stacked diamond lattice in the pressure range of $p = [1.84, 4.34]$. This random stacking can to be removed by using a continuous generalization of the HCRS pair potential. The colors are there to help distinguish the third dimension. All spheres of a given color are in the same plane.

we see that this splits the degeneracy of the random stackings of the lattice and the diamond lattice is recovered as the minimal enthalpy phase. We also have evidence that this phase is stable at small but finite temperatures due to the fact that this structure was self-assembled at finite temperature before the annealing froze out all degrees of freedom.

2.5 Conclusions

We have found that the HCSS family of isotropic repulsive pair potentials exhibits far more complex zero temperature phase behavior than has been previously recognized. This includes some significant differences from previous work, in particular, the work of Jagla, the most comparable mapping of the two dimensional zero temperature phase diagram as a function of κ and pressure. One of the most striking features of HCSS systems is the spontaneous appearance of entropy stabilized quasicrystal phases, liquid crystal phases, and lattices that were previously thought to require highly anisotropic bonding. The location of the interesting systems in parameter space is hinted at by the zero temperature phase diagram. Given that the properties of the HCSS family are largely unexplored, particularly in 3D, these observations only scratch the surface at the underlying richness of this class of systems.

Of particular importance are the opportunities that soft shoulder systems offer for the targeted design of novel self-assembling materials. As shown here and in the previous work, the soft shoulder clustering mechanism leads to fantastically diverse phase behavior.[35, 36, 8] Further, the Likos criterion and its generalization to HCSS systems provide theoretical tools for predicting under what thermodynamic conditions clustering will occur.[46, 27] We have developed tools that can be used for rudimentary targeted self-assembly in colloidal systems interacting via isotropic pairwise interactions with promising results including the forward prediction of a diamond lattice. While the end result is similar to work being pursued by Rechtsman et al., the technique is distinct and the focus is on utilization of clustering instabilities to stabilize specific structures.[60]

It is also of interest to contrast the soft shoulder clustering mechanism with an alternative route to clustering behavior arising from a competition between short-range attraction and long-

range repulsion, which has been widely studied in recent years. The latter clustering mechanism leads to a characteristic cluster size determined by competition between volume and surface energy contributions, whereas the characteristic size of clusters in soft shoulder systems can be identified with the spinodal instability length scale. Both characteristic lengths depend primarily on the form of the pair potential, however, and it is not clear that these two clustering mechanisms are physically distinct. In particular, we have found that adding short-range attraction to a purely repulsive HCSS potential leads to qualitatively similar phase behavior (including a variety of crystalline and liquid crystalline cluster phases), but with significantly enhanced temperature stability of cluster phases. This suggests that it may be possible to treat both clustering mechanisms within a unified theoretical framework.

Chapter 3

Advanced Monte Carlo sampling

As introduced in Section 2.1, a Monte Carlo sampling procedure generates states within the configuration space of a system in proportion to their Boltzmann factors. However, the configuration space of a system can be extremely large and, given an algorithm that poorly samples this space, the samples can give a distinctly unrepresentative picture of the equilibrium distribution. It is the job of the computational physicist to define the transition matrix, $k_{\mu \rightarrow \nu}$, that best explores configuration space within the constraints of ergodicity and detailed balance. In essence, this is the goal of algorithm development for these Monte Carlo calculations.

We have already made extensive use of one simple, and extremely common, transition matrix, the simple Metropolis site displacement move, and an additional set of transitions, the volume change move. When the random displacement is chosen from a spherically symmetric distribution, this transition function very clearly satisfies detailed balance because the displacement δr is just as likely as the displacement $-\delta r$. Further, since this displacement choice can generate any real world trajectory of the system, we have satisfied the ergodicity requirement. These methods are sufficient for many problems, but they become much less efficient in many systems. One place where these methods fail quite spectacularly is when dealing with tightly bound aggregates of particles. For these systems, motion of individual particles does not translate well into large scale motion of the aggregate.

3.1 Monte Carlo algorithms with collective mutation

It is often useful to compare Monte Carlo sampling with the sampling that can be achieved via molecular dynamics. One of the strongest benefits of Monte Carlo sampling over molecular dynamics is that the moves in a Monte Carlo Markov chain need not be similar to real world dynamics in any way. This opens up potential techniques that have no analog in a molecular dynamics simulation including non-local moves, particle swaps, and species mutation. The downside to this flexibility is that the very fact that these moves don't match physically realizable pathways in configuration space means that the Boltzmann factors associated with these moves can be very small and thus the acceptance rates of these moves can become extremely small.

What's more, this loss of acceptance rate can occur in sampling schemes where the moves are seemingly close to real world dynamics. Even algorithms like Metropolis sampling, where the individual moves of each particle seemingly mimic real world motion in that they make no non-local moves, can suffer from this effect. In systems that exhibit collective motion (e.g. systems with strongly bound particles, clustering/aggregation phenomena, or strong inter-particle forces), the Metropolis algorithm sampling scheme becomes inappropriate. This is precisely due to the fact that the moves proposed in the Metropolis method are not good analogues to the motion seen in a corresponding real world system. If we consider systems with aggregation, we see that there are at least two modes of relaxation, reorganization of individual particles within aggregates, and reorganization of the aggregates within the system. The dynamics and time scales of these two modes can be vastly different. Metropolis moves are based around the motion of a single particle in the system, and thus are only able to effectively sample configurations involving reorganization of particles within the aggregate. The larger scale reorganization of aggregates as a whole is a slow process.

This introduces issues with ergodicity. While there exists a finite path from one aggregation configuration to another, the probability of one of these paths being taken in the simulation is vanishingly small for the finite number of Markov transitions taken in the computation.

To make this concrete, consider a system of hard spheres with a attractive well interaction of depth ϵ_w and width σ_w where the particles are tightly bound into aggregates. In such a system, the maximum displacement that you expect a single sphere to be able to move independent of the rest of the aggregate is on the order of the well width in either Metropolis sampling or physical dynamics of the spheres. Any displacement larger than this will incur large potential energy costs. However, When considering the aggregates as a whole, it is clear that σ_w has little to do with the largest displacement of the aggregate with reasonable acceptance. In real world dynamics, each aggregate moves collectively, and a gas of these aggregates can have quite large displacements in comparison to σ_w . When using single site displacements, we are unable to easily realize this kind of collective motion as a series of Monte Carlo moves. Moving a aggregate by moving individual spheres requires you to essentially split the aggregate and the intermediate states will have an increase in system energy proportional to the change in surface area of the aggregate which scales with $R^{d-1} \sim (n)^{\frac{d-1}{d}} \Rightarrow P \sim \exp \left[-\beta \epsilon_w \times n^{\frac{d-1}{d}} \right]$. At its very core, this limitation is due to the fact that our moves in Metropolis sampling are ignorant of the system potential and hence the tendency for collective motion. Since standard Metropolis doesn't incorporate that knowledge of the system into its procedure, it cannot probe the relaxation of aggregates.

This effect alone is already enough to make Metropolis-style sampling unsuitable for systems that exhibit aggregation behavior. There is another issue, however, that we must deal with if we wish to probe systems with significant collective motion. The Metropolis move is simply a random translation. This means is that each transition involves a proposed random move that is independent of the last move. However, in order to move an aggregate with a series of these moves, at the very least each particle must move in roughly the same direction. Dividing the possible displacements into m different possible directions, then an upper bound on the probability of an aggregate of n particles moving together via the Metropolis algorithm is m^{-n} . This applies even if our acceptance probabilities are high because our displacements are of order σ_w , the chances of a group of particles moving in the same direction drops off exponentially with the number of particles.

While moves like site displacements allow the Markov chain to sample some of the important internal configurations for the aggregate,¹ they are wholly unsuited for systems that require collective motion for structural relaxation. Any scheme designed to allow for collective motion must provide a method of dealing with both of these issues: moving the particles in a collective fashion and the ability to resolve or avoid high energy intermediate states.

It should be noted that this does not mean that these individual particle moves are unimportant in these systems. In many ways these small site displacements are a sort of lubrication that allows other algorithms to work. This is clear when you consider the volume change move which scales all distances in the system. If we are only moving aggregates, then the distances between the particles within the aggregates will completely determine the maximum density of the system.

3.2 Adding collective moves to the Metropolis algorithm

To a great extent we can correct for many of these deficiencies by including special new moves that are explicitly designed to move groups of particles. One simple method for this is to choose a seed particle and group particles in the system according to an arbitrary linking function that only depends on the initial state, μ . This group is then moved collectively.[42] An example definition for the linking function, LINK_P, is given, but this should be considered a free parameter for the system to optimize. This easily fixes the two problems with Metropolis sampling. All particles are moved at the same time and thus there are no intermediate states to be worried about. Algorithm 5 demonstrates how this process is implemented.

There are several shortcomings of this technique. First, there is an issue with detailed balance that must be addressed. When dealing with unbiased moves, the defining characteristic of detailed balance is that forward moves are just as likely for proposal as the reverse of that move. Since the linking of two particles is determined by the current state of the system, we must ensure that we will never accept a move that will modify the state in such a way that the cluster formed in μ is different from the cluster formed in ν , given the same seed particle. This is remedied by

¹ But not all, as we will see in Section 3.3.3.

Algorithm 5 The proximity cluster move. This function grows a cluster around a selected seed particle and moves it via a standard site dislocation.

```

1: function PROXIMITY_BUILD_CLUSTER( $\{\vec{r}\}$ ,  $i_{\text{seed}}$ )
2:    $\mathcal{C} \leftarrow [i_{\text{seed}}]$      $\triangleright$  Include the seed in the cluster.     $\triangleright$  Select a small random displacement.
3:   for all  $\vec{r}_j \in \{\vec{r}\}$  do
4:     if LINK_P( $\vec{r}_i, \vec{r}_j$ ) then
5:        $\mathcal{C} \leftarrow \mathcal{C} + j$      $\triangleright$  Include  $j$  in cluster.
6:     end if
7:   end for
8:   return  $\mathcal{C}$ 
9: end function
10:
11: function PROXIMITY_CLUSTER_MOVE( $\{\vec{r}\}^{(\mu)}$ )
12:    $i_{\text{seed}} \leftarrow \text{random}(\text{size}(\{\vec{r}\}^{(\mu)}))$      $\triangleright$  Randomly select seed particle.
13:    $\mathcal{C}^{(\mu)} \leftarrow \text{PROXIMITY\_BUILD\_CLUSTER}(\{\vec{r}\}^{(\mu)}, i_{\text{seed}})$ 
14:    $\Delta\vec{r} \leftarrow \tilde{\delta} \times \text{random\_vector}(1.0)$ 
15:    $\{\vec{r}\}^{(\nu)} \leftarrow \{\vec{r}\}^{(\mu)}_{\vec{r}_i \rightarrow \vec{r}_i + \Delta\vec{r} \forall i \in \mathcal{C}}$      $\triangleright$  Displace each particle in  $\mathcal{C}$  by  $\Delta r$ .
16:   if ( $\text{random}(1.0) < e^{-\beta(E^{(\nu)} - E^{(\mu)})}$ ) then
17:      $\mathcal{C}^{(\nu)} \leftarrow \text{PROXIMITY\_BUILD\_CLUSTER}(\{\vec{r}\}^{(\nu)}, i_{\text{seed}})$      $\triangleright$  Compute cluster in  $\nu$  state.
18:     if  $\mathcal{C}^{(\mu)} = \mathcal{C}^{(\nu)}$  then     $\triangleright$  Check if  $\mu$  cluster equals  $\nu$  cluster.
19:       return ( $\{\vec{r}\}^{(\nu)}, (E^{(\nu)} - E^{(\mu)})$ )     $\triangleright$  Accept, return the new state and energy change.
20:     end if
21:   end if
22: end function
23:
24: function LINK_P( $\vec{r}_i, \vec{r}_j$ )
25:   return  $u(\vec{r}_i, \vec{r}_j) < -\tilde{\epsilon}_{\text{link}}$      $\triangleright$  Link if the binding energy is below  $-\tilde{\epsilon}_{\text{link}}$ .
26: end function

```

line 15, which rejects moves if it would result taking two clusters in μ and merging them into one cluster in ν . The inverse of this problem also exists. Since the linking in LINK_P is absolute, it is impossible for one of these moves to split a cluster in μ into two clusters in ν . Second, this does a relatively poor job of relaxing internal structure. This is due to the fact that particles that are grouped are stuck in their configuration as far as the cluster moves are concerned. This means that the cluster algorithm cannot be used to relax internal structure of that cluster. It should be noted that it is highly beneficial to use cluster algorithms to relax internal structure of clusters as the same limitations of Metropolis with regards to collective motion still apply to relaxations within the cluster. This can be fixed in certain cases when there is a parameter in the linking function that can be tuned to reduce the tendency for links to form. In Algorithm 5, this is done by tuning a maximum energy for links to form. For any potential with the exception of a square well, increasing ϵ_{link} has the effect of reducing linking during the cluster building procedure. The work by Kuriabova, Betterton, and Glaser, used a threshold based on the distance between point binding sites on their particles.[42] Despite these drawbacks, this technique proves to be a useful option for cluster moves. In the next section we will see a more general clustering technique that allows for probabilistic linking according to the system potential.

3.3 Swendsen-Wang style collective moves

There is another class of cluster algorithms that originate from the work of Swendsen and Wang, and later generalized to into a single cluster method by Wolff.[67, 78] These methods were developed as a collective update for the Potts model but, as we will see, have been generalized to many different systems.[20, 33, 47, 48]

While our initial discussions of the Metropolis algorithm focused on fluid systems, the same algorithm can be applied to spin systems. To apply the Metropolis algorithm to spin systems, the moves are proposed by picking a spin by random and choosing a random, new spin value for it. The acceptance criterion is the same as in Algorithm 1. This suffers from the same limitations as described in Section 3.1. When the spin-spin correlation length becomes large, the Metropolis

method becomes wholly unsuitable. The work by Swendsen and Wang and later Wolff shows how a Monte Carlo calculation can change multiple spins in concert to overcome this critical slowdown of the algorithm.[67, 78, 72]

Algorithm 6 The Wolff algorithm move for an m -state Potts model. The procedure WOLFF_BUILD_CLUSTER constructs a cluster by linking pairs spins (represented by adding them to the queue, q) according to the change in energy due to changing the spin's state to s . The loop over q takes into consideration the fact that new elements may be added to the end of the queue during the loops execution. This algorithm is specifically structured to demonstrate the similarity between the Wolff algorithm and the GenGCA and VMMC algorithms.

```

1: procedure WOLFF_MOVE( $\vec{\sigma}$ ,  $m$ )
2:   seed  $\leftarrow$  random(size( $\sigma$ ))                                 $\triangleright$  Select random seed spin.
3:    $s \leftarrow$  random( $m$ ) such that  $s \neq \sigma$  [seed]           $\triangleright$  Select new spin state.
4:   ( $\mathcal{C}$ ,  $\Delta\vec{U}$ )  $\leftarrow$  WOLFF_BUILD_CLUSTER( $\vec{\sigma}$ , seed,  $s$ )   $\triangleright$  Build cluster.
5:   for all  $i \in \mathcal{C}$  do
6:      $\sigma[i] \leftarrow s$                                         $\triangleright$  Update spins.
7:   end for
8:   for all  $i$  such that  $\Delta\vec{U}[i] \neq 0$  and  $i \notin \mathcal{C}$  do
9:      $\Delta E \leftarrow \Delta E + \Delta\vec{U}[i]$                      $\triangleright$  Energy only changes at the edge of the cluster.
10:  end for
11: end procedure
12:
13: function WOLFF_BUILD_CLUSTER( $\vec{\sigma}$ , seed,  $s$ )
14:    $q \leftarrow$  [seed]                                            $\triangleright$  Insert the seed spin into the cluster.
15:    $\Delta\vec{U} \leftarrow \vec{0}$                                         $\triangleright$  Zero out energetic contributions.
16:   for all  $i \in q$  do                                            $\triangleright$  Loop over the spins in the cluster.
17:     for all  $j \in \{1, \dots, N\}/q$  do                           $\triangleright$  Loop over spins that aren't in the cluster.
18:        $\Delta u \leftarrow (\epsilon_{ij} - \epsilon_{ij})_{\sigma_i \rightarrow s}$    $\triangleright$  Energy change due to switching  $\sigma_i$  to  $s$  in  $\epsilon_{ij}$ .
19:       if random(1.0)  $> 1 - e^{-\beta\Delta u}$  then
20:          $q \leftarrow$  append( $q$ , [ $j$ ])                             $\triangleright$  Link: add  $j$  to the cluster.
21:       else
22:          $\Delta\vec{U}[j] \leftarrow \Delta\vec{U}[j] + \Delta u$              $\triangleright$  Do not link. Tally the energetic contribution.
23:       end if
24:     end for
25:   end for
26:   return ( $q, \Delta\vec{U}$ )                                          $\triangleright$  Return the cluster and the energetic contributions.
27: end function

```

In the Wolff algorithm, listed in Algorithm 6, clusters are built by growing from a seed spin. Each of the seed spin's neighboring spins is considered for addition to the cluster by *linking* the two spins. When a link occurs, the procedure adds j to the cluster ensuring that the neighbors of j will be considered for addition to the cluster as well. Spins may be considered for inclusion into

the cluster multiple times, but once it has been included, there is no need to consider further link attempts with it.[72]

Spin i is linked with spin j when considering a transition from state μ to state ν according to the probability $p_{ij}(\mu \rightarrow \nu)$ in (3.1). From a practical point of view, at the time of building the cluster we only know the part of the final state ν that involves the spins that have already been linked into the cluster. Fortunately, while the linking probability, $p_{ij}(\mu \rightarrow \nu)$, depends on the values of the spins i and j in μ , it only depends on the value of i in ν and i is guaranteed to be part of the cluster.

$$p_{ij}(\mu \rightarrow \nu) \equiv \max(0, 1 - e^{-\beta(\epsilon_{i'j}^{(\mu \rightarrow \nu)} - \epsilon_{ij}^{(\mu \rightarrow \nu)})}) \quad (3.1)$$

Niedermayer demonstrated how the principles of Swendsen-Wang style Monte Carlo could be applied using any global symmetry of the system and introduced a generalization very similar to the cluster cleaving method of Whitlam and Geissler, see Section 3.3.3.[54] Work by Heringa and Blöte on the geometric cluster move algorithm (GCA), further formalized the theory of these types of moves and applied technique to hard core lattice gases.[33] Pioneering work by Dress and Krauth on the pivot moves showed that the similar techniques can be used in fluid off-lattice hard sphere systems.[20] Much as the Wolff algorithm is a single cluster version of the Swendsen-Wang algorithm, the Liu and Luijten GenGCA algorithm is a single cluster version of the Dress and Krauth algorithm, though a different transformation is used. In addition, the algorithm was extended to systems with arbitrary pair interactions by applying the same linking method as the Wolff algorithm.[47, 48] The GenGCA, shown in Algorithm 7, uses an inversion transformation of the system relative to a randomly selected point and adds subsequent particles to the cluster from the environment via the same linking probability used in the Wolff algorithm. Like the Wolff algorithm, the GenGCA is rejection-free.

While fluid simulation and spin systems are quite different, comparing Algorithms 6 and 7 it is clear that these methods have a lot in common. In fact, with a few minor modifications, the two could be merged into a single generalized form. We note that the cluster building function

Algorithm 7 The GenGCA algorithm from Liu and Luijten, but generalized to arbitrary self-inverse operation. This takes an operation, \mathcal{O} , and performs a rejection-free cluster move. As in Algorithm 6, new particles can be added to the cluster queue during the iteration over the cluster.

```

1: procedure GENGCA( $\{\vec{x}\}, M$ )
2:    $\mathcal{O} \leftarrow \{\text{inversion}, \vec{x}_p : \text{random\_vector}(M)\}$ 
3:    $\text{seed} \leftarrow \text{random}(\text{size}(\{\vec{x}\}))$ 
4:    $(\mathcal{C}, \Delta\vec{U}) \leftarrow \text{SELF\_INVERSE\_BUILD\_CLUSTER}(\{\vec{x}\}, \text{seed}, \mathcal{O})$ 
5:   for all  $i \in \mathcal{C}$  do
6:      $\{\vec{x}\}[i] \leftarrow \mathcal{O}(\{\vec{x}\}[i])$ 
7:   end for
8:   for all  $i$  such that  $\Delta\vec{U}[i] \neq 0$  and  $i \notin \mathcal{C}$  do
9:      $\Delta E \leftarrow \Delta E + \Delta\vec{U}[i]$   $\triangleright$  Energy only changes between particles not in the cluster.
10:  end for
11: end procedure
12:
13: function SELF_INVERSE_BUILD_CLUSTER( $\{\vec{x}\}, \text{seed}, \mathcal{O}$ )
14:    $q \leftarrow [\text{seed}]$   $\triangleright$  Insert the seed particle into the cluster.
15:    $\Delta\vec{U} \leftarrow \vec{0}$   $\triangleright$  Zero out energetic contributions.
16:   for all  $i \in q$  do  $\triangleright$  Loop over particles in the cluster.
17:     for all  $j \in \{1..N\}$  where  $j \notin q$  do  $\triangleright$  Loop over particles that aren't in the cluster.
18:        $\Delta u \leftarrow (\epsilon_{ij} - \epsilon_{ji})$   $\triangleright$  Energy change due to  $x_i \rightarrow \mathcal{O}x_i$  in  $\epsilon_{ij}$ .
19:       if  $\text{random}(1.0) > 1 - e^{-\beta\Delta u}$  then
20:          $q \leftarrow \text{append}(q, [j])$   $\triangleright$  Link: add  $j$  to the cluster.
21:       else
22:          $\Delta\vec{U}[j] \leftarrow \Delta\vec{U}[j] + \Delta u$   $\triangleright$  Do not link. Tally the energetic contribution.
23:       end if
24:     end for
25:   end for
26:   return  $(q, \Delta U)$   $\triangleright$  Return the cluster and the energetic contributions.
27: end function

```

in the GenGCA takes a symmetry operator as an argument. The cluster building procedure in Algorithm 7 is truly agnostic toward symmetry operations. This operation could just as easily been any global symmetry operation. If a spin flip operation is used, the algorithm becomes the Wolff algorithm as seen in Algorithm 8. Similarly, we can implement a single cluster version of the Dress and Krauth algorithm by passing a rotation operator as shown in the PIVOT_CLUSTER procedure in Algorithm 8. While these algorithms apply to very different systems, the only real difference between the three is the specification of the symmetry operation.

Algorithm 8 An implementation of the Wolff algorithm and the Dress-Krauth pivot cluster algorithm using the self-inverse cluster building procedure from the GenGCA algorithm.

```

1: procedure WOLFF( $\{\sigma\}$ )
2:    $\mathcal{O} \leftarrow \{\sigma_{\text{old}} \leftrightarrow \sigma_{\text{new}}\}$ 
3:   seed  $\leftarrow$  random(size( $\{\sigma\}$ ))
4:    $(\mathcal{C}, \Delta\vec{U}) \leftarrow$  SELF_INVERSE_BUILD_CLUSTER( $\{\sigma\}$ , seed,  $\mathcal{O}$ )
5:   for all  $i \in \mathcal{C}$  do
6:      $\{\sigma\}[i] \leftarrow \mathcal{O}(\{\sigma\}[i])$ 
7:   end for
8:   for all  $i$  such that  $\Delta\vec{U}[i] \neq 0$  and  $i \notin \mathcal{C}$  do
9:      $\Delta E \leftarrow \Delta E + \Delta\vec{U}[i]$        $\triangleright$  Energy only changes between particles not in the cluster.
10:  end for
11: end procedure
12:
13: procedure PIVOT_CLUSTER( $\{\vec{x}\}$ )
14:    $\mathcal{O} \leftarrow \{\pi - \text{rotation}, \vec{x}_p : \text{random\_vector}(M)\}$ 
15:   seed  $\leftarrow$  random(size( $\{\vec{x}\}$ ))
16:    $(\mathcal{C}, \Delta\vec{U}) \leftarrow$  SELF_INVERSE_BUILD_CLUSTER( $\{\vec{x}\}$ , seed,  $\mathcal{O}$ )
17:   for all  $i \in \mathcal{C}$  do
18:      $\{\vec{x}\}[i] \leftarrow \mathcal{O}(\{\vec{x}\}[i])$ 
19:   end for
20:   for all  $i$  such that  $\Delta\vec{U}[i] \neq 0$  and  $i \notin \mathcal{C}$  do
21:      $\Delta E \leftarrow \Delta E + \Delta\vec{U}[i]$        $\triangleright$  Energy only changes between particles not in the cluster.
22:   end for
23: end procedure

```

3.3.1 Biased sampling

The Swendsen-Wang style techniques are a special case of biased sampling. In biased sampling the move proposal rate, $\alpha_{\mu \rightarrow \nu}$ in (2.5), is dependent on the initial and final state of the system. The detailed balance restriction in (2.3) must still be maintained, so the acceptance rate must be

modified to correct for this. The goal of biased sampling is to shift some of the acceptance criterion to the move proposal in the hopes of biasing the sampling toward a particular type of move. This can have a positive effect as including knowledge of the system's state, both initial and final, in the move proposal can greatly improve the sampling, as the nature of the moves can fundamentally change. In this section we demonstrate that this method naturally introduces collective motion into the Monte Carlo calculation.

It is useful to apply the *super-detailed balance* property. Super-detailed balance is a further restriction that states that detailed balance must be satisfied for any transition path between μ and ν . When considering generic Swendsen-Wang style algorithms, super-detailed balance requires that the detailed balance must be maintained for each and every possible way of picking a cluster. This includes which links are formed and which have failed during in the cluster building process, but also includes the seed particle that is chosen and even the order in which the links are formed. In this cluster building procedure, there are many paths of execution that result in the same cluster. This is necessary as different execution paths can result in different biases and thus different corrections in the acceptance probability. One path may start with i , then link $i \rightarrow j$, then link $j \rightarrow k$, which requires a different bias correction than starting with i , then linking $i \rightarrow k$, then linking $k \rightarrow j$, though each result in the same three particle cluster. Each of these possible execution pathways that result in cluster \mathcal{C} is considered a realization, \mathcal{R} , of cluster \mathcal{C} . In the language of the master equation the super-detailed balance condition can be written as (3.2). Any sampling chain that satisfies the super-detailed balance condition also satisfies the detailed balance condition which can be restored by summing over all possible realization of the cluster. In general, we have quite a bit of flexibility regarding how we distribute the transition rate into the move proposal and acceptance criterion. We are able to pick any split of the two, moving from the standard unbiased limit, where Metropolis style detailed balance applies, to the limit of rejection-free, or nearly rejection-free,

methods.

$$P(\mu) \times \alpha_{\mu \rightarrow \nu}(\mathcal{R}) \times \text{acc}_{\mu \rightarrow \nu}(\mathcal{R}) = P(\nu) \times \alpha_{\nu \rightarrow \mu}(\mathcal{R}) \times \text{acc}_{\nu \rightarrow \mu}(\mathcal{R}) \quad (3.2)$$

$$\alpha(\mu, \nu) = (\text{probability of selecting } \mu \rightarrow \nu) \quad (3.3)$$

$$\text{acc}(\mu, \nu) = \frac{(\text{probability of selecting } \nu \rightarrow \mu)}{(\text{probability of selecting } \mu \rightarrow \nu)} \times e^{-\beta(E(\nu) - E(\mu))} \quad (3.4)$$

The probability of the forward move getting proposed is given by (3.5) where $p_{ij}(\mu \rightarrow \nu)$, the same function as in (3.1), is the probability of a link being formed and $q_{ij}(\mu \rightarrow \nu) \equiv 1 - p_{ij}(\mu \rightarrow \nu)$ is the probability of a link failing to form. Note that by definition, i and j in the product over successful links are necessarily both part of the cluster. In the product over failed links, some factors are link failures between the cluster and the environment, $i \in \mathcal{C}$ and $j \notin \mathcal{C}$, whereas other failures are within the cluster, $i, j \in \mathcal{C}$. Both products depend on the realization of the cluster and only run over the actual successful or failed links.

$$\alpha_{\mu \rightarrow \nu}(\mathcal{R}) = \prod_{i,j \in \mathcal{C}}^{\mathcal{R}} p_{ij}(\mu \rightarrow \nu) \prod_{i,j}^{\mathcal{R}} q_{ij}(\mu \rightarrow \nu) \quad (3.5)$$

$$\text{acc}(\mu \rightarrow \nu) = \min \left(1, e^{-\beta(E(\nu) - E(\mu))} \prod_{i,j}^{\mathcal{R}} \frac{q_{ij}(\nu \rightarrow \mu)}{q_{ij}(\mu \rightarrow \nu)} \prod_{i,j \in \mathcal{C}}^{\mathcal{R}} \frac{p_{ij}(\nu \rightarrow \mu)}{p_{ij}(\mu \rightarrow \nu)} \right) \quad (3.6)$$

$$\begin{aligned} \text{acc}(\mu \rightarrow \nu) = \min \left(1, e^{-\beta(E(\nu) - E(\mu))} \prod_{i \in \mathcal{C}, j \notin \mathcal{C}} \frac{q_{ij}(\nu \rightarrow \mu)}{q_{ij}(\mu \rightarrow \nu)} \right. \\ \left. \times \prod_{i,j \in \mathcal{C}}^{\mathcal{R}} \frac{q_{ij}(\nu \rightarrow \mu)}{q_{ij}(\mu \rightarrow \nu)} \prod_{i,j \in \mathcal{C}}^{\mathcal{R}} \frac{p_{ij}(\nu \rightarrow \mu)}{p_{ij}(\mu \rightarrow \nu)} \right) \end{aligned} \quad (3.7)$$

$$\text{acc}(\mu \rightarrow \nu) = \min \left(1, \prod_{i,j \in \mathcal{C}}^{\mathcal{R}} \frac{q_{ij}(\nu \rightarrow \mu)}{q_{ij}(\mu \rightarrow \nu)} \prod_{i,j \in \mathcal{C}}^{\mathcal{R}} \frac{p_{ij}(\nu \rightarrow \mu)}{p_{ij}(\mu \rightarrow \nu)} \right) \quad (3.8)$$

With super-detailed balance, we can say that the same links and link failures are present in the forward and reverse moves, simplifying the derivation of the acceptance criterion.[23, 77] We find that the appropriate acceptance probability is (3.6). This probability contains three factors, the standard ratio of Boltzmann factors, the bias correction for failed links, and the bias correction for successful links. We can simplify the expression by separating the failed links that are within the cluster from those that are between the cluster and the environment, giving (3.7). The failed links

between the cluster and environment exactly cancel the Boltzmann factors, leaving the acceptance probability as (3.8).

