THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

Exactly Solvable Models for Self-Assembly

ERIK EDLUND

Department of Energy and Environment
Division of Physical Resource Theory, Complex Systems Group
CHALMERS UNIVERSITY OF TECHNOLOGY
Göteborg, Sweden 2014

Exactly Solvable Models for Self-Assembly
ERIK EDLUND

ISBN 978-91-7597-036-3

Doktorsavhandlingar vid Chalmers tekniska högskola
Ny serie nr 3717
ISSN 0346-718X

Department of Energy and Environment
Chalmers University of Technology
SE-412 96 Göteborg, Sweden
www.chalmers.se
Tel. +46-(0)31 772 1000

Author email: `erik.edlund@chalmers.se`

Cover: Clockwise from top left: striped morphology arising from competing interactions; energy spectrum for the spherical spin model with four particle types; mapping linear combinations of spherical harmonics to predictions for self-assembling patchy colloids; a self-assembled chiral lattice.

Printed by Reproservice
Göteborg, Sweden 2014

Exactly Solvable Models for Self-Assembly

ERIK EDLUND
Department of Energy and Environment, Chalmers University of Technology

Abstract

The field of self-assembly studies the spontaneous formation of order from preexisting components. It holds the promise of fabricating tomorrows materials and devices not by traditional methods, but by designing building blocks that will act as supramolecular "atoms" and form the desired structures without external input.

Theoretical understanding in self-assembly has previously been achieved mainly through approximative and simulation-based methods. In contrast, this thesis shows that a range of many-particle models of interest in self-assembly can be mapped onto exactly solvable models from statistical mechanics. For such models the behavior of interest, be it ground states or the full partition function, can be obtained through analytical calculations, without approximations.

The thesis consists of an introductory text and seven appended research papers. It introduces two analytical tools for use in self-assembly.

First, the exactly solved spherical spin model is generalized to arbitrary isotropic interactions, various geometries, and multiple particle types. The resulting models are shown to admit exact solutions for their ground states and these turn out to predict pattern formation in corresponding systems. The theory is developed in Papers I, II, and VI, and used in Paper VII to design patchy colloids for self-assembly.

Second, an exact design method for lattice self-assembly is presented in Papers III-V. Given a target lattice structure, it produces an isotropic potential that can be proven to have the desired lattice as its ground state. The method is used to design potentials that in Monte Carlo simulations cause self-assembly into various two- and three-dimensional lattices. In Paper V a method for simplifying designed potentials is discussed.

Keywords: Self-assembly, classical spin models, patterns, colloids, isotropic interactions, patchy colloids

List of publications

This thesis is based on the work contained in the following papers, referred to by Roman numerals in the text:

I. Erik Edlund and Martin Nilsson Jacobi, "Universality of Striped Morphologies," *Physical Review Letters* **105**, 137203 (2010).

II. Erik Edlund, Oskar Lindgren, and Martin Nilsson Jacobi, "Novel Self-Assembled Morphologies from Isotropic Interactions," *Physical Review Letters* **107**, 085501 (2011).

III. Erik Edlund, Oskar Lindgren, and Martin Nilsson Jacobi, "Designing Isotropic Interactions for Self-Assembly of Complex Lattices," *Physical Review Letters* **107**, 085503 (2011).

IV. Erik Edlund, Oskar Lindgren, and Martin Nilsson Jacobi, "Chiral Surfaces Self-Assembling from Isotropic Interactions," *Physical Review Letters* **108**, 165502 (2012).

V. Erik Edlund, Oskar Lindgren, and Martin Nilsson Jacobi, "Using the Uncertainty Principle to Design Simple Interactions for Targeted Self-Assembly," *Journal of Chemical Physics* **139**, 024107 (2013).

VI. Erik Edlund, Oskar Lindgren, and Martin Nilsson Jacobi, "Predicting Self-Assembled Patterns on Spheres with Multicomponent Coatings," *Soft Matter* **10**, 2955-2960 (2014).

VII. Erik Edlund, Oskar Lindgren, and Martin Nilsson Jacobi, "A Design Path for Hierarchical Self-Assembly of Patchy Colloids," *to be submitted. Preprint: http://arxiv.org/abs/1405.2867*

Publications not in this thesis

Erik Edlund and Martin Nilsson Jacobi, "Renormalization of Cellular Automata and Self-Similarity," *Journal of Statistical Physics* **139**, 972 (2010).