We can see this cancellation by considering that the linking probability is pairwise. Since $\nu \rightarrow \mu$ is the inverse move of $\mu \rightarrow \nu$, the values of the pairwise potentials satisfy the relations $\epsilon_{i'j}^{(\mu \rightarrow \nu)} = \epsilon_{ij}^{(\nu \rightarrow \mu)}$ and $\epsilon_{ij}^{(\mu \rightarrow \nu)} = \epsilon_{i'j}^{(\nu \rightarrow \mu)}$ as shown in Figure 3.1. In addition, since $\epsilon_{i'j} - \epsilon_{ij}$ must be either positive or negative (or zero, but in that case the entire factor is simply one), each factor in the product has either one in the numerator and $e^{-\beta(\epsilon_{i'j} - \epsilon_{ij})}$ in the denominator, or one in the denominator and $e^{+\beta(\epsilon_{i'j} - \epsilon_{ij})}$ in the numerator. This allows us to finally write the product over the failed links between the cluster and the environment as (3.12), which is the inverse of the ratio of Boltzmann factors that shows up (3.6).

$$\prod_{i \in \mathcal{C}, j \notin \mathcal{C}} \frac{q_{ij}(\nu \rightarrow \mu)}{q_{ij}(\mu \rightarrow \nu)} = \prod_{i \in \mathcal{C}, j \notin \mathcal{C}} \frac{1 - p_{ij}(\nu \rightarrow \mu)}{1 - p_{ij}(\mu \rightarrow \nu)} \quad (3.9)$$

$$\prod_{i \in \mathcal{C}, j \notin \mathcal{C}} \frac{1 - \max(0, 1 - e^{-\beta(\epsilon_{i'j}^{(\nu \rightarrow \mu)} - \epsilon_{ij}^{(\nu \rightarrow \mu)})})}{1 - \max(0, 1 - e^{-\beta(\epsilon_{i'j}^{(\mu \rightarrow \nu)} - \epsilon_{ij}^{(\mu \rightarrow \nu)})})} = \prod_{i \in \mathcal{C}, j \notin \mathcal{C}} \frac{\min(1, e^{\beta(\epsilon_{i'j}^{(\nu \rightarrow \mu)} - \epsilon_{ij}^{(\nu \rightarrow \mu)})})}{\min(1, e^{-\beta(\epsilon_{i'j}^{(\mu \rightarrow \nu)} - \epsilon_{ij}^{(\mu \rightarrow \nu)})})} \quad (3.10)$$

$$= \prod_{i \in \mathcal{C}, j \notin \mathcal{C}} \frac{\min(1, e^{+\beta(\epsilon_{i'j}^{(\mu \rightarrow \nu)} - \epsilon_{ij}^{(\mu \rightarrow \nu)})})}{\min(1, e^{-\beta(\epsilon_{i'j}^{(\mu \rightarrow \nu)} - \epsilon_{ij}^{(\mu \rightarrow \nu)})})} \quad (3.11)$$

therefore...

$$\prod_{i \in \mathcal{C}, j \notin \mathcal{C}} \frac{1}{e^{-\beta(\epsilon_{i'j}^{(\mu \rightarrow \nu)} - \epsilon_{ij}^{(\mu \rightarrow \nu)})}} \times \prod_{i \in \mathcal{C}, j \notin \mathcal{C}} \frac{e^{+\beta(\epsilon_{i'j}^{(\mu \rightarrow \nu)} - \epsilon_{ij}^{(\mu \rightarrow \nu)})}}{1} \quad (3.12)$$

$$= \prod_{i \in \mathcal{C}, j \notin \mathcal{C}} e^{+\beta(\epsilon_{i'j}^{(\mu \rightarrow \nu)} - \epsilon_{ij}^{(\mu \rightarrow \nu)})} = e^{+\beta(E(\nu) - E(\mu))} \quad (3.13)$$

This simplifies the acceptance probability to that of (3.8). This result is quite general for any cluster with the linking probability in (3.1). For some of the Swendsen-Wang style cluster moves, such as the ones we have seen so far, this simplifies considerably. When the linking method symmetry operation is self-inverse, all of the factors involving particles within the cluster, both the successful links and the portion of failed links that don't involve the environment, all reduce to

one. This can be seen by the observation that in the case of self-inverse operations, $\epsilon_{i'j}^{(\mu \rightarrow \nu)} = \epsilon_{i'j}^{(\nu \rightarrow \mu)}$ and $\epsilon_{ij}^{(\mu \rightarrow \nu)} = \epsilon_{ij}^{(\nu \rightarrow \mu)}$ when $i, j \in \mathcal{C}$ as shown in Figure 3.1. The end result is that the acceptance probability for these self-inverse moves is unity.

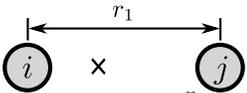
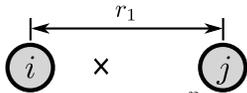
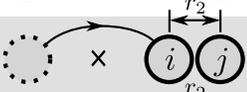
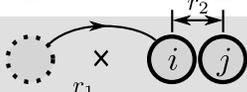
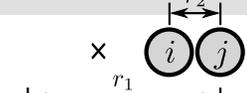
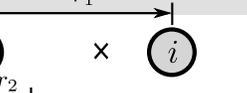
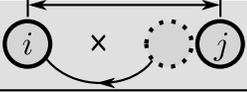
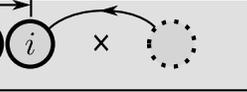
a)	$i \in \mathcal{C} \quad j \notin \mathcal{C}$	Equal to:	$i \in \mathcal{C} \quad j \in \mathcal{C}$	Equal to:	
$\epsilon_{ij}^{(\mu \rightarrow \nu)}$	$\pm_i \pm_j$	$\epsilon_{i'j}^{(\nu \rightarrow \mu)}$	$\pm_i \pm_j$	$\epsilon_{ij}^{(\mu \rightarrow \nu)}$	Wolff
$\epsilon_{i'j}^{(\mu \rightarrow \nu)}$	$\mp_i \pm_j$	$\epsilon_{ij}^{(\nu \rightarrow \mu)}$	$\mp_i \pm_j$	$\epsilon_{i'j}^{(\mu \rightarrow \nu)}$	
$\epsilon_{ij}^{(\nu \rightarrow \mu)}$	$\mp_i \pm_j$	$\epsilon_{i'j}^{(\mu \rightarrow \nu)}$	$\mp_i \mp_j$	$\epsilon_{ij}^{(\nu \rightarrow \mu)}$	
$\epsilon_{i'j}^{(\nu \rightarrow \mu)}$	$\pm_i \pm_j$	$\epsilon_{ij}^{(\mu \rightarrow \nu)}$	$\pm_i \mp_j$	$\epsilon_{i'j}^{(\nu \rightarrow \mu)}$	
b)					GenGCA
$\epsilon_{ij}^{(\mu \rightarrow \nu)}$		$\epsilon_{i'j}^{(\nu \rightarrow \mu)}$		$\epsilon_{ij}^{(\nu \rightarrow \mu)}$	
$\epsilon_{i'j}^{(\mu \rightarrow \nu)}$		$\epsilon_{ij}^{(\nu \rightarrow \mu)}$		$\epsilon_{i'j}^{(\mu \rightarrow \nu)}$	
$\epsilon_{ij}^{(\nu \rightarrow \mu)}$		$\epsilon_{i'j}^{(\mu \rightarrow \nu)}$		$\epsilon_{ij}^{(\mu \rightarrow \nu)}$	
$\epsilon_{i'j}^{(\nu \rightarrow \mu)}$		$\epsilon_{ij}^{(\mu \rightarrow \nu)}$		$\epsilon_{i'j}^{(\nu \rightarrow \mu)}$	

Figure 3.1: An explanation of the pairwise energy terms in the Wolff and GenGCA algorithms. The presentation of these two algorithms is just done for illustrative purposes. These same properties apply to any system that matches the criteria for these Swendsen-Wang style algorithm's restriction of a self-inverse operator that is a symmetry of the Hamiltonian. Note that when considering the special class of failed links where $j \notin \mathcal{C}$, $\epsilon_{ij}^{(\mu \rightarrow \nu)} = \epsilon_{i'j}^{(\nu \rightarrow \mu)}$ and $\epsilon_{i'j}^{(\mu \rightarrow \nu)} = \epsilon_{ij}^{(\nu \rightarrow \mu)}$. This is distinct from the behavior of linked particles where $\epsilon_{ij}^{(\mu \rightarrow \nu)} = \epsilon_{ij}^{(\nu \rightarrow \mu)}$ and $\epsilon_{i'j}^{(\mu \rightarrow \nu)} = \epsilon_{i'j}^{(\nu \rightarrow \mu)}$.

These methods (the Swendsen-Wang algorithm, the Wolff algorithm, the geometric cluster algorithm (GCA), the Dress and Krauth pivot algorithm, and the GenGCA) are all intimately related and operate on the same principle. Indeed, they should be considered the same algorithm. Each exploits an self-inverse symmetry operation, \mathcal{O} of a pairwise system Hamiltonian, $H(\{\Delta\vec{r}\})$ such that $H(\{\Delta\vec{r}\}) = H(\mathcal{O}\{\Delta\vec{r}\})$, to design system transition processes that can do more than

simply fail when there is a poor mutation of the system state. The self-inverse property of the operation states that $\mathcal{O}^2 = I$, where I is the identity operator, and is necessary to prove the relationships between $\epsilon_{ij}^{\mu \rightarrow \nu}$, $\epsilon_{i'j}^{\mu \rightarrow \nu}$, $\epsilon_{ij}^{\nu \rightarrow \mu}$, and $\epsilon_{i'j}^{\nu \rightarrow \mu}$ from Figure 3.1. In the case of Swendsen-Wang and the Wolff algorithms, this symmetry is the inversion of the spins, or in the case of the Potts model, selection of a different spin value. In the case of the Dress and Krauth pivot algorithm, the symmetry is a rotation by π in the 2D systems. In the GCA, the symmetry is left undefined with the only condition that it is a self-inverse operation. In the GenGCA the symmetry is inversion of the system around a point. These are all different symmetries but the underlying idea is the same: move parts of the system particle by particle according to the symmetry operation until the system has reached an acceptable configuration. The ability to save moves from almost certain rejection opens the possibility for arbitrarily large changes to the system configuration without the precipitous drop in acceptance probability that the Metropolis algorithm would incur. Since \mathcal{O} is a symmetry of the Hamiltonian, these methods are all naturally rejection-free. There is guaranteed to be an acceptable configuration in the progress of forming the cluster. In the worst case scenario, this configuration is simply the original configuration with the transformation applied to every particle. The algorithm is succeeding if an intermediate configuration is found where the transformation is only partially applied.

Due to the linking probability's special form, the resulting bias correction exactly cancels the appropriate Metropolis acceptance probability associated with the move. This means that when the algorithm exits, no matter what particles it moved, the resultant configuration is guaranteed to have been generated in correct proportion to the ensemble. The rejection-free nature of this method is not without a price. In many cases the cluster will grow to encompass the entire system. In this situation the symmetry operation is applied globally to the system. While the configuration may technically change, no relaxation of the structure is being performed, nor is the configuration anything that can't be trivially predicted about the system without the aid of Monte Carlo sampling. For all intents and purposes, this should be considered a rejection of the move attempt. This mode of failure is actually much worse than a rejection in the Metropolis algorithm. While rejections

in the Metropolis algorithm also don't sample new configurations, they are exceptionally cheap in comparison.

3.3.2 Virtual move Monte Carlo (VMMC)

One might wonder if there is some way to extend this class of algorithms to operations that are not self-inverse, such as arbitrary translations. Doing so would make the movements within the Markov chain more closely mimic movements in real world systems. This is an attractive feature for some Monte Carlo calculations as methods like this can be used to measure dynamics in the system, something that is not typically possible in standard Monte Carlo.[77] Translations meet part of the requirements on the transformation. In systems with only pair interactions, translating every particle in the system by a distance $\Delta\vec{r}$ leaves the energy invariant. However, translations are not self-inverse, a property that greatly simplified the acceptance probability in the Swendsen-Wang style cluster moves. While it is simple to derive the inverse move of a translation, the bias correction no longer exactly cancels the external Boltzmann factors in the acceptance probability, leaving us with (3.8). Figure 3.2 shows the relations between the pair potential parameters, however the fact that $\epsilon_{ij}^{(\mu\rightarrow\nu)} \neq \epsilon_{ij}^{(\nu\rightarrow\mu)}$, i.e. the fact that translations are not self-inverse, means that no further simplification is possible. There is nothing stopping us from using this algorithm as it is, however. All we need do is retain the information of the pair potential calculations and evaluate the products for failed and successful links in an extra acceptance stage. This algorithm is a slight modification of the *virtual move Monte Carlo* (VMMC) algorithm of Whitelam and Geissler.[77, 76] In the work by Whitelam and Geissler, they are concerned with making the VMMC procedure as close to real world dynamics as possible. There are three main differences. First, they include a term to ensure equal probability for each particle to be moved in a given cluster move. This is not the case in this algorithm as particles that are tightly bound in a group of N particles are N times more likely to be moved than a single particle. Second, they include a rejection probability that is meant to impose the real world diffusion dynamics in the algorithm, selectively rejecting moves based on the size and shape of the cluster. These two features can be safely ignored if we are not interested in

modeling dynamics of the system. The last difference is that their linking probability is zero for any pair of particles that are not “interacting” in the μ configuration. The stated purpose for this is to make moves more local. The benefit of this is not clear compared to a method that links based on the full energy difference. While this can be beneficial for the algorithm, depending on your definition of “interacting”, we leave it out in this analysis as it complicates the acceptance rates and obscures the nature of the underlying algorithm. For a more in depth analysis of this aspect, and other differences between the original VMMC algorithm and the VMMC algorithms in this section, see Appendix C.

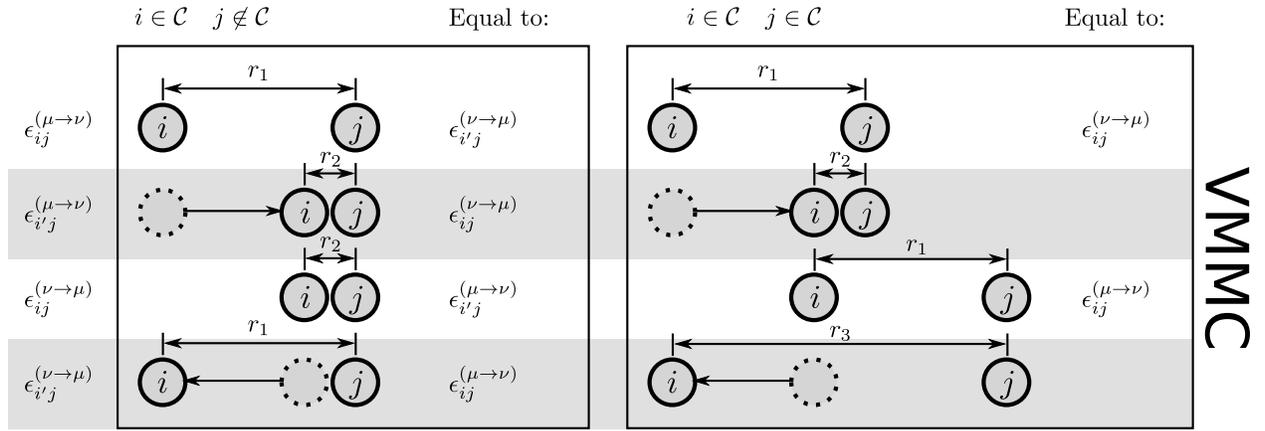


Figure 3.2: An explanation of the pairwise energy terms in the VMMC and related algorithms. This is much like Figure 3.1 except that there are factors in the $i \in \mathcal{C}, j \notin \mathcal{C}$ product that do not cancel. This can be traced back to the fact that the r_3 is not equal to r_2 , which is due solely to the fact that the VMMC moves are not self-inverse.

In the form given, the VMMC procedure is very similar to the Swendsen-Wang style cluster move, using the same linking function and symmetry operation technique, but due to the non-self-inverse symmetry operation, it is distinct. Following the derivation of an alternate version of the VMMC procedure by Whitelam and Geissler, we can see that these methods are actually closer than they initially appear.[76] In this version of the VMMC algorithm, we move more towards the rejection-free limit by shifting as much of the bias correction as possible into the move proposal.

Consider that the point of having a self-inverse symmetry operation is that it ensures that the bias in the case of $i, j \in \mathcal{C}$ cancels out but the bias in the case of $j \notin \mathcal{C}$ does not. We need these

bias factors when j is not in the cluster as these cancel the Boltzmann factors associated with the energy change. Thus we should change the linking probability in the case where links are made such that it naturally corrects for the bias terms that it introduces. We do this by introducing a reverse balance, f_{rev} defined in (3.14), that is included in the linking probability only when two particles link under the standard $p_{ij}(\mu \rightarrow \nu)$ linking probability. The reverse balance's numerator, $p_{ij}(\nu^* \rightarrow \mu)$, depends on a new system state, ν^* , which is the state where j has been added to the cluster, which we haven't determined is the case in ν .

$$f_{\text{rev}} = \frac{p_{ij}(\nu^* \rightarrow \mu)}{p_{ij}(\mu \rightarrow \nu)} \quad (3.14)$$

This adds a new way that a failed link can occur where i and j are linked according to $p_{ij}(\mu \rightarrow \nu)$ but failed to link according to f_{rev} . These links are what Whitelam and Geissler call *frustrated links*.^[76] Frustrated links signify a problem for detailed balance. There are now times when this extra factor will make certain moves impossible to produce in forward and reverse moves. This can be seen when considering a proposed VMMC move in a hard sphere fluid where the translation introduces an overlap between i and j . Due to the infinite energy of the overlap, we require that a link be formed in the $\mu \rightarrow \nu$ stage. However, in the reverse linking according to $p_{ij}(\nu^* \rightarrow \mu)$, there typically isn't an overlap at all meaning that the probability is zero.² This can be remedied by disallowing moves that involve states ν^* where $\nu^* \neq \nu$, i.e. the moves where there is an external frustrated link, a frustrated link between the cluster and the environment.³

With this new linking procedure and the new types of failed links we can write the acceptance probability as (3.15). Following the reasoning by Whitelam and Geissler, this can be simplified considerably.^[76] Once again, the product over external outright failed links with probability $q_{ij} = 1 - p_{ij}$ cancels the Boltzmann factors. We can safely ignore the product over \hat{q}_{ij} with j external to the cluster. This is the probability of frustrated links. As we have seen, any such frustrated link where $j \notin \mathcal{C}$ will necessarily result in a rejection as they explicitly violate detailed balance. The

² This is based on the heuristic that moving closer together can produce an overlap but moving further apart can't. This heuristic is not true when the coordinate has a periodic nature. This can arise in very large translations in systems with periodic boundary conditions and, more likely, when we generalize to rotations and extended objects.

³ Frustrated links internal to the cluster are fine. In fact much of the computation hinges on finding whether there is an external frustrated link or not, see Section C.1.

product over \bar{q}_{ij} , see (3.16), corrects for the probability of failed links internal to the cluster whether they be outright failures or frustrated links.⁴ Each ratio of the reverse to forward probabilities $\bar{q}_{ij}(\nu \rightarrow \mu)/\bar{q}_{ij}(\mu \rightarrow \nu)$ is unity due to the reverse balance. Lastly, the inclusion of the reverse balance ensures that each factor in the product over successful links is unity. This results in the acceptance probability (3.20) where $\delta_{\text{f.e.}}$ is one when there are no external frustrated links and zero otherwise.

$$\begin{aligned} \text{acc}(\mu \rightarrow \nu) &= \min \left(1, e^{-\beta(E(\nu)-E(\mu))} \prod_{i \in \mathcal{C}, j \notin \mathcal{C}} \frac{q_{ij}(\nu \rightarrow \mu)}{q_{ij}(\mu \rightarrow \nu)} \right. \\ &\quad \left. \times \prod_{i \in \mathcal{C}, j \notin \mathcal{C}}^{\mathcal{R}} \frac{\hat{q}_{ij}(\nu \rightarrow \mu)}{\hat{q}_{ij}(\mu \rightarrow \nu)} \prod_{i, j \in \mathcal{C}}^{\mathcal{R}} \frac{\bar{q}_{ij}(\nu \rightarrow \mu)}{\bar{q}_{ij}(\mu \rightarrow \nu)} \prod_{i, j \in \mathcal{C}}^{\mathcal{R}} \frac{p_{ij}(\nu \rightarrow \mu)}{p_{ij}(\mu \rightarrow \nu)} \right) \end{aligned} \quad (3.15)$$

$$\bar{q}_{ij}(\mu \rightarrow \nu) = q_{ij}(\mu \rightarrow \nu) + \hat{q}_{ij}(\mu \rightarrow \nu) \quad (3.16)$$

$$= 1 - p_{ij}(\mu \rightarrow \nu) + p_{ij}(\mu \rightarrow \nu) \left(1 - \min \left[1, \frac{p_{ij}(\nu \rightarrow \mu)}{p_{ij}(\mu \rightarrow \nu)} \right] \right) \quad (3.17)$$

$$= 1 - p_{ij}(\mu \rightarrow \nu) \min \left[1, \frac{p_{ij}(\nu \rightarrow \mu)}{p_{ij}(\mu \rightarrow \nu)} \right] \quad (3.18)$$

$$\Rightarrow \frac{\bar{q}_{ij}(\nu \rightarrow \mu)}{\bar{q}_{ij}(\mu \rightarrow \nu)} = 1 \quad (3.19)$$

$$\text{acc}(\mu \rightarrow \nu) = \delta_{\text{f.e.}} \quad (3.20)$$

An implementation of this algorithm is given in Algorithm 9. The algorithm needs to consider the reverse balance and thus it must consider the reverse inverse move, \mathcal{O}^{-1} . Rather than actually calculate the reverse operation, we note that for pairwise interactions, we can just as easily move the paired particle to produce the same effect as the inverse move as far as the pair potential is concerned. The parameter $\epsilon_{ij'}$ is the pair potential where the j particle is moved via \mathcal{O} . This has the added benefit of simplifying the notation within the algorithm so that we need not keep track of explicit $\mu \rightarrow \nu$, $\nu \rightarrow \mu$, or the special case for the reverse balance, $\nu^* \rightarrow \mu$.

It is important to note that the reverse bias is fundamentally encapsulating the fact that the operation is not self-inverse. In the self-inverse case, we can include this factor in the linking probability, but under those operations, $f_{\text{rev}} \equiv 1$ for the same reasons that the bias corrections involving $i, j \in \mathcal{C}$ cancel. This means that the including the reverse balance term in the linking

⁴ $\bar{q}_{ij} \equiv \bar{q}_{ij}$ in the notation of Whitlam and Geissler.

Algorithm 9 The VMMC algorithm in its Swendsen-Wang, ‘rejection-free’ form. This uses an operation, \mathcal{O} , and performs cluster move. As in Algorithm 6, new particles can be added to the cluster queue during the iteration over the cluster. The parameter $\epsilon_{ij'}$ is used instead of the more explicit $\epsilon_{ij}^{(\nu^* \rightarrow \mu)}$ to designate moving j via \mathcal{O} . This has the same effect as moving i by \mathcal{O}^{-1} for at least most operations. This was done to simplify the algorithm and keep all ϵ terms related to the forward operation.

```

1: procedure VMMC( $\{\vec{x}\}$ )
2:    $\mathcal{O} \leftarrow \{\text{translate}, \Delta x : \tilde{\delta} \times \text{random\_vector}(1.0)\}$ 
3:    $\text{seed} \leftarrow \text{random}(\text{size}(\{\vec{x}\}))$ 
4:    $(\mathcal{C}, \Delta \vec{U}, f) \leftarrow \text{VMMC\_BUILD\_CLUSTER}(\{\vec{x}\}, \text{seed}, \mathcal{O})$ 
5:   if  $\text{size}(f - \mathcal{C}) > 0$  then ▷ Check for frustrated external.
6:      $\delta_{\text{f.e.}} \leftarrow 0$ 
7:   else
8:      $\delta_{\text{f.e.}} \leftarrow 1$ 
9:   end if
10:  if  $1 = \delta_{\text{f.e.}}$  then
11:    for all  $i \in \mathcal{C}$  do ▷ Accept the move.
12:       $\{\vec{x}\}[i] \leftarrow \mathcal{O}(\{\vec{x}\}[i])$ 
13:    end for
14:    for all  $i$  such that  $\Delta \vec{U}[i] \neq 0$  and  $i \notin \mathcal{C}$  do
15:       $\Delta E \leftarrow \Delta E + \Delta \vec{U}[i]$  ▷ Energy only changes between particles not in the cluster.
16:    end for
17:  end if
18:  return  $\{\vec{x}\}$ 
19: end procedure
20:
21: function VMMC_BUILD_CLUSTER( $\{\vec{x}\}, \text{seed}, \mathcal{O}$ )
22:   $q \leftarrow [\text{seed}]$  ▷ Insert the seed particle into the cluster.
23:   $f \leftarrow []$  ▷ Initialize a list of frustrated links.
24:   $\Delta \vec{U} \leftarrow \vec{0}$  ▷ Zero out energetic contributions.
25:  for all  $i \in q$  do ▷ Loop over particles in the cluster.
26:    for all  $j \in \{1..N\}$  where  $j \notin q$  do ▷ Loop over particles that aren't in the cluster.
27:       $\Delta u \leftarrow (\epsilon_{ij} - \epsilon_{ij'})$  ▷ Energy change due to  $x_i \rightarrow \mathcal{O}x_i$  in  $\epsilon_{ij}$ .
28:      if  $\text{random}(1.0) > 1 - e^{-\beta \Delta u}$  then
29:         $\Delta u' \leftarrow (\epsilon_{ij'} - \epsilon_{ij})$ 
30:         $f_{\text{rev}} \leftarrow 1 - e^{-\beta \Delta u'} / 1 - e^{-\beta \Delta u}$ 
31:        if  $\text{random}(1.0) < f_{\text{rev}}$  then
32:           $q \leftarrow \text{append}(q, [j])$  ▷ Link: add  $j$  to the cluster.
33:        else
34:           $f \leftarrow \text{append}(f, [j])$  ▷ Note frustrated link.
35:        end if
36:      else
37:         $\Delta \vec{U}[j] \leftarrow \Delta \vec{U}[j] + \Delta u$  ▷ Do not link. Tally the energetic contribution.
38:      end if
39:    end for
40:  end for
41:  return  $(q, \Delta \vec{U}, f)$  ▷ Return information about the move.
42: end function

```

procedure is actually consistent with the generalized Swendsen-Wang style algorithm. In other words, we can establish that the alternate version of VMMC by Whitelam and Geissler is the same algorithm that we have been talking about.[76] The $\delta_{f.e.}$ term is a direct consequence of the non-self-inverse symmetry operation and in the case of the other methods, this factor is always one. This algorithm can in some ways be thought of as the “rejection-free” version of the VMMC procedure; it is as close as we can get to truly rejection-free with non-self-inverse moves within the Swendsen-Wang style cluster move framework.

It is helpful to understand the types of collective motion made possible by the Swendsen-Wang style cluster moves. The main benefit of these moves, indeed the reason that they can be rejection-free, or nearly rejection-free, is that they propose moves that are sensitive to the potential of the system. Further, when using symmetry operations that make local changes to the system, as VMMC was designed to do, the move proposal method is dependent on an approximate gradient of the system potential. Because of this, the methods show excellent promise in fluid simulation, especially simulations with aggregation phenomena.

However, there are cases where these cluster moves fail to do any good. Of particular interest to this work is the case of purely repulsive potentials, especially isotropic repulsive potentials. Included in this class of potentials are hard sphere systems and the HCSS systems from Chapter 2. These systems don’t show aggregation phenomena, but they do show clustering, and thus it seems possible that applying Swendsen-Wang style moves could be beneficial. If we attempt to apply the translation symmetry operation to the HCRS system, we find that any link that forms will be a frustrated external link. This is due to the fact because the potential is purely repulsive, the reverse balance is zero whenever the forward balance is non-zero. This reveals a property of the Swendsen-Wang style cluster moves: non-frustrated links can only occur between particles i and j if $u(\sigma_i, \sigma_j) < u(\mathcal{O}\sigma_i, \sigma_j)$ and $u(\sigma_i, \sigma_j) < u(\mathcal{O}^{-1}\sigma_i, \sigma_j)$ where $\sigma_{i/j}$ is the state of particle i/j , $u(\sigma_i, \sigma_j)$ is the pair potential between particles i and j , and $\mathcal{O}\sigma_i$ is the state of particle i after transformation by \mathcal{O} . In the case of self-inverse moves this is exactly what you want. Since $u(\mathcal{O}\sigma_i, \sigma_j) = u(\mathcal{O}^{-1}\sigma_i, \sigma_j)$, this simply means that if the operation makes i and j reduce their pair

potential, then there is no link, otherwise there is a chance of a link. In the case of non-self-inverse operations like translation, this means that the pair potential of i and j needs to be approximately at a local minimum or there is a possibility of frustrated external link. In the case of purely repulsive potentials, where there is no local minimum, and therefore any link that forms will necessarily be an external frustrated link. A useful intuition for how these cluster algorithms work, especially when it comes to translations and rotations, is that particle i can only link with particle j if it will *push* (*pull*) it under the forward operation and *pull* (*push*) it under the reverse.

3.3.3 Rejection and efficiency

As we have shown, selecting the linking probability $p_{ij}(\mu \rightarrow \nu)$ as in (3.1) ensures that the move proposal is done in such a way that the method is rejection-free for non-inverse transformations. We have also noted that even in these rejection-free systems, if the symmetry is applied to the entire system then configuration space isn't being explored effectively at all. In fact, rejection-free algorithms are often less efficient at sampling phase space than competing algorithms that allow rejections.

This is due to the fact Monte Carlo necessarily involves a random component to the procedure. This random component means that some of the starting seed moves, the seed particle and the symmetry operation, will result in significant increases in the system energy. After we have evaluated all of the particles interacting with the seed particle, we know what the Metropolis acceptance rate of such a move would have been. In the Swendsen-Wang style cluster moves, the algorithm attempts to fix high energies by linking particles and adding them to the cluster. As we have noted, this can be beneficial, but it is also extra computational work. Sometimes it would simply be faster to outright reject this move and try again with a new random seed and operation on to the next iteration. Indeed, including the possibility of move rejection within a rejection-free method via an early exit scheme is often a good way to speed up sampling by weeding out cluster moves that are taking too long.

Sometimes, however, adding rejection to a rejection-free method can provide completely dif-

ferent modes of relaxation of structure which the rejection-free method cannot. We recall that one of the primary benefits of Monte Carlo sampling over molecular dynamics is that Monte Carlo transitions need not be related to physical transitions. It is not surprising then, that departing from the linking probability that is physically motivated can be beneficial. It is in this vein that Niedermayer and, much later, Whitlam and Geissler designed methods that can trade some rejection for better structural relaxation.[54, 77, 76] The Whitlam and Geissler method is known as *cluster cleaving* and it uses a linking probability according to a false higher temperature (inverse temperature β^*) than the true system temperature (inverse temperature β). This has the effect of reducing all linking probabilities in the system in a very natural way. As this changes the probability of failed links between the cluster and the environment, this modifies the acceptance probability by no longer completely canceling the Boltzmann factors of the move, see (3.21). All of the other factors in the acceptance probability still simplify in the same manner.

$$\text{acc}(\mu \rightarrow \nu) = \min(1, \delta_{\text{f.e.}} e^{-(\beta-\beta^*)(E^{(\nu)}-E^{(\mu)})}) \quad (3.21)$$

Cluster cleaving is similar in purpose to the selection of the clustering cutoff in the proximity clustering algorithm from Section 3.2. In contrast to the proximity methods, the cleaving method still links according to the potential, but with a natural adjustment that statistically limits the size of the cluster. It is also not dependent on the proximity clustering method's assumption that the closer particles are more likely to belong to the same cluster. The cleaving method is a clean way to tune the algorithm for the relaxation of degrees of freedom internal to the cluster. It provides a tunable parameter that shifts the bias correction terms between the move proposal and acceptance probability. Ordinary Metropolis single particle style moves occur when $\beta^* = 0$ and the standard rejection-free or nearly rejection-free methods when $\beta^* = \beta$. It should be noted that while Whitlam and Geissler only described this technique in relation to VMMC, this method can be applied to any of the Swendsen-Wang style moves discussed here.

3.3.4 VMMC with augmented potentials

Whitelam and Geissler noted any linking potential can be used in the VMMC procedure.[77] We sought to use this fact to extend the VMMC procedure to apply to purely repulsive, isotropic potentials. The goal is to encourage clusters of HCSS spheres to move in concert as well as providing a basic mechanism for recovery from overlaps found in simple hard sphere systems. We do this by changing our linking potential to (3.22), resulting in a linking probability of $p_{ij}^*(\mu \rightarrow \nu)$, see (3.23). The augmentation potentials will, in general, be attractive wells that are added to the surface of the hard sphere potential. This will make it possible to propose moves that involve links with other particles even when there is no pair potential minimum between the two particles.

$$\epsilon_{ij}^* = \epsilon_{ij} + w_{ij} \quad (3.22)$$

$$p_{ij}^*(\mu \rightarrow \nu) = 1 - e^{-\beta((\epsilon_{i'j}^{(\mu \rightarrow \nu)} + w_{i'j}^{(\mu \rightarrow \nu)}) - (\epsilon_{ij}^{(\mu \rightarrow \nu)} + w_{ij}^{(\mu \rightarrow \nu)}))} \quad (3.23)$$

$$= 1 - p_{ij}(\mu \rightarrow \nu) \times e^{-\beta(w_{i'j}^{(\mu \rightarrow \nu)} - w_{ij}^{(\mu \rightarrow \nu)})} \quad (3.24)$$

It should be noted that the cluster cleaving method can be seen as a special case of the potential augmentation where the augmentation is to counteract a portion of the actual pair potential. Again, the same analysis of the bias correction applies. The factors in the bias correction that cancel with their reverse move, the factors for failed internal links, $\bar{q}_{ij}(\mu \rightarrow \nu)$, and the factors for successful links, $p_{ij}(\mu \rightarrow \nu) \min\left(1, \frac{p_{ij}(\nu \rightarrow \mu)}{p_{ij}(\mu \rightarrow \nu)}\right)$, will continue to do so. The acceptance probability should be the same with the exception of the corrections due to failed links between the cluster and the environment. These failed links usually cancel the exponential factors in the acceptance probability. When an augmented potential is used, they still cancel those factors, but they introduce additional factors that, necessarily, must reduce the acceptance probability, see (3.26). We can see this by adding an augmentation to (3.12) as shown in (3.25). The effect on the acceptance probability is to increase the rejection of moves that happen to decrease the augmented potential of the system. This means that this algorithm disfavors changes to the system that move it closer to states where it is more likely to move collectively according to the augmented potential. For the typical short range attraction like augmentation, this means that merging clusters is disfavored.

This makes sense as the reverse move, splitting clusters, is proportionately unlikely.

$$\begin{aligned}
& \prod_{i \in \mathcal{C}, j \notin \mathcal{C}}^{\epsilon_{i'j}^{(\mu \rightarrow \nu)} > \epsilon_{ij}^{(\mu \rightarrow \nu)}} \frac{1}{e^{-\beta(\epsilon_{i'j}^{(\mu \rightarrow \nu)} - \epsilon_{ij}^{(\mu \rightarrow \nu)})}} \times \prod_{i \in \mathcal{C}, j \notin \mathcal{C}}^{\epsilon_{i'j}^{(\mu \rightarrow \nu)} < \epsilon_{ij}^{(\mu \rightarrow \nu)}} \frac{e^{+\beta(\epsilon_{i'j}^{(\mu \rightarrow \nu)} - \epsilon_{ij}^{(\mu \rightarrow \nu)})}}{1} \\
& \times \prod_{i \in \mathcal{C}, j \notin \mathcal{C}}^{w_{i'j}^{(\mu \rightarrow \nu)} > w_{ij}^{(\mu \rightarrow \nu)}} \frac{1}{e^{-\beta(w_{i'j}^{(\mu \rightarrow \nu)} - w_{ij}^{(\mu \rightarrow \nu)})}} \times \prod_{i \in \mathcal{C}, j \notin \mathcal{C}}^{w_{i'j}^{(\mu \rightarrow \nu)} < w_{ij}^{(\mu \rightarrow \nu)}} \frac{e^{+\beta(w_{i'j}^{(\mu \rightarrow \nu)} - w_{ij}^{(\mu \rightarrow \nu)})}}{1} \quad (3.25) \\
\text{acc}_{\mu \rightarrow \nu} &= e^{+\beta(E^{(\nu)} - E^{(\mu)})} \times e^{+\beta(w^{(\nu)} - w^{(\mu)})} \quad (3.26)
\end{aligned}$$

The pseudocode for the augmented potential VMMC method is in Algorithms 10 and 11. It uses the same tricks that Algorithm 9 uses to work around calculating inverse operations. As stated earlier, the VMMC procedure only links when particles both push and pull one another under the operation, \mathcal{O} . With self-inverse operations, this is always the case. With non-self-inverse symmetry operations, there is a certain chance that a move that will be rejected due to the presence of frustrated external links, a particle that pushes (or pulls) but doesn't pull (or push). The augmented potential method increases the likelihood of forming links and, when used wisely, can decrease the chances of frustrated links. Indeed, a significant number of the links that the augmented potential forms would have been frustrated external links. This naturally incurs a decrease in acceptance probability as the Boltzmann factor is no longer exactly canceled by the linking bias correction. The augmented potential method should be seen as a way to tune the algorithm via optimization of potential details such as range and depth, to strike a balance between the rejections due to frustrated external links and the uncanceled Boltzmann factors in the acceptance probability. Once again this is done by shifting some knowledge of the system back to the move selection stage, however in this case the knowledge is based on “on the fly” style optimization of parameters that depends on the system specification and its current state in non-trivial ways.

3.4 Application of VMMC and VMMC_{aug}

This was implemented and applied to several systems including a standard two dimensional Lennard-Jones system, the Tavares two dimensional lattice model for patchy colloids, the two

Algorithm 10 The Augmented Potential VMMC algorithm. This is a modification of the ‘rejection-free’ version of the VMMC which is suitable for purely repulsive and other systems where the VMMC procedure is known to be unproductive. This uses an operation, \mathcal{O} , to build a cluster via $\text{VMMC}_{\text{aug_BUILD_CLUSTER}}$ (see 11) and performs cluster move.

```

1: procedure  $\text{VMMC}_{\text{aug}}(\{\vec{x}\})$ 
2:    $\mathcal{O} \leftarrow \{\text{translate}, \vec{\Delta}x : \delta \times \text{random\_vector}(1.0)\}$ 
3:    $\text{seed} \leftarrow \text{random}(\text{size}(\{\vec{x}\}))$ 
4:    $(\mathcal{C}, \Delta\vec{U}, \Delta\vec{w}, f) \leftarrow \text{VMMC}_{\text{aug\_BUILD\_CLUSTER}}(\{\vec{x}\}, \text{seed}, \mathcal{O})$ 
5:   for all  $i$  such that  $\Delta\vec{w}[i] \neq 0$  and  $i \notin \mathcal{C}$  do ▷ Compute augmented bias.
6:      $\Delta W \leftarrow \Delta W + \Delta\vec{w}[i]$  ▷ Only the bias of external failed links matters.
7:   end for
8:   if  $\text{size}(f - \mathcal{C}) > 0$  then ▷ Check for frustrated external.
9:      $\delta_{\text{f.e.}} \leftarrow 0$ 
10:  else
11:     $\delta_{\text{f.e.}} \leftarrow 1$ 
12:  end if
13:  if  $\text{random}(1.0) < \delta_{\text{f.e.}} e^{+\beta\Delta W}$  then
14:    for all  $i \in \mathcal{C}$  do ▷ Accept the move.
15:       $\{\vec{x}\}[i] \leftarrow \mathcal{O}(\{\vec{x}\}[i])$ 
16:    end for
17:    for all  $i$  such that  $\Delta\vec{U}[i] \neq 0$  and  $i \notin \mathcal{C}$  do
18:       $\Delta E \leftarrow \Delta E + \Delta\vec{U}[i]$  ▷ Energy only changes between particles not in the cluster.
19:    end for
20:  end if
21:  return  $\{\vec{x}\}$ 
22: end procedure

```

Algorithm 11 The Augmented Potential VMMC cluster building function. This uses the given operation \mathcal{O} and bias pair potential w_{ij} to build a cluster linked by a fictitious potential. As in Algorithm 6, new particles can be added to the cluster queue during the iteration over the cluster.

```

1: function VMMCaug-BUILD_CLUSTER( $\{\vec{x}\}$ , seed,  $\mathcal{O}$ )
2:    $q \leftarrow$  [seed] ▷ Insert the seed particle into the cluster.
3:    $f \leftarrow []$  ▷ Initialize a list of frustrated links.
4:    $\Delta\vec{U} \leftarrow \vec{0}$  ▷ Zero out energetic contributions.
5:    $\Delta\vec{W} \leftarrow \vec{0}$  ▷ Zero out augmented potential contributions.
6:   for all  $i \in q$  do ▷ Loop over particles in the cluster.
7:     for all  $j \in \{1..N\}$  where  $j \notin q$  do ▷ Loop over particles that aren't in the cluster.
8:        $\Delta u \leftarrow (\epsilon_{ij} - \epsilon_{ij})$  ▷ Energy change due to  $x_i \rightarrow \mathcal{O}x_i$  in  $\epsilon_{ij}$ .
9:        $\Delta w \leftarrow (w_{ij} - w_{ij})$  ▷ Augmented potential change due to  $x_i \rightarrow \mathcal{O}x_i$  in  $w_{ij}$ .
10:      if  $\text{random}(1.0) > 1 - e^{-\beta(\Delta u + \Delta w)}$  then
11:         $\Delta u' \leftarrow (\epsilon_{ij'} - \epsilon_{ij})$ 
12:         $\Delta w' \leftarrow (w_{ij'} - w_{ij})$ 
13:         $f_{\text{rev}} \leftarrow 1 - e^{-\beta(\Delta u' + \Delta w')} / 1 - e^{-\beta(\Delta u + \Delta w)}$ 
14:        if  $\text{random}(1.0) < f_{\text{rev}}$  then
15:           $q \leftarrow \text{append}(q, [j])$  ▷ Link: add  $j$  to the cluster.
16:        else
17:           $f \leftarrow \text{append}(f, [j])$  ▷ Note frustrated link.
18:        end if
19:      else
20:         $\Delta\vec{U}[j] \leftarrow \Delta\vec{U}[j] + \Delta u$  ▷ Do not link. Tally the energetic contribution.
21:         $\Delta\vec{W}[j] \leftarrow \Delta\vec{W}[j] + \Delta w$  ▷ Tally the augmented potential contribution.
22:      end if
23:    end for
24:  end for
25:  return ( $q, \Delta\vec{U}, \Delta\vec{W}, f$ ) ▷ Return information about the move.
26: end function

```

dimensional hard-core repulsive-step system.[69] These systems provide a large variety of systems that test the usefulness of the VMMC and VMMC_{aug} algorithms.

Correctness of the algorithm is verified by comparison of the Lennard-Jones fluid to previous work.[5] The standard VMMC approach described here is used to demonstrate that VMMC enables collective motion which aids the proper sampling of aggregation relaxation. The trace of the energy as the system equilibrates shows the ability for these methods to make major improvements over naive Metropolis methods. With that in mind, I have found no clear evidence of this method performing better than other methods of ad hoc cluster moves, such as those in Section 3.2. The use of VMMC_{aug} configured to act as the cluster cleaving VMMC algorithm, shows enhanced ability for the system to relax internal structure, producing more globular than dendritic configurations. Again, when compared casually to the methods of Section 3.2, it is not clear what outperforms. These achieve the goal of internal relaxation in different ways, the VMMC method being somewhat more generally applicable, but not clearly superior in practice in the systems that were examined.

There is considerable fine-tuning of the moves that must be performed. Standard optimization is often not sufficient to get good sampling results. To illustrate this, in the case of aggregation the system shows multiple important length scales for motion. As discussed at the beginning of this chapter, one length scale is the mean distance between particles within an aggregate while an independent length scale is the mean distance between aggregates. Simple optimization like the method described in Section A.1 is not able to accommodate both of these, leading to one being selected over the other. Indeed, the method of selection of \mathcal{O} itself is not compatible with multiple length scales. This can be fixed in various ways, but it adds to the complexity of the method and this is a general feature of any collective move.

The lattice system of Tavares is a lattice model meant to model systems patchy colloidal systems.[69] In this model, particles have a location in the lattice and an discrete orientation along the x or y axis. Along with volume exclusion, the particle's have attractive sites on their surface.⁵ If the particles are adjacent and aligned such that their attractive patchy spots are adjacent, then

⁵ I also examined generalization of this system to elongated particles that occupy multiple lattice sites, but will

the pair energy is $-\epsilon$. This system can be seen as a kind of self-assembled Zwanzig model.[81, 12] Application of VMMC and VMMC_{aug} to this system has demonstrated the algorithms' ability to sample the system using collective moves. In the case of the VMMC_{aug} algorithm, I used the augmented potential to make the pair potential isotropic using parameter λ which turns on the anisotropy of the linking potential. In this scheme, $\lambda = 1$ means the augmented potential is off and thus normal VMMC is used and $\lambda = 0$ means that the augmented potential has made the system link as if the attraction is isotropic.

When dealing with the selection of a potential augmentation, it is helpful to remind ourselves of why we might think this is a good idea. In this case, we see that in all but one special case, neighboring particles are usually in a non-interacting configuration. This means that VMMC cannot give move these particles in the concert. In order to provide some possibility for these collective motions, I introduce a potential well around the entirety of the particle, increasing the likelihood of a link forming (but subsequently reducing the likelihood that the resulting move will be accepted).

The algorithms are correct⁶ and recreate the cluster size distribution in the system, see Figure 3.4, and are much more effective than standard site displacement sampling schemes. The VMMC_{aug} method shows its expected behavior, it increases the likely hood of links forming and thus tends to increase the size of the clusters that it attempts to move, but at the same time reduces the probability of these moves being accepted, see Figure 3.4. The key think to note here is that, while the acceptance rate is hurt for smaller clusters and the overall rejection rate is higher, the acceptance rate of larger clusters is increased. This indicates a different method of sampling which could, in principle provide better sampling. In these studies, however, we were not able to confirm that this method performed better than standard VMMC methods or even the non-local moves utilized by Tavares et al. in the original paper for that matter.

Lastly, we consider the HCRS system of Chapter 2. We know that the structure of the system

only be discussing the system with isotropic volume exclusion.

⁶ I actually implemented the Whitelam and Geissler algorithms to the minute detail as described in [77, 76] in every way I could conceive of their text mapping to a program as well as a few of my own versions. These all showed the same result demonstrating the techniques as correct.

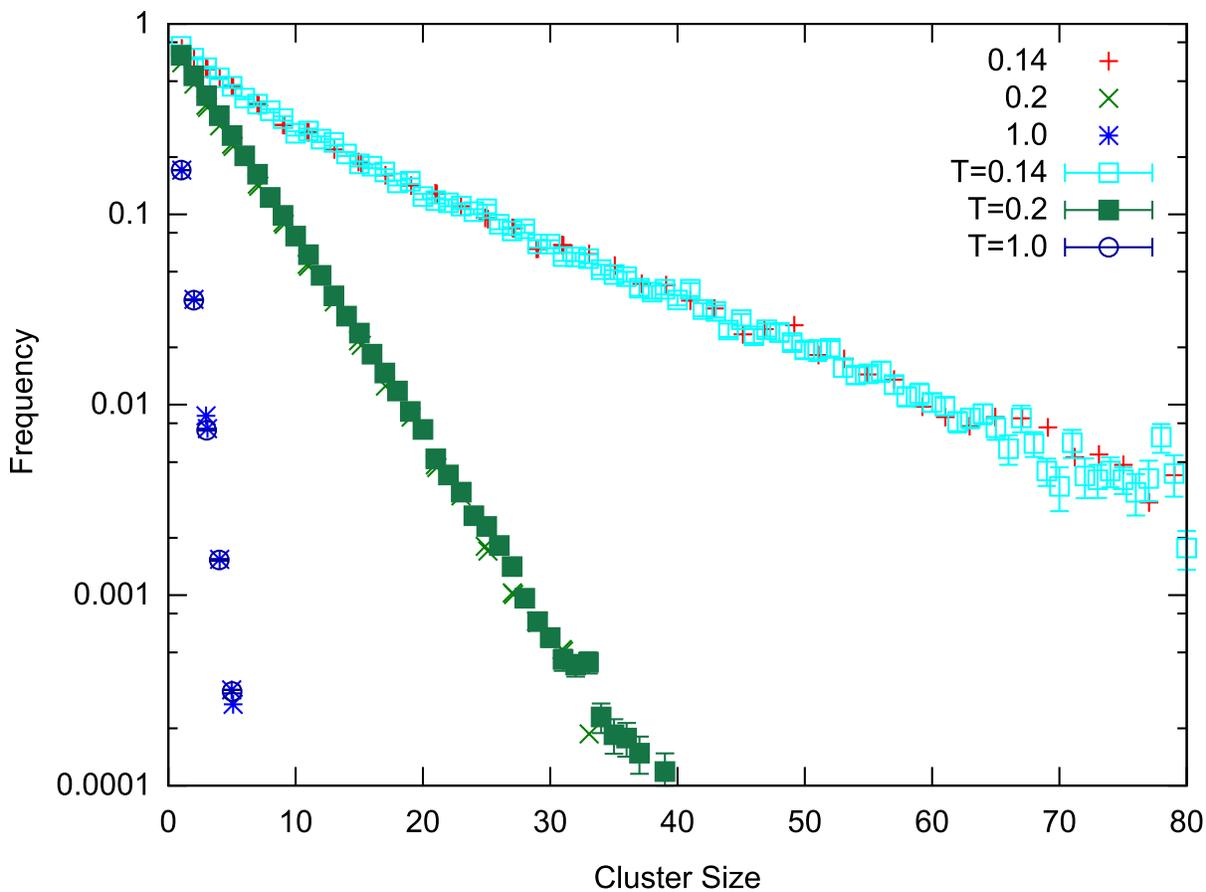


Figure 3.3: A comparison of the results of Tavares et al. to our calculation using various VMCM procedures. Only the lowest temperature distribution is in the model's nematic phase, while the higher two temperatures are in the isotropic phase.

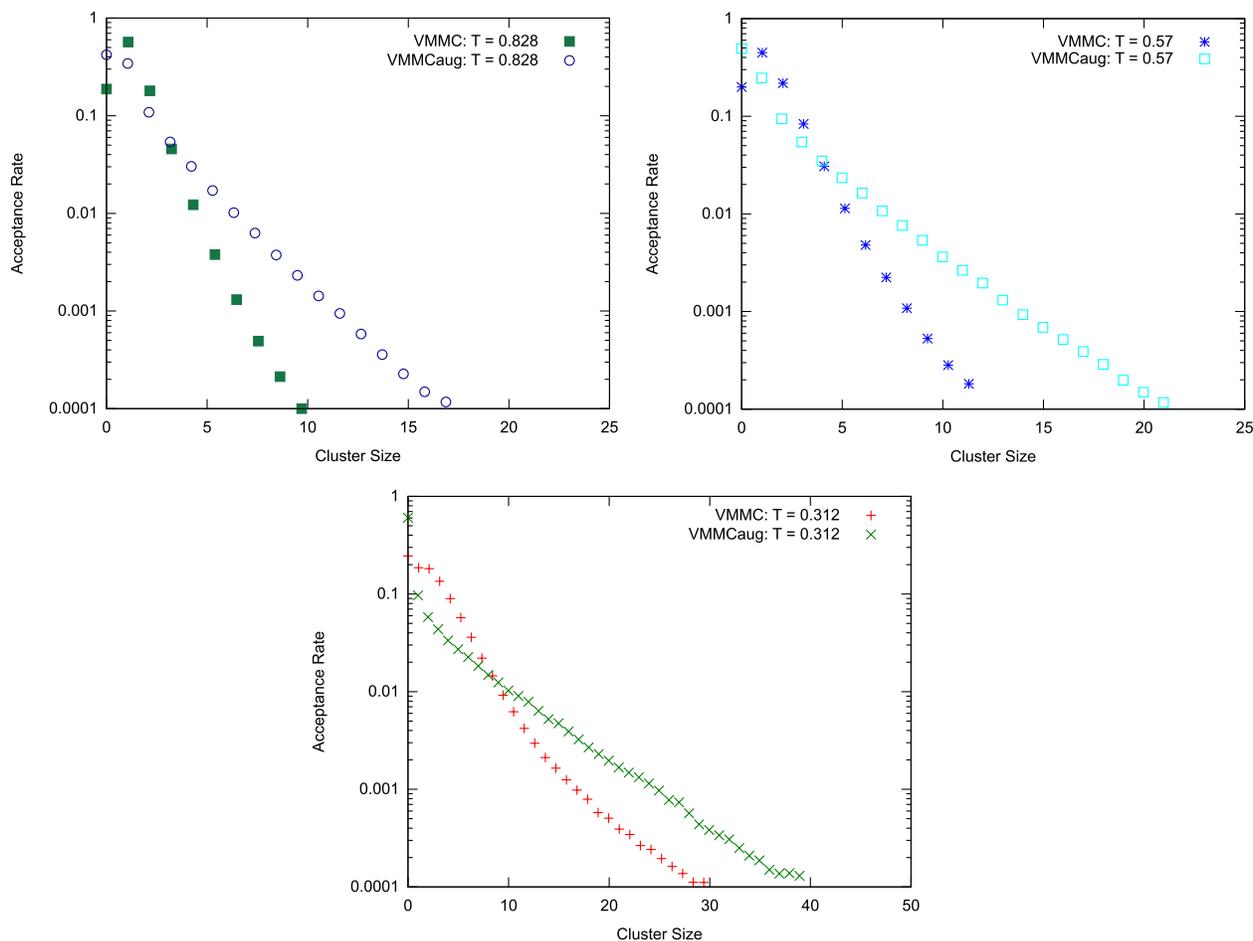


Figure 3.4: A comparison of the acceptance rates of the VMMC_{aug} algorithm using the $\lambda = 0$ augmented potential (isotropic linking potential) versus the normal VMMC procedure. In this histogram, the zero bin indicates moves that were rejected for whatever reason. The VMMC_{aug} method makes it more likely for larger clusters to be successfully moved, but at the cost of lowering the overall acceptance rate. This is seen by the fact that the acceptance rates of the VMMC_{aug} method actually drop below the acceptance rates of the standard VMMC procedure in the small cluster limit. This is particularly true of the lowest temperature data which is in the nematic phase.

includes inhomogeneous density, clustering, and even liquid crystalline like structures, which implies that algorithms that help with the equilibration of self-assembled aggregates might be of service here. However, this system is useful to probe the utility of the augmented approach as it is purely repulsive and thus it is known that VMMC is not a useful technique. As noted at the end of Section 3.3.2, the VMMC algorithm fails completely when used with purely repulsive potentials. Once again, this is a direct result of the non-self-inverse nature of the translation operation. These repulsive potentials tend to form links in one direction but will invariably fail to form the reverse link resulting in frustrated links which almost always lead to a failure at acceptance time. In the end, the process results in an expensive site displacement algorithm.