Acknowledgements

Martin, thank you for making me realize that Academia is where I belong. Thank you for showing me what it means to really focus on the core of something, for teaching me to be precise when precision is needed, and for the admonitions to be lazy when it is not. We have had just the right mix of similarities in personality and differences in strengths to make my PhD a delight.

Thank you Oskar for four years of great collaboration. We and Martin have worked closer than most groups, and it has certainly been a recipe for success.

I would also like to express my appreciation of the rest of the people in the Complex Systems Group as well as the Division of Physical Resource Theory. It is an odd but stimulating environment we have, with lots of room for intelligent discussion about everything and nothing.

Thank you friends and family, for love, support, and loads of fun. Martin, Oskar, and Anton, thank you for helpful comments on this manuscript. Any remaining mistakes are mine, of course.

To my grandfather Georg, thanks for your never-ending interest in physics, and to Svante Silvén, thanks for challenging me. You were my greatest influences in choosing this path.

Helena, thanks for everything. I love you.

Göteborg, May 2014
Erik Edlund

Contents

Chapter 1

Introduction

What would be the optimal fabrication method? Why, the one where we don't need to fabricate at all. There would be no need to painstakingly fit together delicate pieces just so, no need for carefully crafted tools nor arthritic artisans. There would be a great deal less trouble and suffering in a world where building blocks assembled as if by magic. This is the dream of directed self-assembly.

This dream is on its way to become reality thanks to our increasing ability to design the interactions between various types of particles. For if we can fashion parts in a way that only allows for a single configuration, then we can step back and let thermal motion do the rest. They will wiggle around, stick together and break loose, twist and turn, until they fit our scheme, and there the interactions will lock them. By careful design we can get heat, this most useless form of energy, to work for us rather than against us.

Why are we so preoccupied with structure anyway? Well, while the constituents of matter are important for its properties, the arrangement of them is crucial. Carbon in the form of diamonds is the hardest natural material, possible to craft only using other diamonds, while in the form of graphite it easily peels off sheets, allowing us to write with it. The difference between a boiled egg and a raw one is only the folding of the proteins, rearranged after the heat breaks their carefully configured non-covalent bonds.

Self-assembly promises to radically improve our control of the structure, and therefore properties, of matter at the nanometer and micrometer scale. It truly is the materials science of the future.

1.1 Prediction and design

The central task in directed self-assembly is *design*: given some desired structure, how do we design a system to organize into this structure? The straightforward way to arrive at a solution is trial and error, picking some experimental system one hopes has the target structure in its repertoire and varying parameters until the goal is reached. However, this is of course a time-consuming way of working.

A way to reduce the time needed would be to use methods to predict the behavior of the system without having to do the experiment. Such methods range from simulations, allowing us to do the experiments *in silico*, to theory allowing for direct calculations, with or without computer assistance. They allow us to reduce the time needed for each cycle of trial and error from months to days, minutes, or even fractions of seconds, allowing us to get real-time feedback on our parameter guesses.

Even more desirable are direct design methods. They remove the need for errors altogether by giving explicit recipes for constructing systems from descriptions of target structures. Thus they embody an even deeper understanding of their systems, where prediction of behavior becomes implicit.

This thesis is about these twin tasks of prediction and design.

1.2 Exactly solvable models

To predict the behavior of a system, we need a working model of its parts and their governing rules, together with a way to analyze the model. Here we need to be aware of the tension between the accuracy of the model and our ability to analyze it. Do we choose a more accurate model, despite it being more complicated and necessitating the use of approximate solutions or simulations? Or do we choose a more tractable model, making the connection to the system of interest more complicated instead?

The approach followed in this thesis is to study *exactly solvable models*, in the sense of Baxter [1]. This means to choose models that are simple enough that we don't need approximations at all, allowing for exact solutions. Note that these are exact but not necessarily rigorous; proper care with infinities and related joys will not be taken here.

In statistical mechanics the usage of exactly solvable models is motivated by *universality*, the observation that the behavior of systems close to phase transitions (*i.e.*, the most interesting part) only depends on symmetries and not on details. This means that once quantities such as critical exponents are known from an exactly solvable model, they have to be the same for all models and experimental systems in the same universality class.

In our case, we will have no such general justification. Instead it is a practical question of showing that the exact solutions to our models can be fruitfully connected to the systems we are really interested in.

In this context we should also note that the method described in Chapter 4 can be used to design potentials that can be proven to have certain lattice structures as ground states. These are thus exactly solvable models as well; they are important as rigorously proving a certain structure is a ground state of a given potential is normally very hard.