To introduce the possibility of collective motion into this sampling scheme, I applied the VMMC_{aug} algorithm using attractive square wells as the potential augmentation. Choice of the augmenting potential can have a profound affect on the nature of the moves proposed. For instance, using a well just outside the repulsive step diameter will result in particles getting linked that separated by around the shoulder length. This sort of augmentation will quite likely percolate to the entire system in some cases. We can contrast this with the choice of augmenting well just outside the core diameter. This will tend to only link particles that are roughly core diameter apart. This has the nice feature that in you are micellar phases, the cluster selection will rarely, if ever, extend past the micelle.⁷ As seen in the pair distributions of fluid micellar phases in Figure 2.4, the position of HCRS particles in neighboring clusters is only weakly correlated if at all. This indicates that the well outside the core diameter is a better choice. Much as in the patchy colloid lattice model, analysis of VMMC_{aug} with these augmenting potentials, and the combination of the two, shows that the VMMC_{aug} procedure works as expected. At high temperatures or low pressures, the algorithm finds small (but $N > 1$) clusters where none would be created via VMMC. At lower temperatures or higher pressures, the clusters tend to percolate to the system size and produce no relaxation. Sadly, even in the high temperature case, I once again see no evidence that

⁷ The trade-off here is that the move will be more likely to be rejected due to an introduced repulsive step overlap that cannot be accounted for.

System	Method	τ [s]	τ [million interactions]
fluid micellar	Metropolis	0.787	5.76
fluid micellar	VMMC _{aug}	8.84	22.0
crystal micellar	Metropolis	1.35	15.7
crystal micellar	VMMC _{aug}	133	359

Table 3.1: Measured decorrelation times for sample HCRS systems. Decorrelation times are measured from NVT fluid and crystalline micellar systems using 100% Metropolis sampling and 100% VMMC_{aug} sampling. These are scaled to two different metrics of time, the wall time and the number of pair interactions. The pair interactions can be considered more accurate here as they are insensitive to implementation details and more accurately predict the theoretical limit of the algorithm’s implementation. Metropolis wins across the board but the difference is particularly striking when dealing with crystalline systems where the cluster moves percolate to the system size and become unusable.

the algorithm has any efficiency gain over normal site displacements, see Table 3.4.

3.5 Conclusion

I have shown how the GenGCA procedure can be seen as a generalization of the Wolff algorithm and how both are based on the principles of the Swendsen-Wang algorithm. I presented a generalized version of this algorithm that encompasses both algorithms as well as any single cluster Swendsen-Wang style algorithm based on self-inverse moves. I have also presented a generalized version of that algorithm that encompasses non-self-inverse moves such as those used in the VMMC algorithm as well as self-inverse moves, where the algorithm reduces to the self-inverse case when such an operation is specified. This algorithm is recognized as identical to a version of the VMMC algorithm as presented by Whitelam and Geissler.[76] This version of the VMMC algorithm is examined and the similarities between the Swendsen-Wang algorithm are exposed as well as the *nearly* rejection-free status of VMMC. These algorithms are essentially the same but have different properties due to the system symmetry and operations used. For instance, I implicitly show that if an inversion were to be used in this generalized algorithm, then the GenGCA algorithm would be recovered exactly as laid out in the papers by Liu and Luijten.[47, 48]

This latter algorithm is extended to apply to systems with isotropic, purely repulsive pair interactions. The VMMC procedure is shown to be completely useless for these types of interactions,

reducing to its Metropolis style, single particle move limit. This result is explained as well as why this is not the case in the related GenGCA algorithm and is shown to be a limitation of the use of non-self-inverse moves in the algorithm. The generalization of the VMMC procedure utilizes an augmented potential where the linking probability is the sum of the standard pair interaction and a minor attractive component to drive linking. This bias to the move proposal reduces the acceptance probability but simultaneously reduces the probability of frustrated external links, the only thing keeping the VMMC procedure from being a rejection-free method.

Future work on this project would be to test what novel self-inverse operations can be found for sampling interesting systems. It would also be very interesting to evaluate the performance of the original VMMC algorithm compared to the alternate version that we use here. In particular, what benefit do rejection-free methods give you in VMMC? The VMMC procedure could be applied to several interesting projects such as aggregation of systems with highly anisotropic pair potentials. It would be interesting to see what performance gains are made using the GenGCA procedure in systems utilizing the Asakura-Oosawa-Vrij depletion model where depletant particles are modeled as non-mutually interacting spheres but which have excluded volume interactions with the colloids.[70] This method is used fairly commonly and, due to the large size anisotropies associated with the technique, GenGCA should help with significantly.

Future work regarding the augmented potential method is to apply it to more systems with hard core or otherwise purely repulsive interactions and to examine the systems we've looked at in further detail. This method could very well still be applicable as a tool for sampling HCRS systems which exhibit clustering phenomenon but no attractive interactions. In addition to these systems, any system with a pair potential that is gradient free, such as a simple attractive square well, will demonstrate similar problems with standard VMMC as the reverse balance will neglect any possibility of linking for moves significantly smaller than the well width. While this method has been proven correct analytically, it would be nice to see it succeed when applied to some interesting problems.

Chapter 4

Spontaneous twist in bundles of rods

Chirality is an extremely common aspect of the physical world which often shows up in unexpected places. From snail shells to nucleic acid helices to self-assembled helical filaments in liquid crystal phases, chiral structures permeate the study of the physical world. Understanding the mechanisms by which chiral structures form is integral to understanding the physical world around us. In recent years several papers have been published demonstrating the somewhat surprising tendency for simple systems of achiral objects such as assemblies of rod shaped macro-molecules and filaments bundles to prefer chiral structures for their ground state configuration. Bruss et al. demonstrated that filament bundles lower their energy when they twist despite the fact that it takes energy to bend the filaments in order to make the bundle.[6, 7] They attribute the tendency for twist via a competition between several energetic components, namely the packing of the filament cross-sections, the cohesion energy between rods, and the bend energy to deform the filament into a helix. Gibaud et al. demonstrated that self-assembled membranes of 880nm-long rod-like fd virus generate vortex bundles, or disks composed of a single layer of virus displaying macroscopic twist.[26] Yang et al. modeled this system with hard spherocylinder rods with depletion-mediated attraction and identified surface tension as the mechanism producing the twist. The twist is explained as a mechanism for achieving the observed lenticular shape of the vortex bundle which is reported to minimize surface area.[79] Each of these studies identifies this effect as an energetically driven phenomenon.

In this section, we offer analysis of a simple model of bundled hard rods which shows chiral

self-assembly via an as yet unreported, purely entropic, mechanism. We use umbrella sampling techniques to make explicit *in silico* measurements of the free energy profile for bundle twist. We identify the origin of this entropic effect via a cell theory approximation and demonstrate the applicability of this mechanism for twist to a broad variety of related bundle-like systems.

Our interest in this phenomenon began from what appeared to be an error in the simulation of systems of self-assembled colloidal bundles. These bundles made up of hard rods with pairwise distributed attraction tended to form twisted structures even though there were no chiral interactions within the system. This spontaneous chirality was baffling as it seems unfavorable from both an energetic and entropic standpoint. Any distortion of the nematic order necessarily decreases the packing efficiency. For attractive interactions, this can only increase the energy of the structure. Further, decreased packing efficiency also decreases the translational entropy of the system. With this in mind, it was somewhat surprising to us to observe nearly universal twist in the self-assembled bundles.

As it turns out, the behavior we witness is real, simple in origin, applies to all bundle-like systems, and is even something that is encountered in everyday macroscopic life. Everyday objects such as a bundle of pencils held by a rubber band or a sheaf of wheat overwhelmingly show this twisting behavior when agitated. Anybody who has carried a bundle of sticks, pipes, or other objects with large aspect ratios, under their arm while walking is familiar with this effect. The tendency for the objects in the bundle to splay at their ends is immense, making the bundles extremely difficult to carry. The most natural way for this splay to occur is via a macroscopic twist in the bundle as it allows for splay without strongly disrupting the packing. This is a behavior that is common in nature, from the macroscopic to the microscopic scales, and understanding it is essential to fundamentally understanding the world around us.

Our work shows that the twisted state is the ground state for bundles as verified by umbrella sampling calculations and a cell theory approximation. We find that the entropy is maximized and the free energy is minimized at non-zero twist for many all bundles with a large enough aspect ratio. Due to this phenomenon's counterintuitive nature, it is often neglected in the modeling and

analysis of twisting behavior in bundles.

4.1 The model and methodology

To understand the origin of this behavior, we turn to explicit models of bundles rather than their self-assembled counterparts. We begin by examining the behavior of explicitly defined bundles of hard spherocylinder rods. Bundles are defined by constraining the centers of the rods of aspect ratio L/D to a disk of diameter D_b/D within the $z = 0$ plane as illustrated in Figure 4.1. The bundle has a set area density $\rho = \frac{N}{\pi(D_b/2D)^2}$. The constraint of centers of the spherocylinders to the disk can be seen as a type of distributed attraction. In this way, the model is a simplified version of the self-assembled system examined by Yang et al.[79] However, in our system the number of degrees of freedom has been significantly reduced, making computation cheaper and analysis considerably easier.

We use standard NVE Monte Carlo Metropolis sampling to explore the phase behavior of these bundles. These explicit bundles readily form twisted structures similar but qualitatively different from those seen in simulations of self-assembled colloidal systems. A cursory survey of the model parameters reveals that the tendency for twist seems to increase with increasing rod aspect ratio. In fact, it seems that systems with large bundle aspect ratios, the length of the rods, L , divided by the diameter of the bundles, D_b , tended to twist where bundles with small aspect ratios tended not to. The bundle aspect ratio is a useful predictor for twist. Bundles that look like bundles, i.e. are elongated along the bundle axis, tend to twist whereas groups that look like membranes, i.e. are pancake shaped, tend to remain straight. The area density, ρ , also plays a role here and promotes twisted structures at smaller bundle aspect ratios.

Since all energetic components of the free energy have been removed, the mechanism for twist must be entropic. This is interesting as at first it seems that introducing a distortion to the nematic order must necessarily decrease the system's ability to pack efficiently, and thus decrease the entropy. While this is true regarding the translational entropy, there is another source of entropy in these systems, orientational entropy. Indeed, we see that the ends of the rods get quite

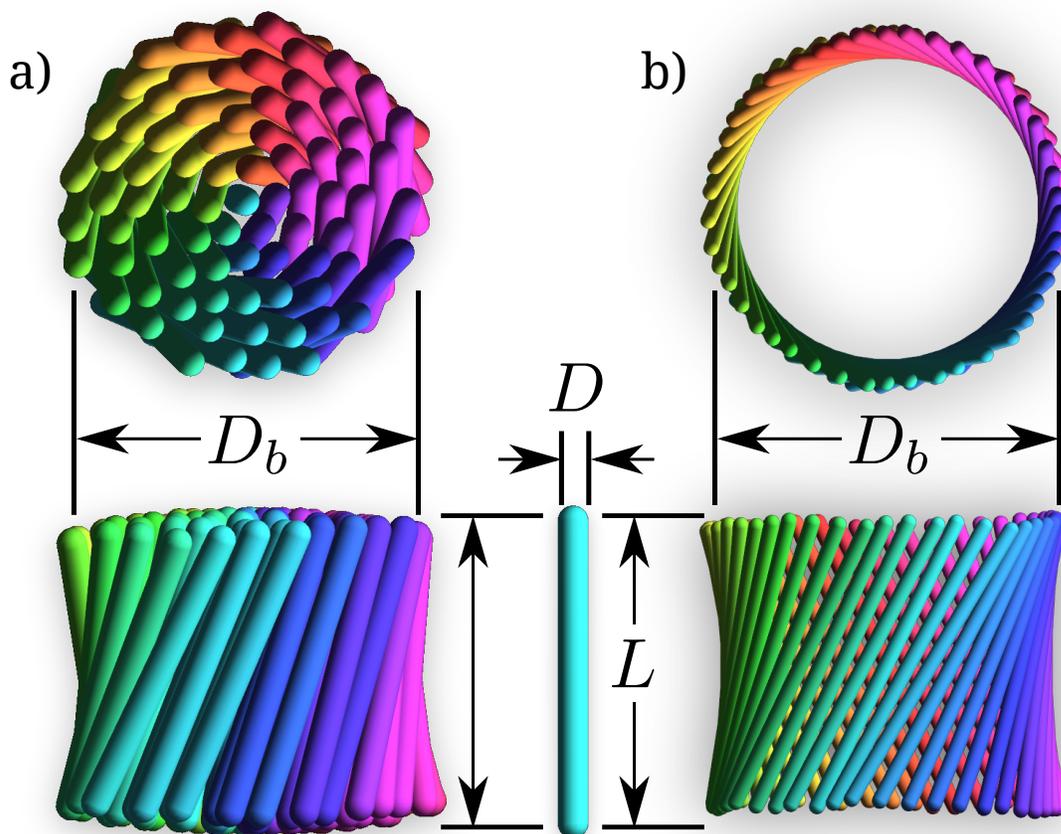
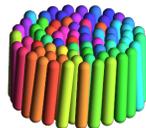


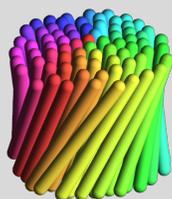
Figure 4.1: A schematic of (a) the explicit bundle and (b) the ring systems. Colors denote the orientation of the tilt direction mapped to a color wheel. The parameters defining a bundle or ring are the aspect ratio of the rod, L/D , the aspect ratio of the bundle, L/D_b , and the density of the bundle (area density in the case of bundles; linear density, rings).

Untwisted

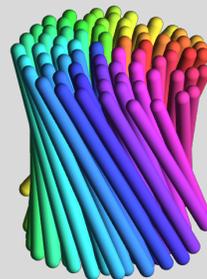


$$\frac{L}{D_b} = \frac{1}{2}$$

Twisted



$$\frac{L}{D_b} = 1$$



$$\frac{L}{D_b} = \frac{3}{2}$$

Figure 4.2: Example snapshots of the explicit bundle system as sampled via NVE Monte Carlo sampling at $\rho = 0.77$ with three different bundle aspect ratios. The system shows the tendency for twist for some system parameters but not others. Bundles with moderate or large bundle aspect ratios tend to display twisting behavior while bundles with small aspect ratios don't.

far apart in a twisted bundle. It seems logical that there is an increase in orientational entropy. In addition, this increase in entropy must be larger than the entropy loss due to the decrease in packing efficiency in order for twist to occur.

We use umbrella sampling to probe this entropy increase, see Section 4.1.1 for details. For this to work we must first define a twist order parameter to act as the sampling coordinate. We measure twist according to \mathcal{T} , see (4.1) which varies from -1.0 to 1.0 for fully left to right handed twist, respectively. \mathcal{T} is the bundle average of individual twist order, \mathcal{T}_i . The individual twist is the sine of θ_c (the rotation of the rod around \hat{r}) or, equivalently, the magnitude of the circumferential component of the rod's orientation vector. This ignores radial components of the rod's orientation and assumes that the center of the twist structure coincides with the center of the bundle.

$$\hat{\mu}_i \equiv \hat{z} \times \hat{r}_i \quad (4.1)$$

$$\mathcal{T}_i \equiv \hat{u}_i \cdot \hat{\mu}_i \quad (4.2)$$

$$= \sin(\theta_{c,i}) \quad (4.3)$$

$$\mathcal{T} \equiv N^{-1} \sum_{i=1}^N \mathcal{T}_i \quad (4.4)$$

If we look at the value of \mathcal{T}_i for individual rods, we see that the twist is largest at the edges of the bundle. At the center of the bundle, the twist goes to zero which is expected due to symmetry.

4.1.1 Umbrella sampling, multiple histogram, and flat histogram methods

In theory, finding the free energy along any coordinate is simple. A sampling run will eventually sample every state in the correct proportion given by the Boltzmann distribution as shown in (4.5). We can simply run the sampling, build up a histogram along some coordinate, x , and then solve for the free energy as in (4.6).

$$\rho(x) = Z^{-1} e^{-\beta G(x)} \quad (4.5)$$

$$G(x) = -k_B T \ln(\rho(\Delta\phi)) + \text{const.} \quad (4.6)$$

This does not work however, as the finite nature of an calculation means that certain low probability states will be sampled either not at all or so poorly that we will have zero knowledge about what is happening there. In order to get this to work, we must use a more advanced technique, umbrella sampling.

Umbrella sampling is a method used to calculate free energy profiles along arbitrary system coordinate pathways. We use this method to calculate the free energy, and thus the entropy, as a function of our twist order parameter. This typically involves defining windows for your system that will limit the range of phase space that the sampling must explore before converging to an accurate representation of the constrained system. These windows can be infinite well, harmonic, or any potential that confines the sampling to a particular area. Multiple histogram techniques are then used to reconstruct the free energy along the pathway defined by the windows. These techniques involve transferring ensemble samples from one Monte Carlo calculation (or molecular dynamics simulation) at a given set of system parameters to another set of system parameters. This is done by correcting the counts in a computed histogram by reweighing each bin according to a bias.[24] In the case of umbrella sampling, this bias is rather explicit in that it is a potential that you apply to force the system into areas of phase space that would otherwise be poorly sampled in any finite calculation. The only real downside to umbrella sampling is that it must be applied with care to your problem. Determining the size of the windows requires some knowledge of the system's behavior. Correctly choosing these window sizes has been addressed in Frenkel and Smit's text where they recommend choosing window sizes that take approximately as long to sample as it takes to sample the phase space perpendicular to the order parameter.[24]

Another technique for the calculation of free energy is the *flat histogram method* by Wang and Landau.[45] Developed for spin systems, this method performs a Monte Carlo sampling procedure but with a dynamically generated bias potential. The bias potential is determined by building up a histogram of visited macrostates in the system and applies the same inversion of that histogram as in (4.6) to determine the free energy landscape of the system. This provides a working approximation of the free energy at any point in the sampling procedure and, as shown by Wang and Landau, this

approximation will converge to the true free energy as the histogram of states becomes “flat”, or invariant to within some tolerance (typically 10%).

Compared to the umbrella sampling approach to free energy calculation, the flat histogram method is considerably easier to use as it requires no manual selection of the window parameters or bias function. Further, as these parameters are completely determined by the natural sampling of the system, they are in some sense optimal for biasing the system into new regions of configuration space.

In spin systems, the total energy of the system is bounded. It is not clear how well the flat histogram method operates in fluid systems, particularly ones in which the histogram cannot be flat due to hard core interactions. In the system at hand, the hard spherocylinder interactions limit the values of \mathcal{T} that can be accessed. At the same time, it would be beneficial to use a method of umbrella sampling that borrows some of positive features of the flat histogram method. We developed such a method that combines the proven umbrella sampling method and the basic dynamic bias potential generation.^[41] This method iteratively performs blocks of Monte Carlo sampling and on each iteration sums the bias reweighted histograms generated so far. This combined histogram is used to compute the working approximation of the free energy which is used to bias the next block of sampling.

The method of Kost-Smith et al. retains the self optimizing nature of the flat histogram method but also resembles the standard windowed umbrella sampling method so closely that most standard tools can be applied. Unlike in the flat histogram method, the block sampling tends to ensure that the computed histogram will never become flat. Each sampling block tends to be short enough that the trajectories do not fully explore the accessible range of the order parameter and thus tend to produce a “lump” of counts on the histogram of the order parameter. This lump means that our approximation of the free energy is inherently inaccurate. This inaccuracy, however, can be beneficial, as it promotes exploration of different areas of the order parameter. Thus, instead of continuing the process until the combined histogram is flat, we continue these iterations until the each bin in the range of interest for the order parameter shows sufficient counts

to support the accuracy we seek. Because we have not converged to the actual bias potential (i.e. the histogram is not flat) we must perform the multiple histogram analysis procedure to determine the final approximation of the free energy profile.[24]

4.2 The entropy of twist

Using the umbrella sampling method, we can directly measure the entropy gain associated with twist. Sampling over the twist order parameter extracts the free energy as a function of \mathcal{T} , see Figure 4.3, and confirms that the twisted configuration is the ground state for certain parameters of the system. In the small aspect ratio bundle, the entropy is maximized at zero twist. At larger aspect ratios, the entropy is maximized at non-zero twist. This analysis shows that twist is more common in bundles composed of high aspect ratio rods, and bundles of low aspect ratio rods tend not to form twisted structures. The results confirm our initial observation of two different bundle regimes. Aggregates of rods that look like a classic bundle, which have a large bundle aspect ratio, tend to be twisted while aggregates that look like membranes, which have a small bundle aspect ratio, are straight.

This calculation gives us the relative likelihood of the twisted versus untwisted state. For instance, in this bundle with aspect ratio one, we see a difference in free energy of $30k_B T$. This means that the twisted state is around 10^{13} times more likely than the untwisted state. And this free energy difference increases with aspect ratio. This is for a modest bundle aspect ratio. In the case of something like a sheaf of wheat with a much larger bundle aspect ratio, this tendency to twist is astronomically large.¹

It should be noted that while the untwisted state that is achieved in the umbrella sampling method is often the straight configuration of the bundle at low bundle aspect ratio, there are other possible configurations that the system can assume which are preferred when the bundle aspect

¹ You might ask what does all of this have to do with a sheaf of wheat or a bundle of pencils? After all, there are no thermal fluctuations in a sheaf of wheat that drives this entropic effect. Entropy, however, is in essence a measure of relative probability. During the process of forming a bundle, there is agitation that gives rise to randomization, where entropy plays a role.

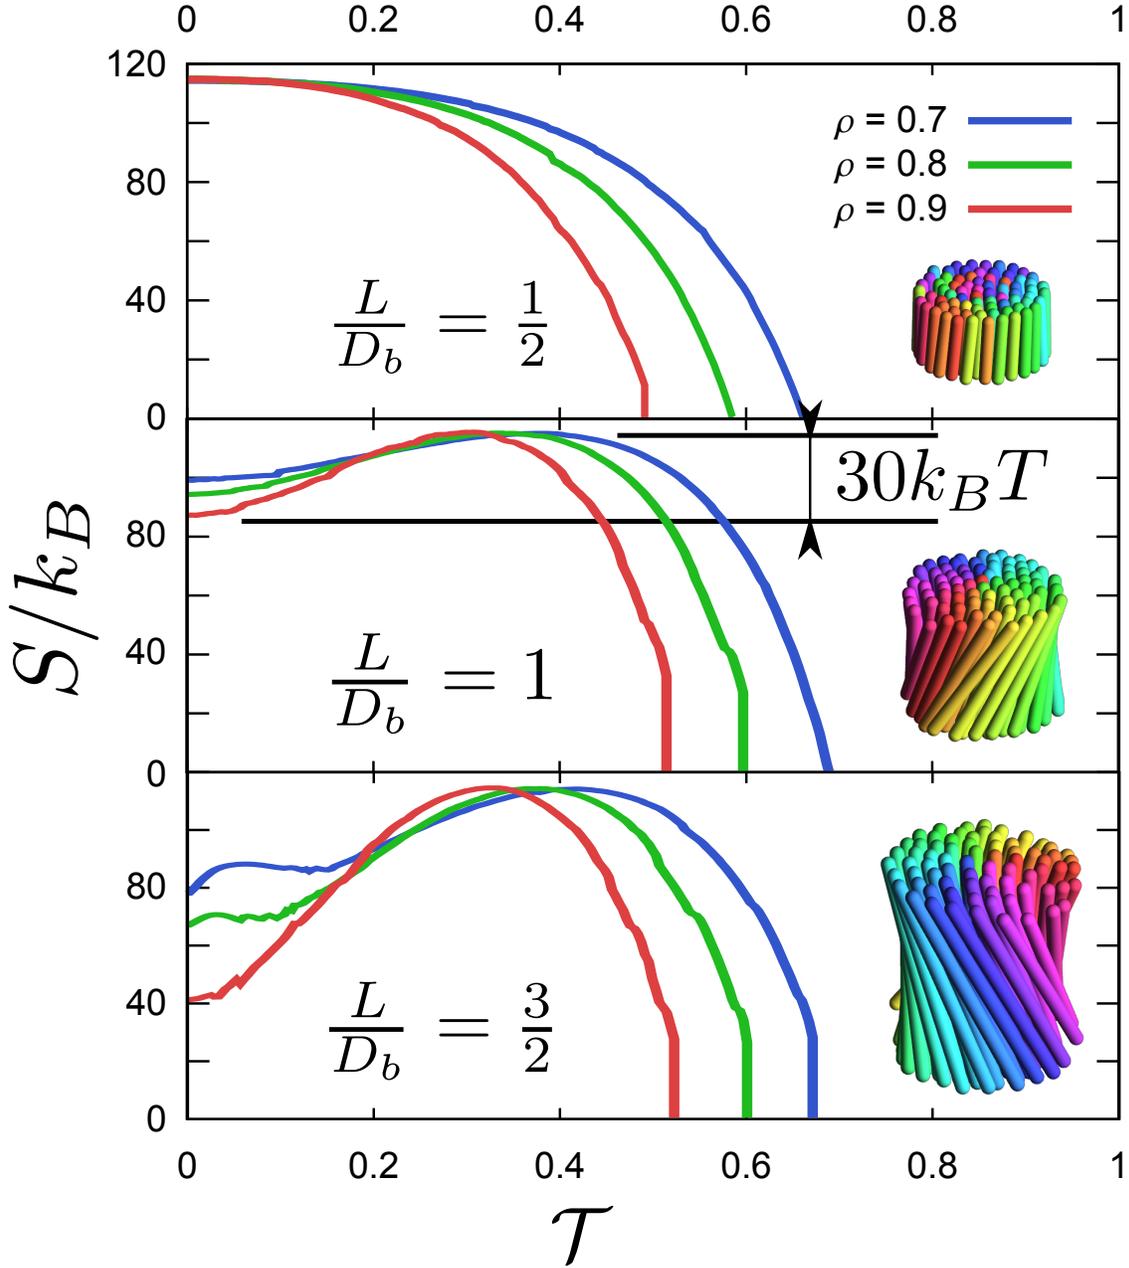


Figure 4.3: Entropy curves as calculated by umbrella sampling for three different bundle aspect ratios, $L/D_b = \{\frac{1}{2}, 1, \frac{3}{2}\}$, each at three different area densities, $\rho = \{0.7, 0.8, 0.9\}$. For small bundle aspect ratio ($L/D_b = \frac{1}{2}$), the entropy of twist is maximized at $\mathcal{T} = 0$. At larger bundle aspect ratio the entropy is maximized at a finite angle whose value is largely a function of only density, as demonstrated by comparing the $L/D_b = 1$ and $\frac{3}{2}$ bundles. The free energy barrier between left and right handed twist is quite large even for moderate bundle aspect ratios, on the order of $30 k_B T$, which indicates that the untwisted state is extremely unlikely in bundle systems. The low density, low \mathcal{T} behavior in the $L/D_b = \frac{3}{2}$ case is due to the system developing multiple twist domains that average to zero twist rather than sampling the straight configuration.

ratio gets significantly larger than one. When looking at the $L/D_b = \frac{3}{2}$ results we note that the entropy is not as low as we would expect in the untwisted state. This is because the untwisted state still has significant individual rod twist, \mathcal{T}_i . The untwisted state in this case is still chiral. This type of behavior is what is known as a *topological rubber glove*.^[21] The transition state between the left and right handed configurations is shown in Figure 4.4.

In this chiral untwisted state there are two domains that have opposite handedness. They often form as concentric or nearly concentric bundles where the inside will have a certain handedness and the outside will have the opposite, causing the twist to average to zero. These states seem to only appear where there is enough free translational volume to accommodate the boundary between the two domains. This is why the $\rho = 0.9$ system shows the normal drop in entropy and has an achiral untwisted state. This should be seen as further evidence of the strength of the entropy driven twist. The system would rather incur the entropic penalties of a domain boundary due to the poor packing that is associated with it than lose the entropy gain of twisting. We now ask, where does this orientational entropy increase come from exactly?

4.3 Identifying the source of twist

As noted, the act of defining an explicit bundle was one of removing degrees of freedom from the system. This simplified the model and made it more amenable to analysis techniques. The assumption that the source of the twisting phenomenon doesn't lie in the removed degrees of freedom is an ansatz that is verified by the continued tendency for twist in the simplified system. This tells us that at least part of the origin of this twist phenomenon lies in the remaining elements of the system and that the removed elements are not essential for the qualitative behavior. We continue this process of simplification by removing degrees of freedom until we isolate the source of the twist inducing entropy by either observing a failed ansatz or by simplifying the model to the point where it can be easily understood.

The first simplification is to remove x and y translations by fixing the centers of the rods to a triangular lattice defining a lattice bundle. This removes the entropy associated with translation.

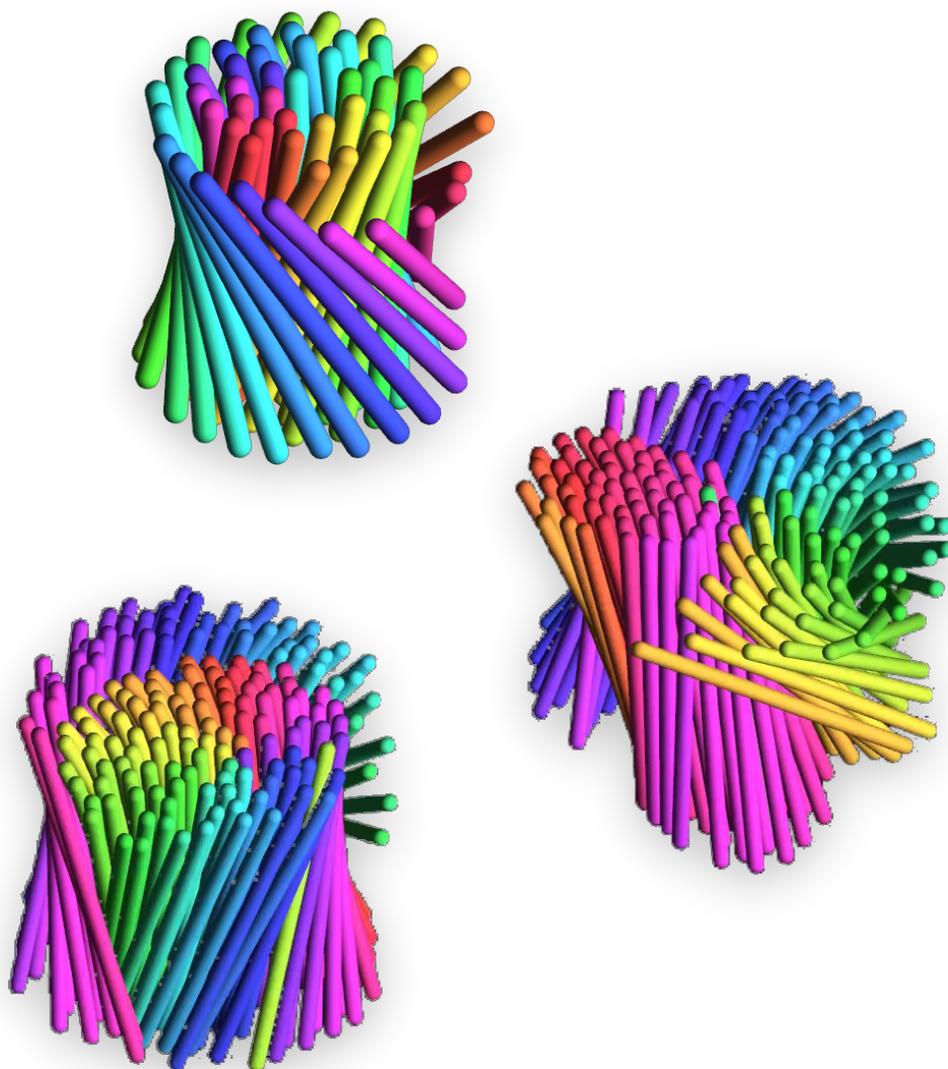


Figure 4.4: Examples of topological rubber glove behavior in the twisted bundle system. These transition states between the left and right handed configurations are induced during the umbrella sampling procedure by biasing the system against twist. The states have zero average twist but remain chiral. These types of configurations are preferred to the achiral transition configuration when twist is highly favored and occur when the density is small enough that there is enough room for the structure and domain boundary to form.

We don't expect for this to qualitatively matter as the translational entropy works against the tendency to twist. This lattice bundle continues to exhibit twist, so we turn to the orientational degrees of freedom. We define our orientation space using a cylindrical coordinate system where the polar rotation is the rotation of the rods toward and away from the center of the bundle (around the $\hat{\phi}$ axis) and the circumferential rotation is the rotation tangential to the edge of the bundle (around the \hat{r} axis). We know that the twist itself is completely defined within the circumferential orientation of the rods. Thus we constrain the rods' rotations to be purely circumferential. Implicit in this constraint is the restriction of the center of the bundle vortex to lie at the center of mass of the bundle.

The continued tendency to twist is evidence that the entropy of the system lies in the circumferential orientations of the rods. At this point we are left with only a single degree of freedom for each rod. Since the rotations in this system are only circumferential, each rod tends to interact primarily with its neighbors at the same radius. Due to this, we can gain some insight on the behavior of the circumferential lattice bundle by approximating it as a set of weakly interacting concentric lattice rings, and limiting our analysis to a single lattice ring. Much as in the bundles, we see that rings with large aspect ratios, L/D_b , spontaneously twist while rings with small aspect ratios do not.

4.4 Cell theory and orientational caging

The lattice ring system is simple enough to get a clear understanding of the entropic contribution of the rods' circumferential orientations. We turn to cell theory to quantify this contribution. In cell theory we approximate the entropy of the system by calculating the free volume that a particle can explore within the *cage* or *cell* formed by its neighbors.[34] In hard core systems, the entropy of the system is intimately related to the amount of free volume each particle can explore in the system. That volume is directly proportional to the number of microstates of the system and thus the entropy should be proportional to the logarithm of the free volume. In this case the volume we are dealing with is an orientational volume.

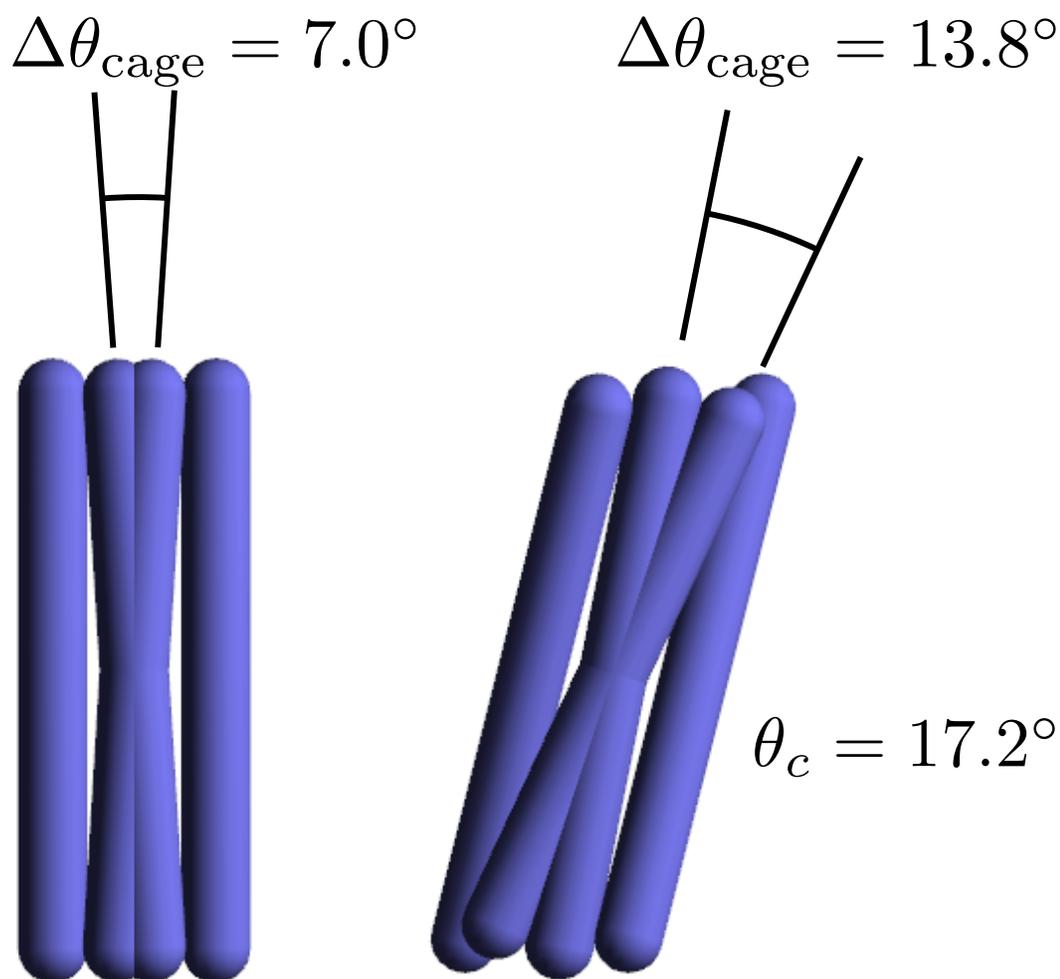


Figure 4.5: Demonstration of the increase of the available angular range as twist is introduced to the ring model. In the straight (untwisted) configuration at this rod aspect ratio, L/D , linear density, ρ , and bundle diameter, D_b/D , the probe rod has 7° of room in its orientational cage. When we increase the twist of the ring to tilt 17.2° , the orientational cage nearly doubles in size. The cage grows because the curvature of the ring allows for circumferential escape allowing the rods to partially miss each other, as is seen on the right. This increase in accessible angular range corresponds to an increase in orientational entropy making the twisted configuration the ground state.

Cell theory is a mean field level theory for approximating that free orientational volume available to each rod. This approximation equates the entropy to the free orientational volume available to a single *probe rod* in a system where the other rods are fixed to twist angle θ_c , as pictured in Figure 4.5. That probe rod is used to find the full angular range available, $\Delta\theta_{\text{cage}}$, which is assumed to be available on average to each of the rods in the ring. In practice, the circumferential range of motion is calculated numerically as a function of θ_c . This involves looking at an orientational cage or cell around the probe rod. We use this to approximate the single rod entropy of the system as $s = s_0 + k_B \ln \Delta\theta_{\text{cage}}$, where s_0 is a constant entropy that we can ignore without loss of generality. We formally describe this process in (4.7 - 4.11) where Z is the N -rod partition function, and z is the single rod partition function. The parameters θ_{max} and θ_{min} are the angles where contact is made between the probe rod and the rods that make up its orientational cage.

$$Z = z^N \tag{4.7}$$

$$z = \alpha \int_{\theta_{\text{min}}(\theta_c)}^{\theta_{\text{max}}(\theta_c)} d\theta_{\text{cage}} \tag{4.8}$$

$$= \alpha \Delta\theta_{\text{cage}} \tag{4.9}$$

$$s = s_0 + k_B \ln \Delta\theta_{\text{cage}} \tag{4.10}$$

$$\theta_{\text{exp}} = \theta_c \text{ that maximizes } \Delta\theta_{\text{cage}} \tag{4.11}$$

Equation 4.10 is the entropy as a function of the twist angle. In the case of pure hard rod systems, the free energy only involves the entropy and thus the free energy is equal to $-sT$. The free energy minimum marks the most likely twist angle for the system, θ_{exp} .² The accuracy of this mean field treatment hinges on the unimportance of correlations in rotation in the twisting behavior. Ultimately, we rely on the umbrella sampling to validate our use of cell theory for this system. As shown in Figure 4.6, the cell theory approximation of the system is remarkably close to the true umbrella sampling based measurement.

² Notably, this is different from the expected mean of the systems twist. These systems are finite and, as such, are not subject to the same simplifying effects of the thermodynamic limit.

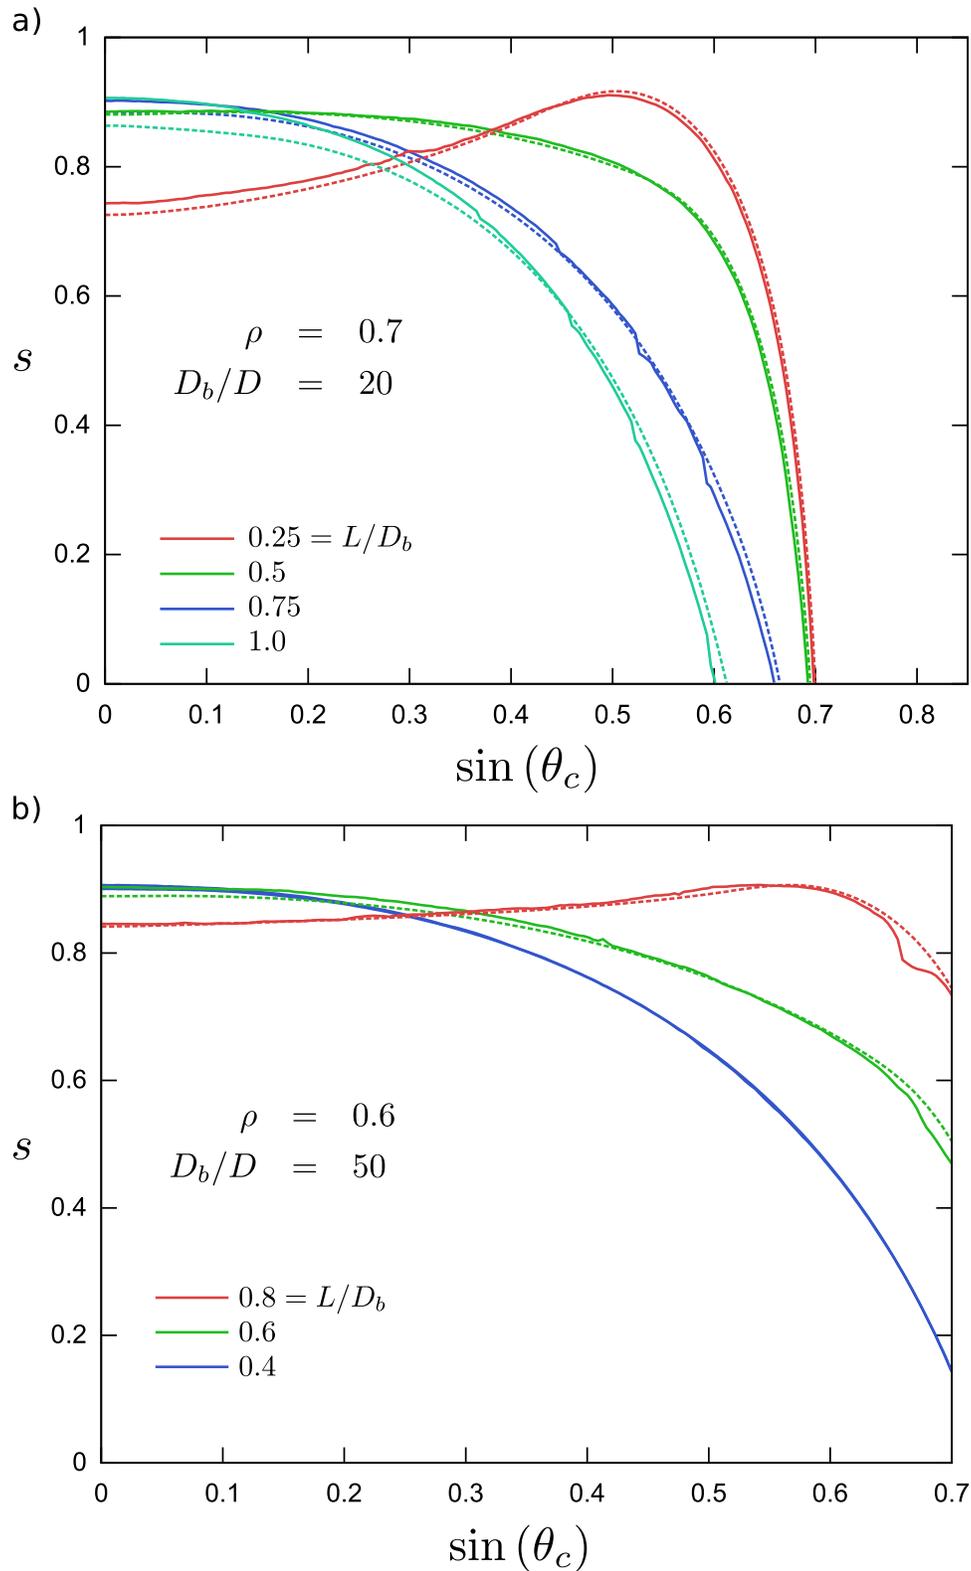


Figure 4.6: A comparison of the umbrella sampling (solid) and cell theory calculations (dashed) of the entropy for a sampling of system parameters. The calculations have no free parameters other than the scaling factor α in (4.8) which is manifested as an additive constant, s_0 in (4.10). This parameter is fit to the umbrella sampling data revealing the accurate nature of the cell theory approximation.

This cell theory result reveals the origin of the twist in the ring system. The twist behavior is a result of the increase in entropy associated with the an increase of the average twist tilt in the ring. This expansion of the orientational cage is a bit surprising. If we think of a linear lattice of molecules, tilting them only reduces the angular cage. The difference between the linear lattice and the ring lattice is the curvature. Curvature allows the rods to ‘get out of the way of’ each other so that they actually increase their orientational volume in a phenomenon we call *orientational escape*. The curvature enables the rods to partially miss each other, as seen in Figure 4.5.

As stated, this doesn’t happen in the linear lattice. There is a transition that happens as a function of curvature. We measure curvature by the diameter of the ring, D_b/D . We also recall that the tendency to twist grows with increasing rod aspect ratio, L/D . The theta range as a function of θ_c is plotted for several different bundle aspect ratios in Figure 4.7. We see that the high twist behavior is the same regardless of the bundle aspect ratio, but the small twist behavior changes dramatically. We see that for $\rho = 0.7$, the bundle transitions from a straight configuration to a twisted one at $L/D_b \approx 0.75$.

We can see that the same principle of orientational escape applies to bundles as well. We can apply a similar cell theory calculation for our lattice bundles with only circumferential rotations. Just as in the ring example, all we need to do is define a cage for the probe rod and find its theta range. This will give us a cell theory representation of a rod at any particular spot in the bundle. Bundles have the added complication in that there are internal rods which have a cage component located at radii larger than their radius relative to the bundle’s center. Figure 4.8 shows the behavior of theta range for rods at the surface of the bundles for several different densities and radii. Figure 4.9 shows the same information but for rods in the interior of the bundle. It is clear that the orientational escape effect still plays a role in bundles even on the interior, though the effect is less there and the transition to preferring untwisted configurations happens at smaller bundle aspect ratios.

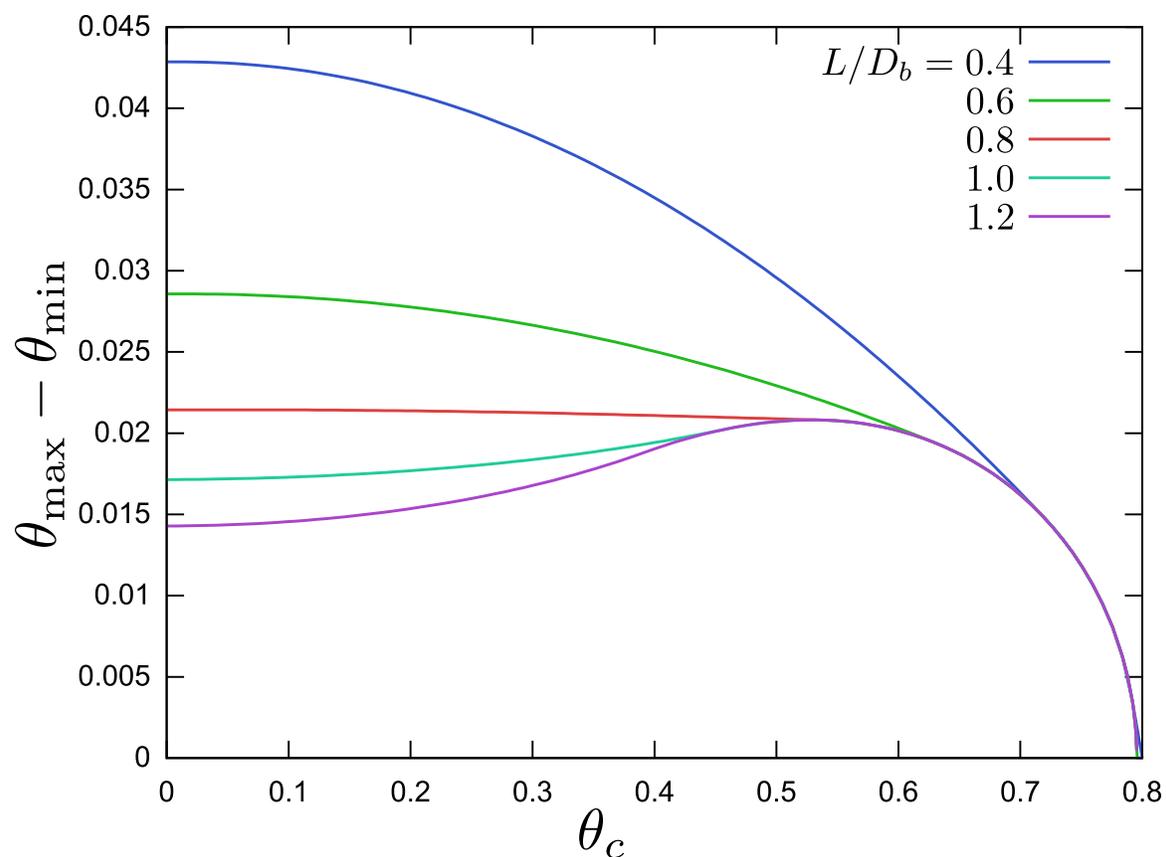


Figure 4.7: The angular range of the orientational cage as a function of a 100 diameter ring's twist tilt, θ_c , for several different bundle aspect ratios, L/D_b . We see a clear transition from a twisted state to the straight configuration near $L/D_b = 0.8$.

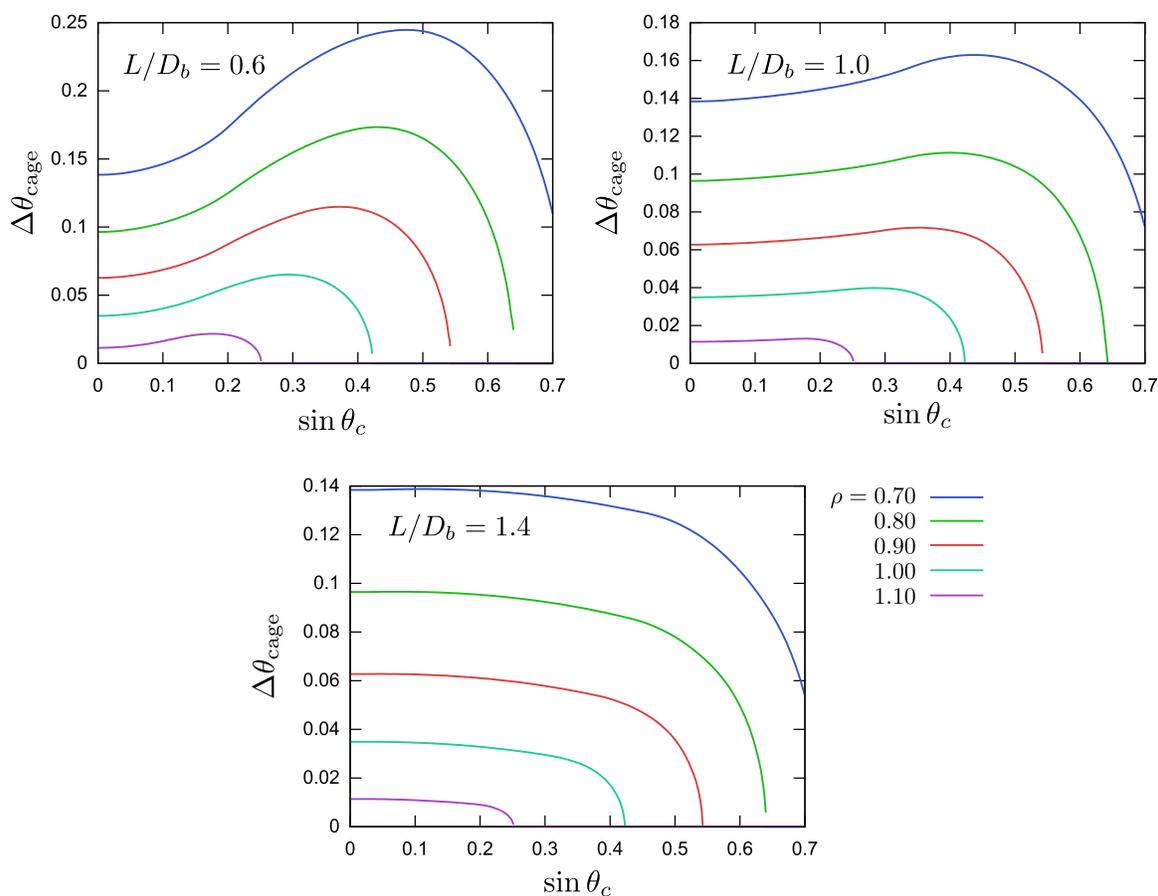


Figure 4.8: The angular range at the edge of bundle of length 10 spherocylinders as a function of the bundle's twist tilt, θ_c , for several different bundle aspect ratios, L/D_b , and densities, ρ . We see a clear transition from a local tendency for twisted behavior to a local tendency for untwisted behavior as a function of the bundle aspect ratio, L/D_b .

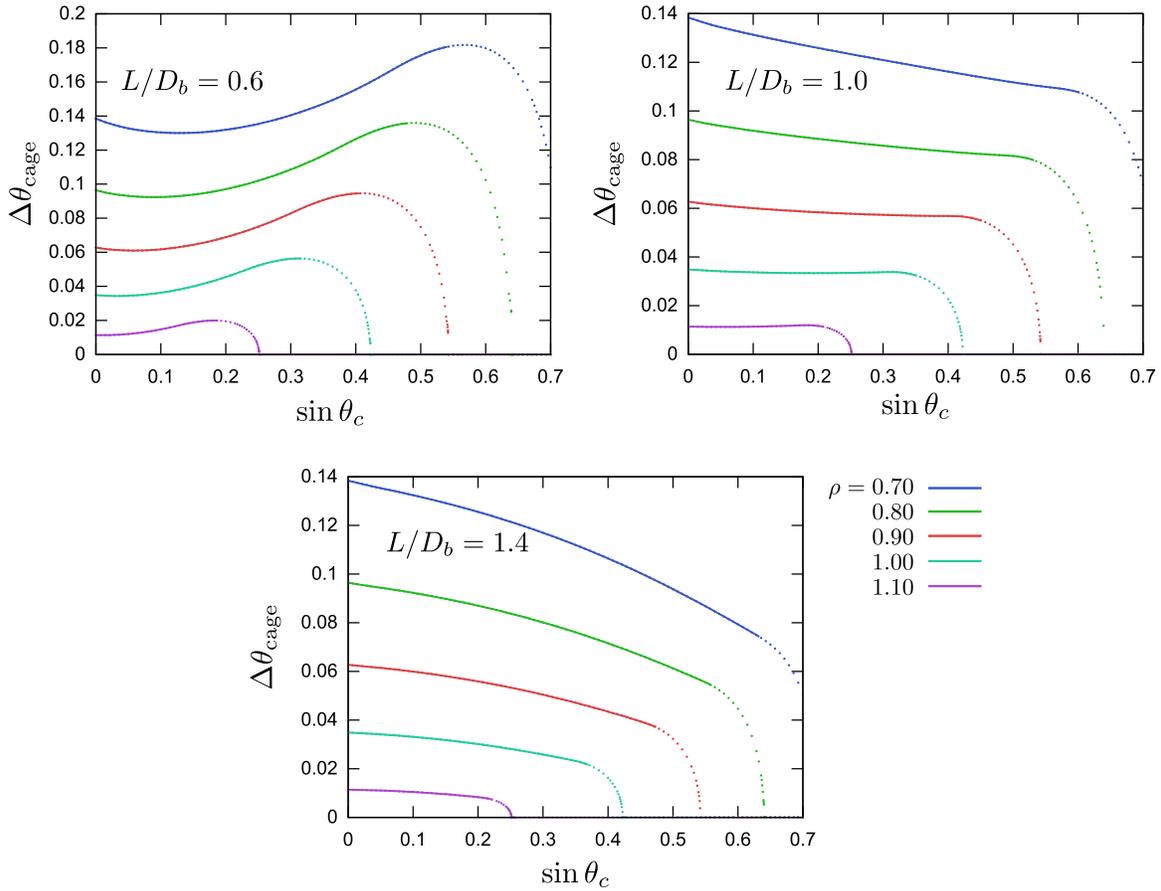


Figure 4.9: The orientational cage of the rods in a bundle as a function of the bundle's twist tilt, θ_c , at several densities, ρ , and radial locations within the bundle. This is like Figure 4.8 except that the rods' cages also involve rods at larger radii than itself. The dotted lines indicate unphysical configurations where the cage rods would intersect other cage rods in order to achieve the desired θ_c . The transition as a function of bundle aspect ratio, L/D_b , is less clear here, but the expansion of the interior cage at small radius is still apparent.

4.5 Phase behavior

Using the cell theory for the lattice ring system, we derive the phase diagram as a function of the ring aspect ratio, L/D_b , and the packing fraction on the ring in the large ring diameter, D_b/D , limit, see Figure 4.10. When the ring aspect ratio is moderate to large, the phase behavior simplifies to depending on only these two parameters. We see two clear phases, one with high twist and one with no twist. The transition between them is discontinuous. In addition to these two phases, we see an intermediate phase with weak twist. This phase exhibits a small but finite entropy maximizing twist, but the handedness of these rings tends to flip readily in Monte Carlo simulations. This effect, however, becomes much less important in the limit of large ring diameters. When dealing with smaller diameters, the phase behavior becomes more complex. This complexity is presumably due to some effect of the shape of the excluded volume of the rod. The nature of this behavior is still an unanswered question, but we suspect that the effect here is related to interaction between ends of the rods. However, even in the more complex cases, the ring aspect ratio and density are still the predominate deciding factors in whether a ring will twist.³

The very fact that the aspect ratio is the important parameter for deciding twist tells us something about the mechanisms at play. This parameter contains two competing terms. The rod length, L , is a measure of the affect a rotation of the rod will have. The size of the displacement of the end of the rod, and thus the probable points of contact, scales as approximately $L \times \Delta\theta$. The reciprocal ring diameter, D_b^{-1} , is a measure of the curvature. The importance of the aspect ratio encapsulates the idea of orientational escape, the increase in the motion of the point of contact is canceled by the curvature at the surface.