1.3 Isotropic interactions

The models considered in this thesis all have in common that their constituents interact with general isotropic potentials. Here, *general* means that a solution to the model should be able to handle any interaction we throw at it, *potential* carries with it an assumption of pair-wise additivity, and *isotropic* means that the interactions do not depend on orientations.

The fact that our prediction methods do not make any additional assumptions on the interactions is a strength. For example, while many studies of spin models exhibiting striped morphologies exists, Paper I gives the first explanation for general interactions, without depending on a specific one. On the other hand, the assumption of general isotropic interactions being available is a weakness of our direct design method as it typically results in very complex designs. We thus spend parts of Papers III and V reintroducing constraints to simplify the interactions and make experimental implementations more feasible.

Is the restriction to isotropic potentials a major one? Yes and no. The classical formulation of the fundamental forces of physics such as Newton's law of gravitation and Coulomb's electrostatic interaction is indeed in terms of pair-wise forces that only depend on the distances between bodies, not on orientation, and one might expect that we should be able to build up macroscopic models using nothing but such. There are two problems with this.

First, the world is not classical. When quantum mechanics is involved, things become more complicated. Indeed, careful analysis show that even noble gases break pair-wise additivity, exhibiting three-body Axilrod–Teller interactions [2] from the third order expansion of the van der Waals interactions. Second, a tractable description of any interesting system is necessarily coarse-grained, using as building blocks objects that are themselves composites. Interactions between them are typically anisotropic; water molecules are not point particles.

On the other hand, we can often construct *effective interactions* between such objects that are isotropic. For one thing, many colloidal particles of practical importance are approximately spherical. Further, one can show that in any system interacting with isotropic interactions on the micro-level, it is at least formally possible to integrate out fast degrees of freedom to obtain an effective Hamiltonian which only involves slow degrees of freedom [3]. This is often used in soft matter physics to either exactly or approximately calculate effective interactions, a generally successful approach.

On the third hand (contrary to Truman's wishes, a good scientist always has a spare one), there certainly are cases where such descriptions are not practical to construct or does not tell us what happens when parameters are varied. This is certainly the case with the interactions between patchy colloids, an approach to self-assembly where the anisotropy is crucial to the design. We will look closer at these in Chapter 4.

1.4 Overview of the thesis

This thesis consists of an introductory text and seven reproduced papers. The introductory remarks of the present chapter are followed by a selective review of pattern formation and self-assembly, the sciences of prediction and design of morphology. We then take a look at the exactly solvable spherical spin model, describing how it has been studied earlier and discussing our contribution. The fourth chapter describes two methods for designing self-assembling systems: one based on the solution to the spherical model, used to design patchy colloids that then assemble into various structures, and a direct method for self-assembling lattices from isotropic potentials.

Chapter 5 gives an overview of the included papers briefly outlined here as well. In Paper I we show how solutions to a generalization of the spherical spin model can be used to explain the ubiquity of striped morphologies and to predict their main feature, their wavelength. In Paper II we adapt the model to describe finite aggregated clusters. This produces a limited alphabet of structures that aggregating isotropic interactions tend to produce and predicts their wavelength as well. Papers III and IV introduce our spectral method for designing isotropic potentials for lattice self-assembly, assembling Kagome, snub square, and snub hexagonal lattices, among others. Paper V describes a method for simplifying such designed potentials, dramatically reducing the number of features and interaction lengths needed. In Paper VI we extend the spherical spin model to handle multiple particle types and spherical geometries. This is applied to a model proposed as a way to self-assemble patchy colloids. Paper VII finishes by using said extension for hierarchical self-assembly: therein we design potentials that cause patches to form on spheres which in turn cause the spheres to self-assemble into various structures.

The main text is concluded with a short summary and outlook in Chapter 6, giving some speculative thoughts about what could and/or should be done in various directions.

Chapter 2

Pattern formation and self-assembly

“Form ever follows function,” cried Louis Sullivan’s credo [4]. It achieved its impact precisely because the immediate causality runs in the opposite direction, from form to function. Indeed, an object’s desired function determines its form through the actions of a designer—be it a craftsman or natural selection—only to the extent that the form facilitates this function.

The study of the appearance and properties of form in nature is the task of the field of pattern formation. Its earliest steps was arguably taken in biology, but today it talks a language inherited from condensed matter physics, a language of instabilities, bifurcations, and spontaneous symmetry breaking. The field of self-assembly, while speaking a different language, can be seen as an effort to use the observations of pattern formation to bring form into being, often on scales or in materials we cannot shape otherwise.