The phase transition between twisted and untwisted states is discontinuous as seen in Figure 4.11. In addition, the twist of the ring does not change with ring aspect ratio once the transition occurs. This feature is easily understood by observing that in the maximal entropy state, rods only interact via side-to-side interaction where the length of the rods don't matter. In the high packing

³ This trend holds down to $D_b/D \sim 20$ at which point the trend is still present, but the ring's weakly twisted phase behavior is quite prominent.

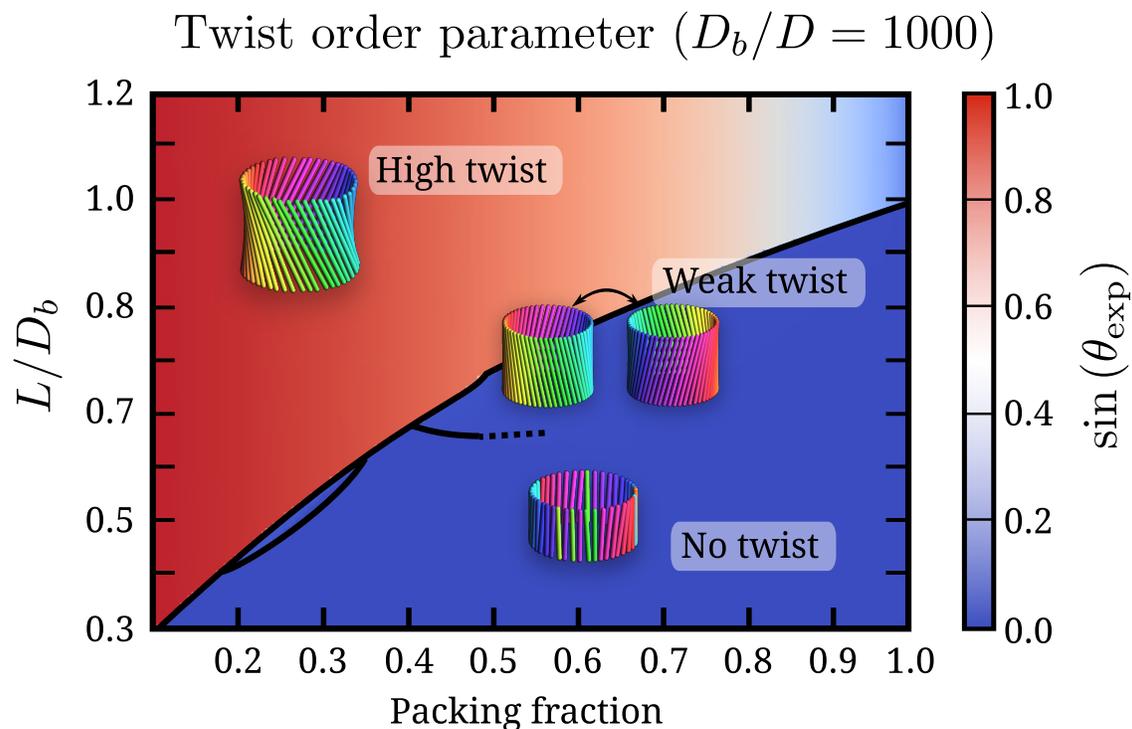


Figure 4.10: Phase diagram for the ring system of density, ρ , versus bundle/ring aspect ratio, L/D_b for large bundle diameters, D_b/D . There is a clear transition between the twisted and untwisted states as a function of aspect ratio. In the close packed limit, the transition aspect ratio is one. There is also a weakly twisted phase that shows up in pockets near the twisted to untwisted phase transition. This phase easily switches its handedness in finite temperature Monte Carlo simulations. The effect of the weakly twisted phase appears to lessen in the large D_b/D limit.

fraction side of the diagram, we note that the entropy maximizing twist smoothly approaches zero as the orientational cage’s volume is reduced to zero. In addition, the transition point between the two phases approaches the aspect ratio of $L/D_b = 1$. For lower packing fraction rings, the transition aspect ratio is lower. The extent to which we can explore low packing fraction or bundle aspect ratio is fundamentally limited by the model.⁴

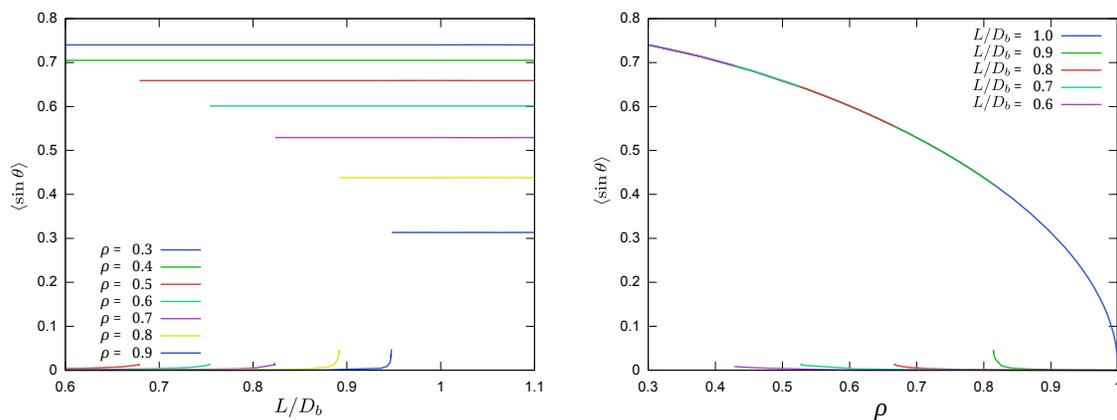


Figure 4.11: The entropy maximizing twist versus ring aspect ratio L/D_b for several densities, ρ , and versus density for several aspect ratios in the large D/D_b limit. There are scans of the phase diagram in Figure 4.10. Aside from the transition, the entropy maximizing twist doesn’t vary much with bundle aspect ratio and the twisted state curves collapse upon each other in the large D/D_b limit. The transition is clearly discontinuous.

4.6 Conclusion

This twist driven by orientational escape is a new mechanism for twist ordering based on orientational entropy. It is seemingly distinct from the behavior described by Yang et al. but it shows intriguingly similar results. It is also consistent with an overall trend of results showing twisted structures as the ground state in simple otherwise achiral systems.

This phenomenon is very common and something that is present in macroscopic systems.

⁴ At some point, the very concept of a “lattice ring” is severely degraded. At minimum, this leads to completely different modes of interaction with between neighbors, such as the *top* of a rod interacting with the *bottom* of a neighbor, which don’t help us understand the twisted bundle phenomenon we sought to explore as it is qualitatively different. At low enough ring aspect ratio or packing fraction, the rods cannot physically intersect with their neighbors at all. While worthy of further exploration, it is not clear that these cases are physical. For instance, it is not clear how these structures would be stable as bundles with reasonable interactions.

This is the reason that a bundle of pencils will develop a twist if shaken or a natural way that a sheaf of wheat can develop a twist due to agitation during the bundling process. Nature is full of twisted or otherwise chiral structures and this phenomenon could be a significant contribution to the explanation of these behaviors. However, due to its counterintuitive nature, this route to spontaneous chirality is perhaps overlooked.

There are several interesting directions to take this research. We defined a series of systems with varying numbers of degrees of freedom some of which are explored here. Other systems of interest include annuli of rods where there is an interior radii that the rods cannot move inside and nematic tubes where the constraint to the $z = 0$ plane is removed.

Discussions with Zvonimir Dogic have raised the question of coupling between twist and curvature at the surface of the bundle. Since free energy is reduced by twist when the curvature is above a threshold, it stands to reason that there is a coupling between the twist and the shape of the interface. If the surface of the bundle is free to move, the system should show some tendency to develop higher curvature if it will produce twist. Monte Carlo calculations of the interface will help to clarify the nature of this coupling. This could help to explain some experimental results that the Dogic group sees in the fluctuations of the interface.

The next steps in this project are to finalize the analysis of the system. The nature of the weakly twisted phase in lattice ring systems needs to be further explored. If the weakly twisted phase is indeed related to the shape of the excluded volume of the rods, use of models of rods other than spherocylinders should shed some light on this phenomenon. Initial investigations with cylinders showed little difference, but exploring the phase diagram of ellipsoids will most likely be fruitful.

When considering self-assembled bundle systems, it is clear that the distributed attraction plays an important role. While our explicit bundle systems captured the entropic twisting phenomenon, in systems with attraction, this effect is partially suppressed as the entropic gain due to the twist is in direct competition with the associated energetic cost. We are looking into how we can reasonably include the distributed attraction in the cell theory of bundles. More directly, we are

returning to the systems that prompted this analysis in the first place, the self-assembled bundle systems. Both the systems with explicit distributed interactions via the square well line potential and the ghost particle system where the attraction is modeled via an explicit depletant, serve as good systems to probe the twisting phenomenon in self-assembled bundles. With an appropriately generalized twist order parameter, we can perform the same calculation of free energy via umbrella sampling.

These self-assembled systems are much more computationally intensive. In the case of the square well line attraction the pair potential calculations become expensive. For the ghost particle method the systems must be much larger and involve significant size asymmetry for the implicit depletion interaction to work properly. The computational efficiency of these systems is much more of a concern. These systems may be suitable candidates for the techniques examined in Chapter 3. For the square well line attraction, it is quite possible that the VMMC procedure could provide some benefit. In the case of the ghost particle system, there is a set of large rod particles that interact with much smaller sphere depletants. This is precisely the type of simulation where the GenGCA algorithm, the Swendsen-Wang style cluster move using the inversion symmetry operation, is known to excel, see Section 3.3. This is likely a very effective method for this kind of calculation and much more suitable than Metropolis.

Chapter 5

de Vries liquid crystals

In 1977, de Vries reported a new type of liquid crystal SmA phase that showed a set of qualitatively different features from the conventional SmA phase.[15, 17, 16] Initially this phase was distinguished from other smectics by a first order SmA-SmC phase transition with an anomalously small reduction in layer spacing.[43] Subsequent work showed the characteristics of the de Vries SmA phase also includes large electroclinic responses, large field driven changes in birefringence in chiral materials, and a lack of a nematic phase in the phase sequence.[39, 65] One of de Vries' early models, now called the hollow cone model, proposed that these phases are smectics with a molecular orientational distribution which is uniform in azimuthal angle ϕ but has a preferred tilt away from the layer normal, θ_A , sweeping out the surface of a cone, Figure 5.1. While this model remains controversial, it neatly explains the electro-optic response of chiral de Vries systems as well as the small layer contraction as the effects of averaging the molecular properties over azimuthal angle.[11]

Although many of the features of the de Vries SmA phase can be explained by applying the hollow cone model, the observations of a first order SmA-SmC transition remain unexplained despite substantial study of these systems.[15, 39] In this chapter we show using simulations and mean field theory that, much like the anomalous layer spacing and electro-optic response, the first order SmA-SmC phase transition seen in de Vries smectics can be understood as a direct result of the hollow cone model.

Materials exhibiting de Vries phases sparked interest for use in ferroelectric liquid crystal

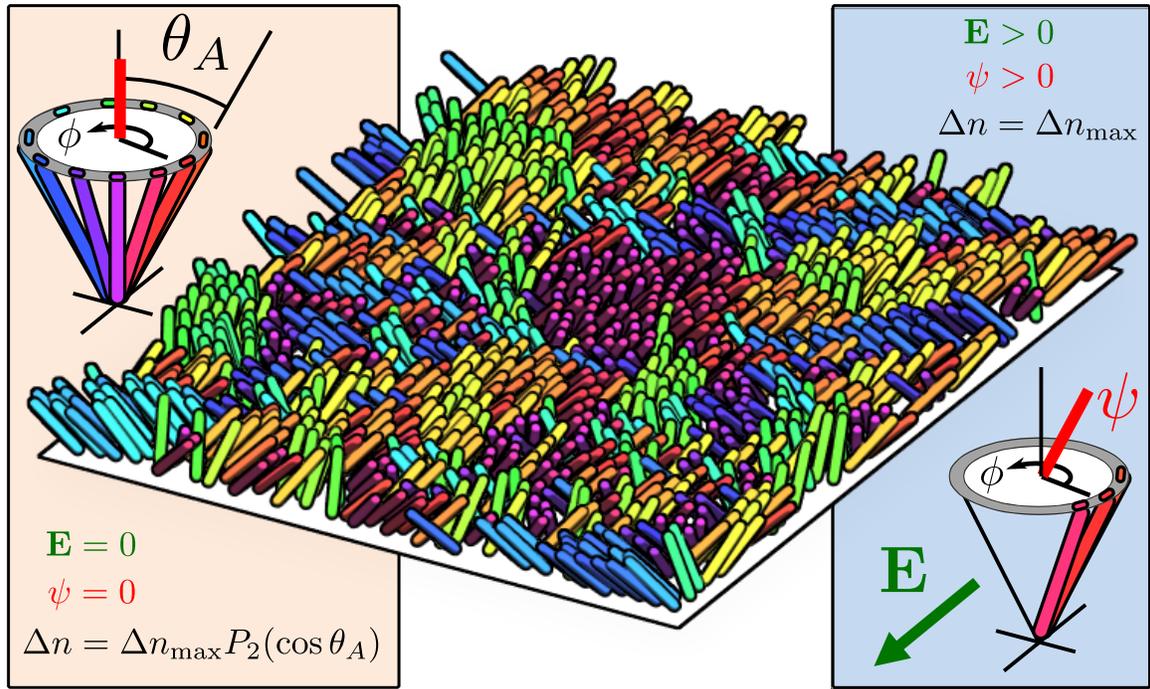


Figure 5.1: A snapshot from a Monte Carlo simulation of a single layer hard spherocylinder smectic with the hollow cone orientation distribution (color represents azimuthal orientation). This NPT simulation was performed with cone angle $\theta_A = 45^\circ$ and spherocylinder length $L/D = 5$ in the SmA phase near the SmA-SmC phase transition. Molecules in the hollow cone model have the orientation distribution in tilt, θ , and azimuthal angle, ϕ , of the form $f(\theta, \phi) = (2\pi)^{-1}g(\theta)$. The distribution is uniform in ϕ while $g(\theta)$ is narrowly peaked around the cone angle, θ_A , which we take as the delta function, $\delta(\theta - \theta_A)$. Hollow cone smectics show no global polar order in the SmA phase, but produce finite correlated regions. Macroscopic tilt, ψ , is the result of global bias of molecules to one side of the ϕ distribution.

displays [71]. The small change in layer spacing makes these good candidates for ferroelectric liquid crystal (FLC) displays where large changes to that spacing can cause the formation of “zig-zag” defects during manufacturing [68]. In addition, the large electroclinic effect of de Vries SmA phases is well suited for sensitive chirality detection [39]. These potential applications have led to substantial research attempting to characterize and understand de Vries behavior which in turn has led to empirical exploration of the properties of de Vries materials. For instance, liquid crystal chemists discovered that chemical and structural motifs that promote layering, such as polyphilic or bulky tails, tend to produce de Vries-like behavior. This has given us some degree of predictive power in the design of de Vries mesogens [61]. The correlation between strong layering and de Vries materials also manifests as a direct isotropic to SmA phase transition, bypassing the nematic phase entirely. Meanwhile, high-resolution x-ray studies of the SmA-SmC phase transition have directly observed a discontinuous change in layer spacing in several materials [66, 39]. Exploration of the electro-optic behavior reveals a distinct sigmoidal polarization response, or double peaked polarization current response [4, 11, 73, 58, 25] in a broad range of de Vries materials, also indicative of a first order SmA-SmC phase transition. This first order phase transition is in sharp contrast to the second order SmA-SmC transition common in conventional smectics [50, 15].

Several groups have performed theoretical studies of the de Vries SmA phase and the SmA-SmC transition at a phenomenological level. Bahr et al. used a simple Landau theory to model de Vries like electroclinic response seen in material C7 [4]. The mean field theory developed by Saunders et al. shows that coupling between tilt and biaxiality in smectics can produce a first order SmA-SmC phase transition, but doesn't point to a microscopic origin for this coupling [63, 62]. In a similar vein, work by Gorkunov et al. shows that the addition of higher order coupling terms into a mean field theory recreates the anomalously small change in layer spacing through the SmA-SmC phase transition, and shows how these couplings might arise based on a model of intermolecular interaction [29, 30]. At the microscopic level, Lagerwall et al. have proposed an alternate to the hollow cone model based on a conventional SmA with abnormally low nematic order [44].

Our work starts from the microscopic foundation of the hollow cone model and explores the

implications of this model for de Vries smectics using Monte Carlo simulations, umbrella sampling, coarse graining, and mean field techniques. This bottom-up analysis culminates in the realization that the first order SmA-SmC transition can be viewed as a consequence of the hollow cone model via a defect condensation mechanism.

5.1 The Implications of the Hollow Cone Model for de Vries SmA

We begin by considering the behavior of a hollow cone fluid of a single smectic layer of hard spherocylinders. We performed NPT Monte Carlo simulations of a fluid of spherocylinders which are confined to the $z = 0$ plane and tilted from the layer normal by fixed cone angle θ_A , but allowed to freely rotate in azimuthal angle ϕ . This rigid realization of the hollow cone model where the distribution of spherocylinders in z and θ_A are explicitly delta functions is chosen to reduce the number of free parameters in the model. A more realistic model would include out-of-layer fluctuations and a finite distribution in angle θ centered around the cone angle, but such generalizations should not change the qualitative conclusions of the analysis.

Simulations of this idealized system at sufficiently large cone angles show a first order phase transition between a quasi-long ranged SmC-like phase at high pressure and a disordered phase at low pressure where finite ϕ -correlated domains average to SmA symmetry as demonstrated in Figure 5.1. These domains occur in hollow cone smectics due to the effective potential experienced by tilted spherocylinders interacting with their neighbors. The existence of these correlated domains predicts very different behavior from conventional SmA phases. Specifically, whereas fields on conventional SmA phases act on single molecules resulting in weak electroclinic coupling, the electroclinic effect due to reorienting the ϕ value of correlated domains around the cone is much larger. This is particularly noticeable near the SmA-SmC phase transition, which is consistent with experimental measurements of de Vries systems [39].

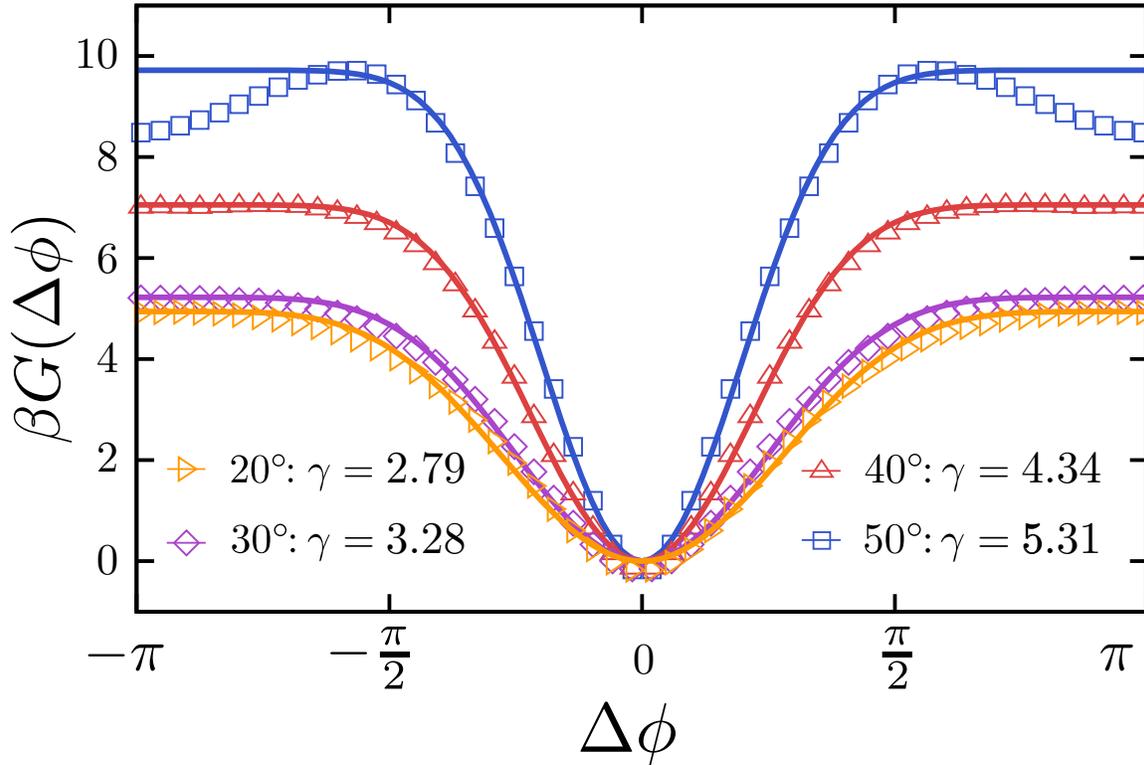


Figure 5.2: Umbrella sampling calculations of the Gibbs free energy of azimuthal rotation within a smectic layer of a single spherocylinder relative to the polarization vector averaged among surrounding spherocylinders within radius $a = 5\sigma$ where σ is the spherocylinder diameter. These calculations are performed within the SmC phase near the SmA-SmC phase transition at different cone angles, θ_A . At larger cone angles a secondary minimum appears at $\phi = \pi$. Superimposed are fits of the γ -potential, see (5.1), where the secondary minimum is excluded from the fitting procedure if present by separately fitting G_{\max} as demonstrated in the 50° fit. As the cone angle is increased the width of the well is reduced, or equivalently, γ increases.

5.2 Coarse Graining Transformation

We measure the effective azimuthal potential by computing the Gibbs free energy cost of rotating a single spherocylinder by angle $\Delta\phi$ away from the local azimuthal order averaged over a domain of radius a in the fluid layer, $G(\Delta\phi)$. An umbrella sampling scheme ensures proper sampling over the full range of the rotational potential. We make iterative approximations to $G(\Delta\phi)$ in order to bias the system away from well sampled regions. Each iteration improves the approximation by using the sampled histogram, $\rho^{(i)}(\Delta\phi)$, via the equation

$$G^{(i)}(\Delta\phi) = -k_B T \ln(\rho^{(i-1)}(\Delta\phi)) \quad ,$$

where the superscripts denote the iteration. The free energy profile, demonstrated in Figure 5.2 for various cone angles in the SmC phase near the SmA-SmC transition, shows a narrow minimum at $\Delta\phi = 0$ whose width decreases with increasing cone angle. This result is robust to changes in the hollow cone θ -distribution. For instance, defining a diffuse hollow cone via a harmonic potential in the θ coordinate produces qualitatively similar results.

Using the free energy profile and our observation that this system is well represented as a collection of ϕ -correlated domains, we define a coarse graining transformation which maps our off-lattice hollow cone smectic onto a generalized XY spin model with lattice constant a , the coarse graining length-scale. The average ϕ coordinate of the correlated domains within the smectic layer maps to the spin orientations in the XY model and the free energy profile maps to the in-layer nearest neighbor interaction. Motivated by this measured azimuthal free energy curve, we choose an in-layer potential of the form

$$U_{ij} = -J_{xy} \left[\frac{1}{2}(\cos(\phi_i - \phi_j) + 1) \right]^\gamma \quad . \quad (5.1)$$

This potential, which we call the γ -potential, features an energy minimum at $\Delta\phi = 0$ whose width is tuned via parameter γ . The γ -potential fits well to the free energy profile at small to moderate cone angles near the SmA-SmC phase transition as shown in Figure 5.2. At larger cone angles we observe a secondary minimum at $\Delta\phi = \pi$ which we choose to exclude from the fitting

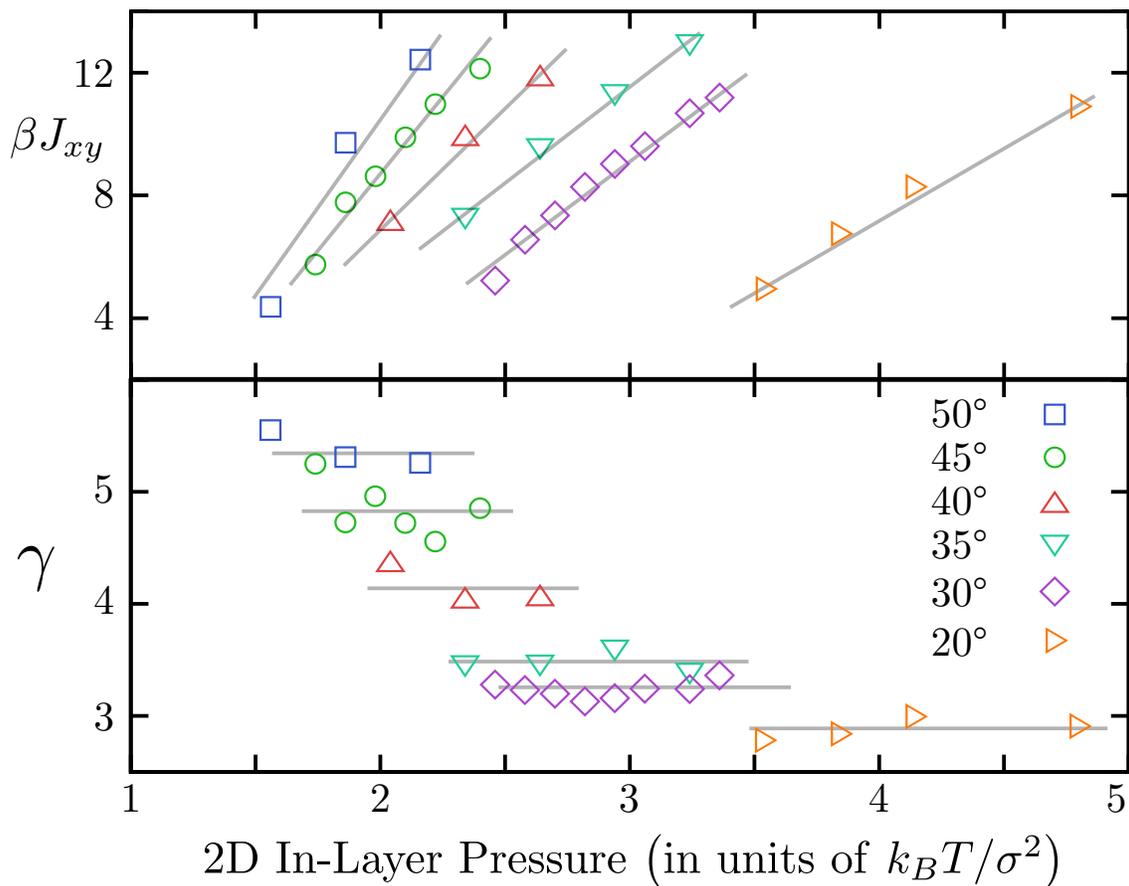


Figure 5.3: The mapping between the hard spherocylinder hollow cone system and the γ -potential XY model when using $a = 5\sigma$ as the coarse graining length scale where σ is the diameter of the spherocylinders. Spin coupling strength, J_{xy} , increases approximately linearly with pressure while the γ values are nearly independent of pressure but increases with larger cone angles, θ_A . While the specific values of J_{xy} and γ depend on a , these trends do not.

procedure as our analysis suggests that the width of the minimum at $\Delta\phi = 0$ is the essential feature of the interaction. At sufficiently large cone angles, however, we predict that de Vries systems might additionally exhibit a smectic phase with two-dimensional nematic order of the c -director analogous to the phase reported in spin systems with a secondary minimum [53].

These fits provide us with a map, summarized in Figure 5.3, between our hard spherocylinder system governed by pressure, spherocylinder length, and cone angle to a spin system of lattice constant a governed by the angular width of the minimum in the neighbor interaction potential (encoded in γ) and the dimensionless energy ratio βJ_{xy} . With respect to pressure, the map reveals a roughly linear increase of the coupling strength, J_{xy} , and approximately constant γ independent of cone angle. We note that while the energy ratio βJ_{xy} corresponding to the transition shows no real trend with respect to cone angle, the value of γ shows a distinct increase as the angle increases.

5.3 Modeling

The γ -potential can be viewed as a continuous version of the Potts model interaction, $u(\phi_i, \phi_j) = -J\delta_{\phi_i, \phi_j}$ where δ_{ϕ_i, ϕ_j} is the Kronecker delta over possible discrete ϕ states. Previous work by Domany et al. on two-dimensional spin systems used the γ -potential to explore the cross-over between the continuous Kosterlitz-Thouless phase transition of the planar rotor XY model with interaction $u(\phi_i, \phi_j) = -J \cos(\phi_i - \phi_j)$ at $\gamma = 1$ and the first order phase transition of the n -state standard Potts model where $n > 4$ [18].

Our de Vries system is composed of smectic layers which, like the two dimensional spin systems of Domany et al., are represented as planes of spin coupled by the γ -potential of strength J_{xy} in the x and y directions. The spins in these layers interact with adjacent layers in the z direction via a planar rotor style coupling of strength J_z . The Hamiltonian for this anisotropic cubic XY model is given by (5.2) where we have assumed integer γ , expressed the γ -potential as a finite Fourier sum with known coefficients b_k , and sum the last term over all in-layer neighbors.

$$H = -pE \sum_i \cos(\phi_i) - J_z \sum_i \cos(\phi_i - \phi_{i+1}) - J_{xy} \sum_{(i,j)} \sum_{k=1}^{\gamma} b_k \cos(k(\phi_i - \phi_j)) \quad (5.2)$$

We characterize the phase behavior of our related system using self-consistent variational mean field theory.[9] In the mean field approximation we absorb the lattice's geometric factors into our coupling constants and combine the planar rotor inter-layer coupling with the first term of the in-layer coupling. We assume that the N -spin matrix can be written as the product of single-spin density matrices, (5.3 - 5.4). Sets of order parameters, c_k , are found that satisfy the γ different self-consistency constraints in (5.5). The stable phase is parametrized by the set of order parameters minimizing the mean free energy, see (5.6).

$$\rho_1 = \frac{1}{Z_1} \exp[\beta(pE + \sum_k J_k c_k \cos(k\phi))] \quad (5.3)$$

$$Z_1 = \int_0^{2\pi} \exp[\beta(pE + \sum_k J_k c_k \cos(k\phi))] d\phi \quad (5.4)$$

$$c_k = \int_0^{2\pi} \rho(\phi) \cos(k\phi) d\phi \quad (5.5)$$

$$F = \frac{1}{2}N \sum_k J_k c_k^2 - \frac{N}{\beta} \ln(Z_1) \quad (5.6)$$

The system shows a transition between the paramagnetic (SmA) phase and ferromagnetic (SmC) phase as shown in the $\gamma - T$ phase diagram in Figure 5.4. Much like Domany et al. we see a continuous phase transition at small γ and a first order transition at large γ . [18] Comparing with our parameter mapping, systems corresponding to smectics with large cone angles show first order SmA-SmC phase transitions while values of γ corresponding to small cone angles show continuous behavior. Our analysis locates the tricritical point in the two-dimensional system between the γ values of 4 and 5 which corresponds to a cone angle of approximately 40° , though this is dependent on the coarse graining length scale. The phase behavior proves to be robust to weak inter-layer planar rotor interactions. To first order, the inter-layer coupling stabilizes the SmC phase to higher temperatures but doesn't significantly change the location of the tricritical point.

The first order phase transition within the spin system, and thus the spherocylinder system, can be understood in the context of defect or vacancy condensation much like phase transition in the standard Potts model [55]. The system free energy is reduced by overlapping disordered regions, producing a depletion-style attraction between disordered domains. This suggests that the

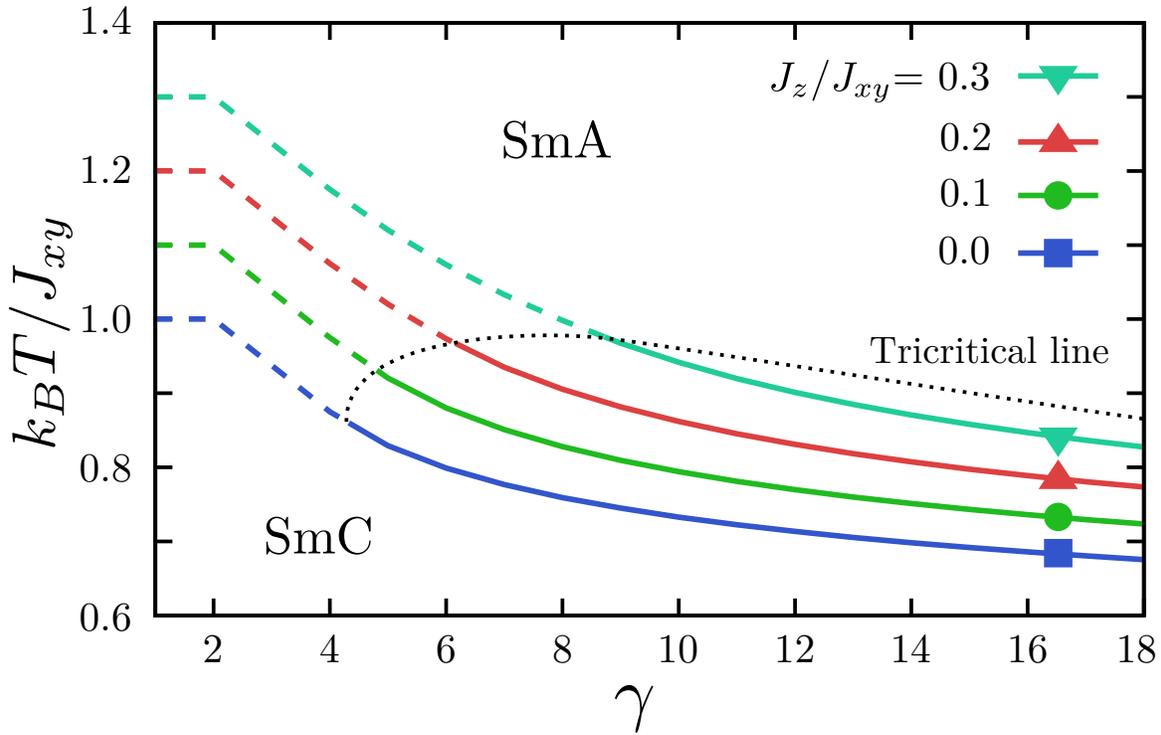


Figure 5.4: The $\gamma - T$ phase diagram of the γ -potential spin system showing tricritical behavior. The dashed and solid phase boundaries denote second and first order phase transitions, respectively. In the two dimensional system (■) and systems with weak inter-layer coupling, $J_z/J_{xy} < 0.1$, the phase transition crosses over from continuous to first order between $\gamma = 4$ and $\gamma = 5$. At stronger couplings, we see a sharp shift to larger values of γ as the cumulative potential begins to resemble the planar rotor.

first order phase transition seen in de Vries smectics also originates from a disorder condensation mechanism.

5.4 Experimental measurements

Experimental measurements of the field dependent polarization, the electroclinic effect, and the field dependence of the birefringence were performed by various other members of the Liquid Crystal Materials Research center.[73, 66] These measurements were done on liquid crystal cells in the common bookshelf cell geometry where smectic layers are perpendicular to the plane of the cell (see Figure 5.5). In these cells, we apply fields along the cell's normal which is in the plane of the smectic layers. Applying this field removes the rotational symmetry, giving rise to the electroclinic effect, or tilting of the optical axis as an electric field is applied. The field dependence of the birefringence, optical tilt, and polarization was collected for W530 and W599 (see Figures 5.6 and 5.8) when subject to in-plane electric fields. Both materials are chiral and exhibit de Vries smectic A phases. In particular we will focus on the behavior in the SmA phase near the SmA to SmC transition. In this region we see behavior seemingly common to chiral de Vries materials, which provides insight into the microscopic organization of de Vries materials in general.

Polarization data is achieved by measuring the induced polarization current from the field applied as a triangle wave as seen in Figure 5.7. This signal has roughly three background components: a capacitive component that contributes a jump when the direction of the triangle wave changes, a resistive component which contributes a linear baseline, and an ion current that contributes a peak at large fields. This background needs to be removed before comparison to any predictions from our model. The background removal process is done by hand and is relatively sensitive to subjective factors like identifying and removing the ion peak when determining the linear baseline. The inset of Figure 5.7 is an example of the data once the background has been removed.

By the optical tilt and birefringence and polarization shown in Figure 5.8, we see that W530 and W599 are de Vries materials based on many criteria. They both show very large electroclinic

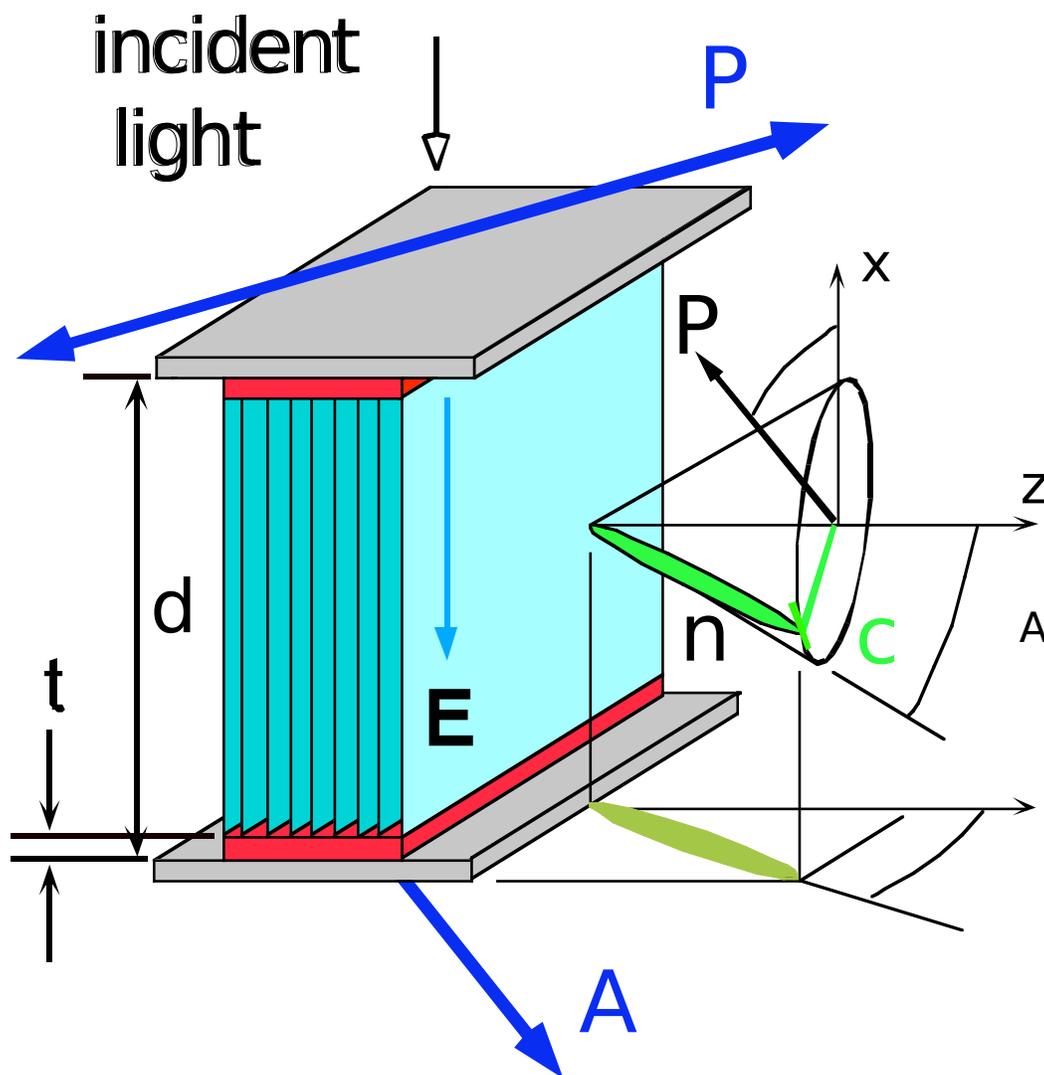


Figure 5.5: The basic bookshelf cell geometry. Smectic layers are perpendicular to the cell substrate. Fields are applied along the cell normal, \hat{x} . The molecules lie on the surface of a cone and are tilted away from the layer normal, \hat{z} , by angle θ_A . The angle the polarization vector makes with the cell normal is ϕ , and ψ is the projection of the molecule onto the $y - z$ plane.

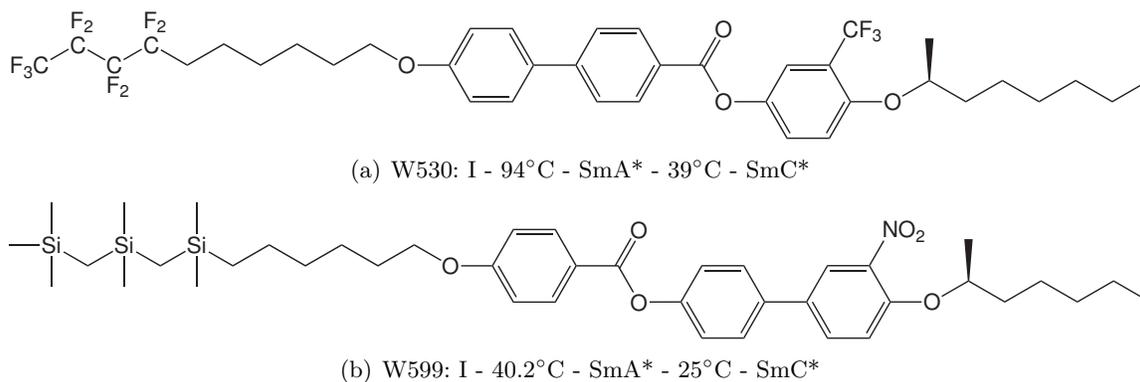


Figure 5.6: Materials W530 and W599 with phase sequences. These chiral de Vries materials show a double peak in current polarization current (or a sigmoidal polarization) which is characteristic of the de Vries SmA phase and of reminiscent of SmC_A phases. These materials both have special tails, perfluorinated in W530 and silocarbane in W599, which are meant to support strong layering, a characteristic known to promote de Vries behavior.[63]

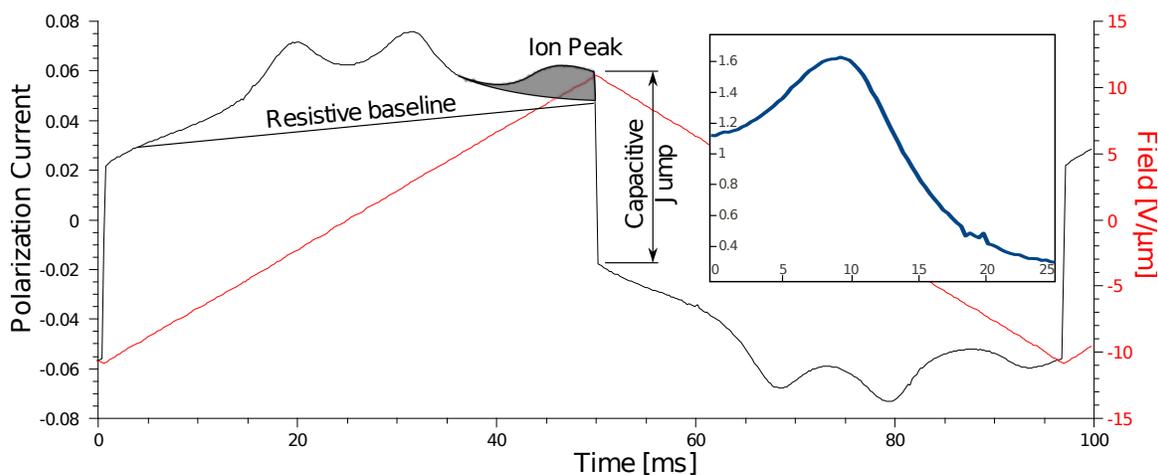


Figure 5.7: An example of the raw data from our polarization current measurements. The current consists of several background contributions (a linear resistive baseline, a capacitive jump, and an ‘ion peak’ whose origins are not entirely understood) which have to be removed before analysis. The inset shows the polarization current once the background has been removed.

response, and polarization current response, as well as a large increase in birefringence with field and a small decrease in layer thickness upon entering the SmC phase.[73, 66] In addition, we see the characteristic double peak switching behavior of the polarization current found in many de Vries smectics, as seen in Figure 5.7. This kind of behavior has been seen by several groups in a variety of different de Vries materials.[73, 58] Double peak behavior often indicates a barrier between the ground state and the ferroelectric state. This can arise for several reasons, but is commonly seen in antiferroelectric systems. In de Vries systems, however, this double peak behavior is due to the proximity of a first order phase transition.

5.5 Discussion

In addition to exploring the phase behavior of the system, we examined the polarization field response of the model. Figure 5.9 demonstrates the qualitative agreement between the spin model and a series of polarization response curves for de Vries material W530 measured via the polarization reversal current [73, 66]. These fits capture the basic sigmoidal behavior over a substantial range of temperatures above the SmA-SmC phase transition using only three free parameters.

In summary, we've shown that the implications of a hollow cone smectic go further than the layer spacing and electro-optics to which it has been previously applied. The steric interactions inherent in a hollow cone smectic imply ϕ -correlated domains and, at sufficiently large cone angles, a first order SmA-SmC phase transition which leads to the observed sigmoidal field response. Notably this analysis involves many simplifying assumptions, including a rigid cone, no out-of-layer fluctuations, and an idealized molecular model that excludes many important temperature dependent features. Even without these features present in the theory we find that the hollow cone model is a microscopic picture of de Vries smectics that encompasses most of the experimentally observed characteristics of the phase including the common first order phase transition and qualitative agreement with the electro-optics.

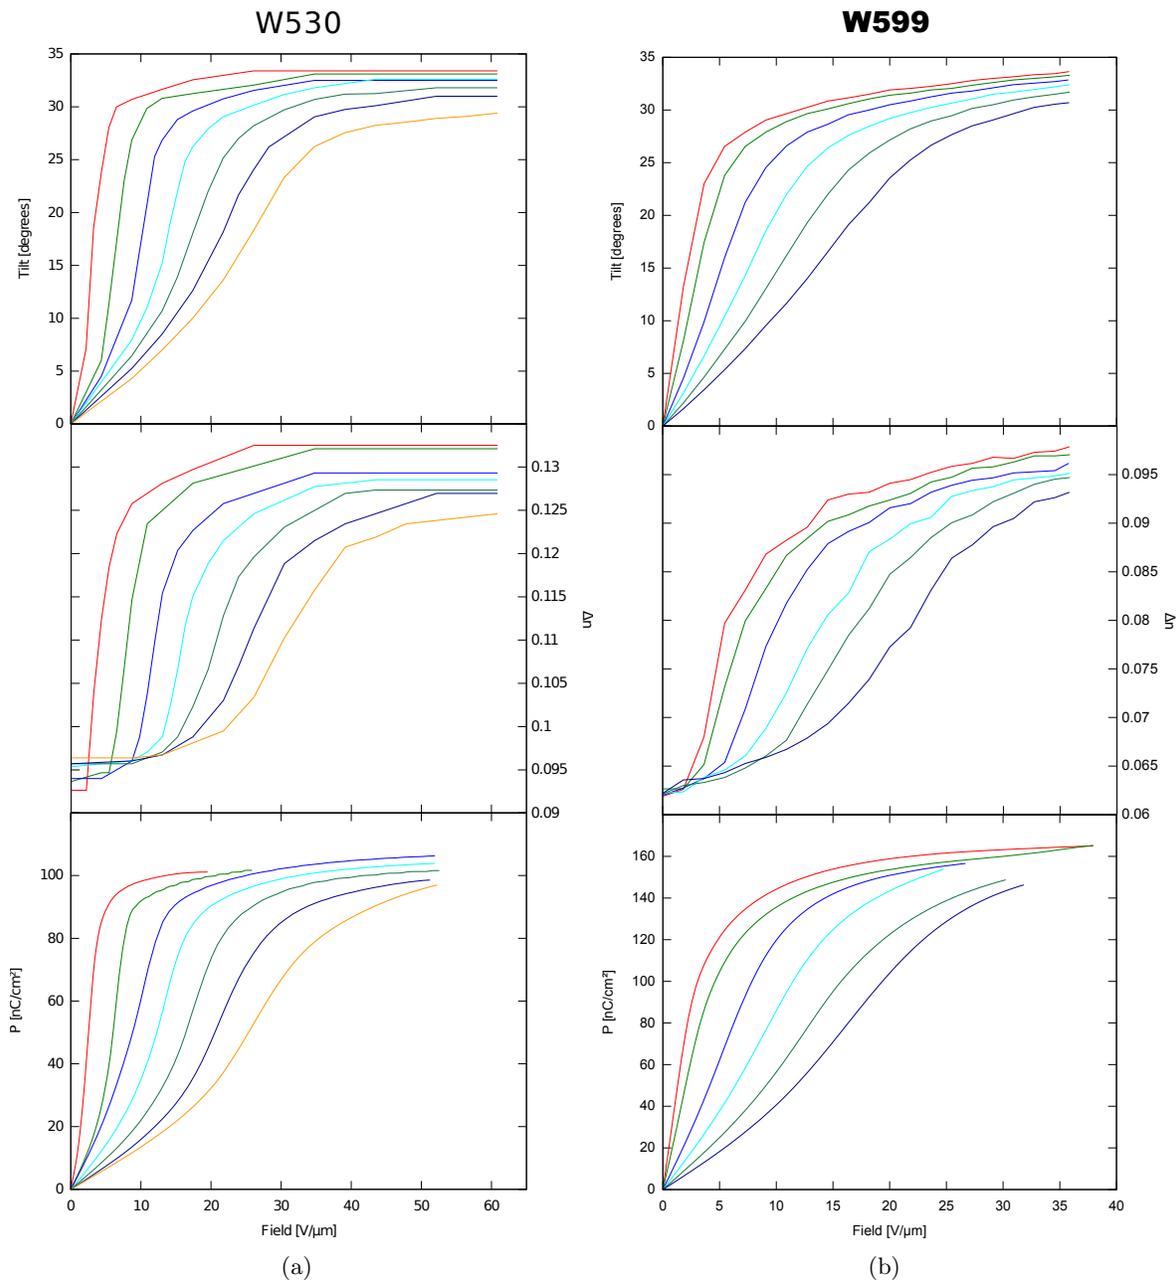


Figure 5.8: Behavior of materials W530 and W599 near their respective SmA \rightarrow SmC transitions at a series of temperatures. We see the characteristically large change in tilt and birefringence in the top two graphs. The bottom graphs show the polarization response which has a sigmoidal, or 'S' shaped, character. This is equivalent to having double peaked polarization current.

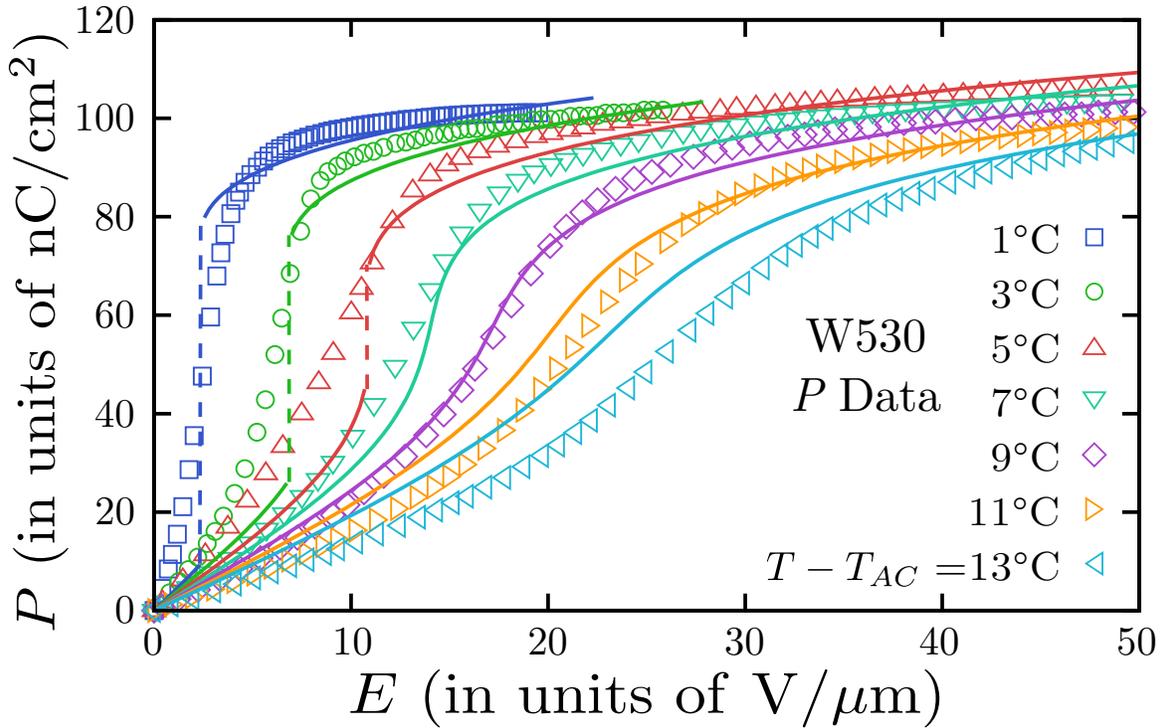


Figure 5.9: Comparison of the polarization density response of de Vries material W530 (symbols) when an electric field is applied across a liquid crystal cell with bookshelf alignment [73] and the model's predictions for $\gamma = 6$ (curves). The model's energy units at a given value of γ are set by equating the model's SmA-SmC transition temperature, T_{AC} , to the transition temperature for W530, 39°C . The polarization density saturation and susceptibility is scaled to demonstrate simultaneous qualitative agreement with several polarization curves near the SmA-SmC transition. At lower temperatures we predict a discontinuous change in P which is not seen experimentally. The continuous experimental behavior may be due to quenched surface disorder in the cell.

Chapter 6

Conclusion

The minimal models presented and characterized here show emergent behavior, behavior that is not trivially reducible to microscopic components of the system. In each of the three systems examined, the underlying model is exceptionally simple, involving hard objects and simple interactions, yet giving rise to complex, counterintuitive, and compelling results, respectively. The HCSS systems exhibit extreme phase diversity due to a competition between a clustering instability and excluded volume interactions. But this complexity provides for opportunity in terms of directed self-assembly. The spontaneous chirality in bundles is related to the orientational entropy increase associated with macroscopic twist. This counterintuitive result is largely overlooked in literature. The first order phase transition in the de Vries system, and in fact nearly all observable de Vries features, can be seen as an result of the hollow cone model. This makes for a compelling argument for the hollow cone model over competing microscopic models of the de Vries smectic phase.

Minimal models are crucial to our understanding of the phenomena at work. These simple systems allow us to understand the fundamental behavior without getting trapped by inconsequential details or computational limitations. Even with these minimal models, the complexity and computational difficulties that can arise requires significant effort to overcome often requiring algorithm development and application of advanced computational techniques. Further, the minimal nature of the HCSS model and the emergence of the clustering instability allows us to utilize the simple directed design methods of Section 2.4. In the chiral bundle system, the minimal nature of the model allowed us to identify the precise source of the increase of entropy to account for the

observed effect. In the de Vries system, the minimal spherocylinder based model made it simple to perform the necessary mapping from the off lattice model to the generalized XY model. The XY model, a minimal model in its own right, allowed us to identify a defect condensation mechanism for the first order SmA-SmC phase transition. These minimal models allow us to develop intuition for each system while probing the emergent phenomena that governs their properties.

Bibliography

- [1] M. P. Allen and D. J. Tildesley. Computer simulation of liquids. Clarendon Press ; Oxford University Press, Oxford [England]; New York, 1989.
- [2] P. W. Anderson. More is different. Science, 177(4047):393–396, August 1972. PMID: 17796623.
- [3] Natalie Arkus, Vinothan N. Manoharan, and Michael P. Brenner. Minimal energy clusters of hard spheres with short range attractions. Physical Review Letters, 103(11):118303, September 2009.
- [4] Ch. Bahr and G. Heppke. Influence of electric field on a first-order smectic-ferroelectric-smectic-c liquid-crystal phase transition: A field-induced critical point. Physical Review A, 41(8):4335–4342, April 1990.
- [5] J. A. Barker, D. Henderson, and F. F. Abraham. Phase diagram of the two-dimensional lennard-jones system; evidence for first-order transitions. Physica A: Statistical Mechanics and its Applications, 106(12):226–238, March 1981.
- [6] Isaac R. Bruss and Gregory M. Grason. Non-euclidean geometry of twisted filament bundle packing. Proceedings of the National Academy of Sciences of the United States of America, 109(27):10781–10786, July 2012. PMID: 22711799 PMCID: PMC3390840.
- [7] Isaac R. Bruss and Gregory M. Grason. Topological defects, surface geometry and cohesive energy of twisted filament bundles. arXiv e-print 1303.1566, March 2013.
- [8] Philip J. Camp. Structure and phase behavior of a two-dimensional system with core-softened and long-range repulsive interactions. Physical Review E, 68(6):061506, December 2003.
- [9] P. M. Chaikin and T. C. Lubensky. Principles of Condensed Matter Physics. Cambridge University Press, Cambridge; New York, NY, USA, reprint edition edition, October 2000.
- [10] Qian Chen, Sung Chul Bae, and Steve Granick. Directed self-assembly of a colloidal kagome lattice. Nature, 469(7330):381–384, January 2011. PMID: 21248847.
- [11] N. A. Clark, T. Bellini, R.-F. Shao, D. Coleman, S. Bardon, D. R. Link, J. E. MacLennan, X.-H. Chen, M. D. Wand, D. M. Walba, P. Rudquist, and S. T. Lagerwall. Electro-optic characteristics of de vries tilted smectic liquid crystals: Analog behavior in the smectic a^[sup] and smectic c^[sup] phases. Applied Physics Letters, 80(22):4097, 2002.

- [12] Nigel Clarke, Jose A. Cuesta, Richard Sear, Peter Sollich, and Alessandro Speranza. Phase equilibria in the polydisperse zwanzig model of hard rods. The Journal of Chemical Physics, 113(14):5817, 2000.
- [13] Piers Coleman. Introduction to Many Body Physics.
- [14] Thomas H Cormen. Introduction to algorithms. MIT Press, Cambridge, Mass., 2009.
- [15] Adriaan De Vries. Experimental evidence concerning two different kinds of smectic c to smectic a transitions. Molecular Crystals and Liquid Crystals, 41(2):27–31, 1977.
- [16] Adriaan De Vries. The implications of the diffuse-cone model for smectic a and c phases and a-c phase transitions. Molecular Crystals and Liquid Crystals, 49:179, 1979.
- [17] Adriaan De Vries, Arun Ekachai, and N. Spielberg. Why the molecules are tilted in all smectic a phases, and how the layer thickness can be used to measure orientational disorder*. Molecular Crystals and Liquid Crystals, 49:143, 1979.
- [18] Eytan Domany, Michael Schick, and Robert H. Swendsen. First-order transition in an xy model with nearest-neighbor interactions. Physical Review Letters, 52(17):1535, April 1984.
- [19] T. Dotera, T. Oshiro, and P. Ziherl. Mosaic two-lengthscale quasicrystals. Nature, 506(7487):208–211, February 2014.
- [20] C. Dress and W. Krauth. Cluster algorithm for hard spheres and related systems. Journal of Physics A: Mathematical and General, 28(23):L597, December 1995.
- [21] Erica Flapan. Topological rubber gloves. Journal of Mathematical Chemistry, 23(1-2):31–49, April 1998.
- [22] Henrik Flyvbjerg. Error estimates on averages of correlated data. In Jnos Kertesz and Imre Kondor, editors, Advances in Computer Simulation, volume 501 of Lecture Notes in Physics, pages 88–103. Springer Berlin / Heidelberg, 1998. 10.1007/BFb0105461.
- [23] Daan Frenkel. Speed-up of monte carlo simulations by sampling of rejected states. Proceedings of the National Academy of Sciences of the United States of America, 101(51):17571–17575, December 2004.
- [24] Daan Frenkel and Berend Smit. Understanding molecular simulation from algorithms to applications. Academic Press, San Diego, 2 edition, 2002.
- [25] S. Ghosh, P. Nayek, S. K. Roy, T. Pal Majumder, M. Zurowska, and R. Dabrowski. Double-peak polarization current response in the unusual SmA^{*} phase of a fluorinated high-tilt anti-ferroelectric liquid crystal. EPL (Europhysics Letters), 89(1):16001, January 2010.
- [26] Thomas Gibaud, Edward Barry, Mark J. Zakhary, Mir Henglin, Andrew Ward, Yasheng Yang, Cristina Berciu, Rudolf Oldenbourg, Michael F. Hagan, Daniela Nicastro, Robert B. Meyer, and Zvonimir Dogic. Reconfigurable self-assembly through chiral control of interfacial tension. Nature, 481(7381):348–351, January 2012.
- [27] Matthew A Glaser, Gregory M Grason, Randall D Kamien, A. Kosmrlj, Christian D Santangelo, and P. Ziherl. Soft spheres make more mesophases. cond-mat/0609570, September 2006. Europhys. Lett. 78 (2007) 46004.

- [28] Sharon C. Glotzer and Michael J. Solomon. Anisotropy of building blocks and their assembly into complex structures. Nature Materials, 6(8):557–562, August 2007.
- [29] M. V. Gorkunov, F. Giesselmann, J. P. F. Lagerwall, T. J. Sluckin, and M. A. Osipov. Molecular model for de vries type smectic- a smectic- c phase transition in liquid crystals. Physical Review E, 75(6):060701, June 2007.
- [30] M. V. Gorkunov, M. A. Osipov, J. P. F. Lagerwall, and F. Giesselmann. Order-disorder molecular model of the smectic- a smectic- c phase transition in materials with conventional and anomalously weak layer contraction. Physical Review E, 76(5):051706, November 2007.
- [31] V. Granville, M. Krivanek, and J.-P. Rasson. Simulated annealing: a proof of convergence. IEEE Transactions on Pattern Analysis and Machine Intelligence, 16(6):652–656, June 1994.
- [32] N. Grewe and W. Klein. The KirkwoodSalsburg equations for a bounded stable kac potential. i. general theory and asymptotic solutions. Journal of Mathematical Physics, 18(9):1729–1734, 1977.
- [33] J. R. Heringa and H. W. J. Blte. Geometric cluster monte carlo simulation. Physical Review E, 57(5):4976, May 1998.
- [34] Terrell L. Hill. An Introduction to Statistical Thermodynamics. Courier Dover Publications, 1960.
- [35] E. A. Jagla. Phase behavior of a system of particles with core collapse. Physical Review E, 58(2):1478, 1998.
- [36] E. A. Jagla. Minimum energy configurations of repelling particles in two dimensions. The Journal of Chemical Physics, 110(1):451, 1999.
- [37] R. A. Jarvis. On the identification of the convex hull of a finite set of points in the plane. Information Processing Letters, 2(1):18–21, March 1973.
- [38] Matthew R Jones, Robert J Macfarlane, Byeongdu Lee, Jian Zhang, Kaylie L Young, Andrew J Senesi, and Chad A Mirkin. DNA-nanoparticle superlattices formed from anisotropic building blocks. Nature materials, 9(11):913–917, November 2010. PMID: 20890281.
- [39] Nadia Kapernaum, David M Walba, Eva Korblova, Chenhui Zhu, Chris Jones, Yongqian Shen, Noel A Clark, and Frank Giesselmann. On the origin of the Giant electroclinic effect in a De VriesType ferroelectric liquid crystal material for chirality sensing applications. ChemPhysChem, 10(6):890–892, April 2009.
- [40] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. Science, 220(4598):671–680, May 1983. PMID: 17813860.
- [41] Zachary V. Kost-Smith, Paul D. Beale, Noel A. Clark, and Matthew A. Glaser. Microscopic origins of first-order sm-ASm-C phase behavior in de vries smectic liquid crystals. Physical Review E, 87(5):050502, May 2013.
- [42] Tatiana Kuriabova, M. D. Betterton, and Matthew A. Glaser. Linear aggregation and liquid-crystalline order: comparison of monte carlo simulation and analytic theory. Journal of Materials Chemistry, 20(46):10366–10383, November 2010.

- [43] Jan P. F. Lagerwall and Frank Giesselmann. Current topics in smectic liquid crystal research. ChemPhysChem, 7(1):20–45, January 2006.
- [44] Sven T. Lagerwall, Per Rudquist, and Frank Giesselmann. The orientational order in so-called de vries materials. Molecular Crystals and Liquid Crystals, 510(1):148/[1282]–157/[1291], 2009.
- [45] D.P. Landau and F. Wang. Determining the density of states for classical statistical models by a flat-histogram random walk. Computer Physics Communications, 147(1-2):674–677, August 2002.
- [46] C. N. Likos, A. Lang, M. Watzlawek, and H. Lwen. Criterion for determining clustering versus reentrant melting behavior for bounded interaction potentials. Physical Review E, 63(3):031206, February 2001.
- [47] Jiwen Liu and Erik Luijten. Rejection-free geometric cluster algorithm for complex fluids. cond-mat/0412744, December 2004. Phys. Rev. Lett. 92, 035504 (2004).
- [48] Jiwen Liu and Erik Luijten. Generalized geometric cluster algorithm for fluid simulation. cond-mat/0503448, March 2005. Phys. Rev. E 71, 066701 (2005).
- [49] Gianpietro Malescio and Giuseppe Pellicane. Stripe phases from isotropic repulsive interactions. Nature Materials, 2(2):97–100, February 2003.
- [50] W. L. McMillan. Simple molecular model for the smectic a phase of liquid crystals. Physical Review A, 4(3):1238, 1971.
- [51] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. The Journal of Chemical Physics, 21(6):1087–1092, June 1953.
- [52] Bianca M. Mladek, Patrick Charbonneau, Christos N. Likos, Daan Frenkel, and Gerhard Kahl. Multiple occupancy crystals formed by purely repulsive soft particles. Journal of Physics: Condensed Matter, 20(49):494245, December 2008.
- [53] H. Nagata, M. Zukovic, and T. Idogaki. Monte carlo simulation of the three-dimensional XY model with bilinear-biquadratic exchange interaction. Journal of Magnetism and Magnetic Materials, 234(2):320–330, September 2001.
- [54] Ferenc Niedermayer. General cluster updating method for monte carlo simulations. Physical Review Letters, 61(18):2026–2029, October 1988.
- [55] B. Nienhuis, A. N. Berker, Eberhard K. Riedel, and M. Schick. First- and second-order phase transitions in potts models: Renormalization-group solution. Physical Review Letters, 43(11):737–740, September 1979.
- [56] Peter Norvig. Paradigms of artificial intelligence programming: case studies in Common Lisp. Morgan Kaufman Publishers, San Francisco, Calif., 1992.
- [57] Gernot J. Pauschenwein and Gerhard Kahl. Zero temperature phase diagram of the square-shoulder system. The Journal of Chemical Physics, 129(17):174107, November 2008.

- [58] S. Krishna Prasad, D. S. Shankar Rao, S. Sridevi, Chethan V. Lobo, B. R. Ratna, Jawad Naciri, and R. Shashidhar. Unusual dielectric and electrical switching behavior in the de vries smectic a phase of two organosiloxane derivatives. Physical Review Letters, 102(14):147802, April 2009.
- [59] P. Prybytak, W. J. Frith, and D. J. Cleaver. Hierarchical self-assembly of chiral fibres from achiral particles. Interface Focus, page rsfs20110104, March 2012. PMID: 24098850.
- [60] Mikael C. Rechtsman, Frank H. Stillinger, and Salvatore Torquato. Optimized interactions for targeted self-assembly: Application to a honeycomb lattice. Physical Review Letters, 95(22):228301, November 2005.
- [61] Jeffrey C. Roberts, Nadia Kapernaum, Qingxiang Song, Dorothee Nonnenmacher, Khurshid Ayub, Frank Giesselmann, and Robert P. Lemieux. Design of liquid crystals with de vries-like properties: Frustration between SmA- and SmC-Promoting elements. Journal of the American Chemical Society, 132(1):364–370, January 2010.
- [62] Karl Saunders. de vries behavior in smectic liquid crystals near a biaxiality-induced smectic-a smectic- c tricritical point. Physical Review E, 77(6):061708, June 2008.
- [63] Karl Saunders, Daniel Hernandez, Staci Pearson, and John Toner. Disorder to order: de vries behavior from a landau theory for smectic phases. Physical Review Letters, 98(19):197801, May 2007.
- [64] Francesco Sciortino, Stefano Mossa, Emanuela Zaccarelli, and Piero Tartaglia. Equilibrium cluster phases and low-density arrested disordered states: The role of short-range attraction and long-range repulsion. Physical Review Letters, 93(5):055701, July 2004.
- [65] Jonathan V. Selinger, Peter J. Collings, and R. Shashidhar. Field-dependent tilt and birefringence of electroclinic liquid crystals: Theory and experiment. Physical Review E, 64(6):061705, November 2001.
- [66] Yongqiang Shen, Lixing Wang, Renfan Shao, Tao Gong, Chenhui Zhu, Hong Yang, Joseph E. Maclennan, David M. Walba, and Noel A. Clark. Generalized langevin-debye model of the field dependence of tilt, birefringence, and polarization current near the de vries smectic-a* to smectic-c* liquid crystal phase transition. Physical Review E, 88(6):062504, December 2013.
- [67] Robert Swendsen and Jian-Sheng Wang. Nonuniversal critical dynamics in monte carlo simulations. Physical Review Letters, 58(2):86–88, January 1987.
- [68] Yoichi Takanishi, Yukio Ouchi, Hideo Takezoe, and Atsuo Fukuda. Chevron layer structure in the smectic a phase of 8CB. Japanese Journal of Applied Physics, 28(Part 2, No. 3):L487–L489, 1989.
- [69] J. M. Tavares, B. Holder, and M. M. Telo da Gama. Structure and phase diagram of self-assembled rigid rods: Equilibrium polydispersity and nematic ordering in two dimensions. Physical Review E, 79(2):021505, February 2009.
- [70] A. Vrij. Polymers at interfaces and the interactions in colloidal dispersions. Pure and Applied Chemistry, 48(4), January 1976.

- [71] David M. Walba, Eva Korblova, Lior Eshdat, Michael C. Biewer, Hong Yang, Christopher Jones, Michi Nakata, Mara Talarico, Renfan Shao, and Noel A. Clark. Chiral SmA* materials for display applications? Journal of the Society for Information Display, 15(8):585588, 2007.
- [72] Jian-Sheng Wang and Robert H. Swendsen. Cluster monte carlo algorithms. Physica A: Statistical Mechanics and its Applications, 167(3):565–579, September 1990.
- [73] L. Wang, D. M. Walba, R.-F. Shao, D. Coleman, M. Nakata, T Shimmachi, J. E. MacLennan, and N. A. Clark. In Proceedings of the 21st International Liquid Crystal Conference. Taylor & Francis, 2006.
- [74] Yu Wang, Yufeng Wang, Xiaolong Zheng, Gi-Ra Yi, Stefano Sacanna, David J. Pine, and Marcus Weck. Three-dimensional lock and key colloids. Journal of the American Chemical Society, 136(19):6866–6869, May 2014.
- [75] Yufeng Wang, Yu Wang, Dana R Breed, Vinothan N Manoharan, Lang Feng, Andrew D Hollingsworth, Marcus Weck, and David J Pine. Colloids with valence and specific directional bonding. Nature, 491(7422):51–55, November 2012. PMID: 23128225.
- [76] Stephen Whitelam, Edward H Feng, Michael F Hagan, and Phillip L Geissler. The role of collective motion in examples of coarsening and self-assembly. 0806.3903, June 2008. Soft Matter, 2009, 5, 1251 - 1262.
- [77] Stephen Whitelam and Phillip L. Geissler. Avoiding unphysical kinetic traps in monte carlo simulations of strongly attractive particles. The Journal of Chemical Physics, 127(15):154101, 2007.
- [78] Ulli Wolff. Collective monte carlo updating for spin systems. Physical Review Letters, 62(4):361–364, January 1989.
- [79] Yasheng Yang, Edward Barry, Zvonimir Dogic, and Michael F. Hagan. Self-assembly of 2D membranes from mixtures of hard rods and depleting polymers. Soft Matter, 8(3):707, 2012.
- [80] D. A. Young and B. J. Alder. Melting-curve extrema from a repulsive "Step" potential. Physical Review Letters, 38(21):1213–1216, May 1977.
- [81] Robert Zwanzig. First-order phase transition in a gas of long thin rods. The Journal of Chemical Physics, 39(7):1714, 1963.

Appendix A

Efficient sampling in Monte Carlo

Special care must be placed on the time complexity of algorithms used in Monte Carlo. This is not critical, but it seems to usually be the case that Monte Carlo sampling and simulation in general should be computable in a $O(N)$ time. By this, I mean that the time it takes to advance, on average, every element in the system should scale linearly with the size of the system. When this is not the case, we usually see an unacceptable slow in the calculation. In molecular dynamics, this manifests as an $O(N)$ calculation of system forces and torques followed by an update of every element in the system. In Monte Carlo, where Monte Carlo moves typically only move a single particle at a time, this scaling constraint limits us to moves that take only constant ($O(1)$) time to accept or reject. There are well defined ways to achieve constant time energy and force calculation per particle. In Monte Carlo, this usually needs to be done via cell lists as these are easier to update than neighbor lists and are more suitable for non-local moves.[1, 24]

This expectation of linear scaling limits the influences the techniques that we can use in simulation. This means that in the case of collective moves, the efficiency must be evaluated in terms of number of elements moved, not the size of the system. In the case where the cluster percolates to the entire system size, this cannot be avoided (although, in this case we are technically still linear in the number of particles moved).

A.1 Optimization of algorithm parameters

In the algorithms listed in this thesis there are often free parameters that are unspecified. These free parameters are parameters that can be tuned to optimize the sampling for the system at hand.

In the case of Monte Carlo techniques, it is absolutely necessary to find approximate optimal parameters for efficient sampling of the phase space. When we wish to actually perform the sampling, the only significant factor is the sampling rate divided by the wall time for equivalent computational resource. Even here this is tricky to define. Sampling efficiency is deeply dependent on the computation you wish to perform. Further, modern systems have many components that can operate concurrently. To truly determine the best computational method, one would need a sophisticated cost-benefit analysis that is beyond the scope of this document.

Often, the most effective move in a simulation requires knowledge about the structure of the system and thus cannot usually be derived a priori even when the system at hand is known. For instance, take our system of HCRS spheres with the standard Metropolis algorithm. One free parameter of this system is the maximum distance of an attempted particle displacement. A random displacement of a particle is not optimal if that random displacement tends to be so small that the moves don't change the configuration very much. Nor is it optimal if the displacement is too big, resulting in an increased chance of overlaps and subsequent high rejection rates. Thus, there tends to be an optimal, or multiple local optimal, values for the maximum attempted displacement.

We cannot use standard optimization techniques as they generally assume that the target function is fixed. Our target function is simultaneously dynamic due to the process of equilibration, and fluctuating due to the statistical nature of the Markov sampling process. The similarity of this optimization problem with optimization in the field of machine learning allows us to utilize some of these techniques here.

A.2 Numerical steepest descent

The simplest method that can be employed for this purpose is making a crude estimate of the derivative of the efficiency function, $\eta(\vec{\xi})$ where $\vec{\xi}$ is the vector of free parameters, based on several blocks of data. This is similar to techniques used in “online learning” in the field of machine learning. The big difference is that the derivatives in machine learning are analytically calculated. Due to the numerical component of this process, this does not scale to large, or even moderate, parameter vectors. The numerical calculation of each partial derivative of the target function is not possible within the time-frame of a typical simulation.

Appendix B

Odds and ends of soft matter computation

For completeness, a discussion of the several topics that don't fit into the overall narrative is warranted. The section here contain information necessary tools and techniques for computational work, including some new developments by myself and the group.

B.1 Treating excessive unit cell shear in small systems

When generating minimal enthalpy structures in Chapter 2, it is almost imperative to use small system sizes. This removes the much of the possibility for defects in the annealed structures which we know will not be the part of the minimum enthalpy configuration. One of the problems with this is that a general assumption of fluid simulation with periodic boundary conditions is that each particle in the system never interacts with a periodic image of itself. This assumption is mostly made for convenience as it simplifies the summations involved in the total system potential energy calculations. This assumption is guaranteed to be true if the perpendicular dimensions of the cell are larger than the interaction range. This is not true in these small systems.

Johannes Hausinger, a student in the group under the oversight of Dr. Glaser, developed the concept of *supercells*, which make replicas of the computational unit cell, partially unrolling the periodic boundary conditions, in order to make this assumption true for the duration of the potential energy and (if applicable) force calculations. These supercells allow systems to be as small as one particle and are used throughout the HCSS study.

At high temperatures, the unit cell of these systems can deform greatly as the density of

the system is exceptionally low and the potential energy is negligible. The unit cell can become excessively elongated or highly sheared. The elongation in any particular direction is not overly troubling as it is limited by the fact that the unit cell volume is fixed by the structure. The shear, on the other hand, is not limited by this minimum volume as shear moves can operate without changing the volume of the unit cell. This becomes very problematic as the perpendicular dimensions of the unit cell does change with shear and approaches zero as the unit cell vectors approach co-linearity. As these perpendicular dimensions become small, the number of supercells required to computed the potential diverges and with it the computational effort.

This is a real world problem that only really exists when applying Monte Carlo sampling or molecular dynamics simulation to small systems at high temperature. This can lead to computations that mysteriously hang and balloon to immense resident memory footprints.

I developed a method to fix this issue by redefining the unit cell to an equivalent, but unskewed, version. This is done by noting that the computational unit cell defines a potential crystal unit cell that is only one of an infinite number of possible unit cell. By using the lattice vectors, of the unit cell, $\{\vec{a}_i\}$, you can define a new set of lattice vectors, $\{\vec{a}'_i\}$, that are a linear combination of the original vectors where the coefficients are integers. This needs to be done in such a way that minimizes the unit cell skew.

A good heuristic for this is to maximize the minimum of the absolute values of the normalized dot products, i.e. maximize the quantity $\min_{i \neq j} |a'_i \cdot a'_j|$.

One nice thing about this transformation is that this does not change the system in any physical way, so this can be done at any time without concern for the Markov chain or detailed balance.

Algorithm 12 Transform the unit cell of the into an equivalent representation that is less skewed. The time scaling of this algorithm is $O(N)$ due to the need to apply periodic boundary conditions over the system at the end.

```

1: function UNSKEW_CONFIG( $\{\vec{r}\}, M$ )
2:    $M' \leftarrow$  UNSKEW_MAT( $M$ )
3:   return (PBC( $\{\vec{r}\}, M'$ ),  $M'$ )
4: end function
5:
6: function UNSKEW_MAT( $M$ )
7:    $(\vec{a}_i, \vec{a}_j) \leftarrow$  maximize  $\vec{a}_i \cdot \vec{a}_j$  ▷ Find maximal dot product.
       where  $\vec{a}_{i,j} \in$  column_vectors( $M$ )
       and  $i \neq j$ 
       and  $|\vec{a}_i| > |\vec{a}_j|$ 
8:    $\vec{a}'_i \leftarrow \vec{a}_i - \text{sign}(\vec{a}_i \cdot \vec{a}_j) \times \vec{a}_j$  ▷ Redefine unit vector.
9:   if  $|\vec{a}'_i| < |\vec{a}_i|$  then
10:    return UNSKEW_MAT( $M_{\vec{a}_i \rightarrow \vec{a}'_i}$ ) ▷ Recurse until a stationary matrix is found.
11:  else
12:    return  $M$  ▷ Stationary matrix found.
13:  end if
14: end function

```

Appendix C

VMMC implementation details

C.1 The search order

The order in which you search for particles to link in the VMMC procedure can be quite important. While it does not effect the correctness of the method, it can effect the efficiency of the algorithm. While it is not clear from the algorithm defined in Whitelam and Geissler’s papers, the use of the term “iteratively” indicates that VMMC was designed as a breadth first clustering search. In such a search order, if i will perform a link against j , all neighbors of i are considered for inclusion in the cluster before the any of the neighbors of a j will be. This could be done in other ways, however, by simply changing the way that new particles are added to the cluster in line 20 in Algorithm 7 and its corresponding lines in Algorithms 9 and 10.¹ Possible search methods include breadth-first, where potential additions are considered in the order that they are found, depth-first, where the most recent potential additions are considered first, and more exotic search methods such as *best first search* where “best” is defined by some ability to prove whether the cluster will be ultimately rejected or not.[56] These different search orders will give different cluster realizations, though crucially the same sites are considered for inclusion. Once the seed is picked, the structure of the cluster is basically defined, and the search shouldn’t have any affect (statistically) on the cluster actually produced, only on the realization.

The depth-first search path will tend to consider more distant particles for inclusion before the near by particles. This could be a useful in an early rejection scheme if there is a set maximum

¹ Care should be taken to ensure that all of the neighbors of particles in the cluster are considered for inclusion.

distance between the seed particle and the particles in the cluster² The breadth-first search path will consider all local particles first. This can be beneficial for non-self-inverse algorithms such as VMMC as, in principle, one could define a procedure for early rejection by proving the existence of a frustrated external link.

C.2 Comparison to the original VMMC algorithm

When working with this algorithm, I limited myself to the “rejection-free” version of VMMC described in Section 3.3.2, only mentioning the jumping off point for the original method.[76] This is due to its similarity to the Swendsen-Wang class of cluster moves. In this section, I look at some of the differences between the two versions.

C.2.1 Reverse balance

The original algorithm did not include the reverse balance in the linking method. This had the effect of a more complicated acceptance probability with several. Namely, the factors in the probability that correspond to failed internal links remain in the product. Here I discuss the possible differences in efficiency between the two methods.

When evaluating a VMMC move for any two pairs of particles in the system, we are concerned with **three** interactions. If we have picked particle i to undergo the VMMC procedure and are considering adding particle j to the cluster move, the success of the move naturally depends on the energy in the initial and final state, or separation, of these two particles. However, depending on the probabilistic linking of i and j , the final separation might be the different or the same as the initial. Whatever the outcome of the linking, we denote the initial state with the symbol μ and the final state ν . It is important to note that the final state is the state of the system if the entire cluster, which has yet to be determined, is moved and accepted.

In the standard VMMC procedure where the reverse move is not considered as part of the

² You often need to impose such a rejection as certain symmetry operations are only locally symmetric. One such operation is rotation by a small angle.

linking method, we can actually use this information to our advantage. If we form a link and we know, we can test the reverse linking probability at that time. If this reverse linking probability cannot happen, then we can exit immediately. Note that this is different from the rejection-free version of VMMC. In this case we have detected a situation where the probability for acceptance is zero, in the rejection-free method we are removing some of the bias of the move proposal method. In strongly anisotropic cases it seems quite likely that early rejection is a win as it is rare for particles to be linked if a failure was encountered.

C.3 The indicator function

In the original VMMC algorithm, the link probability included a term that limited the linking checks to particles that interact with the particle in the initial μ state. This was chosen by Whitelam and Geissler for the goal of keeping moves local, contrasting the algorithm with the GenGCA. Given that the proposed maps in VMMC are local in nature, since they are usually translations, this seems to be of little benefit. The result of this choice by Whitelam and Geissler is, again, a more complicated acceptance probability due to neglected failed external links. The execution of the algorithm is naturally different, including some efficiency considerations.

When we consider systems where whatever move maps are used don't necessarily have the property that if transforming i 's coordinates over a move map brings two particles, i and j into interaction range, the inverse move map acting on i will produce a configuration where they are not interacting. This becomes very clear when looking at systems with long distance translational maps and periodic boundary conditions, systems with rotational maps and extended particles (not point particles), and a plethora of possibilities when dealing with concavity in your extended objects. In these cases, prudent choice the meaning of the word interacting becomes important. If, for instance, we are simulating banana like molecules, a liquid crystal favorite, the indicator function is likely to miss opportunities of linking in clusters of inter-digitated banana molecules when translating along the director. Another more general example, consider simulation of rods in a nematic phase. Even slight rotations of one rod can produce overlaps. These overlaps produce can occur in both

directions of rotation. This means that the original Whitelam-Geissler procedure should fail to exploit this opportunity for collective motion, unless the indicator function is chosen to include more than just what is interacting.

One extremely useful feature of the indicator function is simplifying the implementation of the algorithm. VMMC by definition involves analyzing the system in two different situations. The way simulations are programmed, and quite possibly the only efficient way to do so, is to have a global mutable state for the simulation. This means that moving a particle, even in a virtual sense, requires mutation of that state, which can be an expensive operation. For instance, motion of a particle in a simulation can involve the update of the system neighbor lists, or partial update of the cell lists. The indicator function allows for a way to remove this expense. The indicator function, as described in the procedure by Whitelam and Geissler limited the calculation to the state of the original system. This means that no, possibly expensive update need be performed. Besides providing a possible (but dubious) reduction in computation and non-locality of the move, it allows for algorithms implementing VMMC to be much simpler, particularly when dealing with cluster generating moves other than translation. The flip side of this argument is that you must keep track of considerably more data when there is a bias correction in the acceptance criterion.

C.4 Benefits of VMMC over GenGCA

VMMC provides a way to make local changes to the configuration, which means that you are more likely to get accepted mutations than with many non-local techniques. GenGCA, however, can be extended to do the same. If we make the minor change to the algorithm so that we first select a particle and then select a pivot point within a certain radius of the particle, the moves can be made to be arbitrarily local like in VMMC. This type of inversion operation is a local mutation for small collective moves, but becomes non-local as the cluster of particles grows further from the pivot point. This effect is not an issue with translations in VMMC. However, rotations, which are typically used in VMMC, will suffer from the same issues.