I find it useful to think of pattern formation and self-assembly as twin fields trying to understand and utilize form, respectively. The rest of this chapter gives short introductions to the fields in more classical terms and finishes with a description intended to put the specific subjects of our work into context.

2.1 Pattern formation

Among nature’s most fascinating phenomena is the spontaneous formation of structure. Water gathers into trickles, streams, and rivers, forming complex branching patterns; the wind creates ripples and dunes in sand, resembling waves; and of course life is a matter of ever assembling building blocks into larger complexes.

One might intuitively expect that the formation of complex patterns transpires through complex processes. Especially in the case of patterns in biology, such as the stripes of a zebra or the shape of a snails cone, it is tempting to invoke evolution as an explanation for how this or that structure must be just so.

But it turns out that many simple processes exhibit complex patterns. This shows that no such convoluted explanations are necessary. Further, these patterns tend to follow a large but limited set of motifs, and these motifs are the same that occur in more complex processes [5]. This suggests that there might be common explanations for their existence. In biology such simple, physical explanations are arguably underappreciated, in favor of evolutionary ones that only relocate the question. As famously argued by Thompson [6], many features of the form of organisms are nowhere to be found in their genes but within the laws of physics, determining their growth processes.

Understanding such emergence of structure is the goal of the field of pattern formation. In contrast with self-assembly, it focuses on explaining common or universal patterns with simple models. It thus is the physics to self-assembly's more biological or engineering approach that focuses on designing special systems that exhibit non-typical patterns.

Pattern formation in equilibrium is commonly studied within condensed matter and statistical physics, while the majority of the phenomena studied within the field of pattern formation proper occur out of equilibrium. However, the approaches have much in common, focusing on spontaneous breaking of spatial symmetry.

The basic story, following Cross and Hohenberg [7], is that of growing instabilities. In the beginning there is a spatially featureless system. A blank canvas of disordered magnetic spins, a well-stirred soup of chemicals, or a homogeneous electron gas. However, the rest of the world makes its presence known and small perturbations from this perfect state occurs. These will die out for some values of the system's parameters—be they reaction rates, temperatures, or gradients—leaving the system in a stable state. For other values the system will be unstable and the perturbations will grow, breaking the symmetry and creating a pattern.

Somewhere in between these parameter values is the point, a bifurcation point, where the homogeneous state first becomes unstable and the pattern stable. Bifurcation theory forms the first half of the modern theory of pattern formation. There is seldom only one stable pattern but several, one for each way of breaking the symmetry. This focus on symmetry naturally leads to group theory, the second half of the theory. Combining them gives, among other things, the equivariant branching lemma [8, 9], which makes predictions about the symmetries of patterns given the symmetries of the bifurcations. Importantly, there is a limited number of bifurcation classes, in part explaining why so disparate systems show similar patterns.

There are complications, of course. The unsatisfying one is the fact that bifurcation theory only applies close to the occurrence of the instability, while the similarities between patterns persist far from it. No coherent story patching this gap exists today. More constructive complications, of the job-creating kind, results in for example patterns that can be modulated on long scales, studied as Eckhaus, zigzag, cross-roll or Benjamin-Feir phase instabilities, among others. These can all be studied using various tools, see [10] for an introduction.

2.1.1 Striped patterns in physics

In this thesis the pattern that mostly interests us is what is arguably the simplest non-trivial pattern, the stripe. It is characterized by a single wavelength, with which it varies periodically in one direction while being constant in other directions. It is of course of interest in biology, but many physical systems exhibit stripes too. Interesting examples include lipid monolayers [11], adsorbates on metals [12], and various magnetic fluids [13, 14]. Striped phases are also hypothesized to play a role in the high-temperature superconductivity of transition-metal oxides [15, 16].

These phenomena are well-studied, theoretically typically with spin models involving a competition between short- and long-ranged interactions [17–21]. These interactions are quite specific with physically motivated choices for, *e.g.*, the long range term (Coulomb, asymptotic dipole etc.) as well as the state space (using Ising spin-$1/2$, spin-1 etc.). However, these specifics are unlikely to be important for the good qualitative success of the models and conversely, as noted by Zaanen in the context of Mott insulators [22], their success tells us very little about the original systems' underlying physics. Indeed, a universal phenomenon cannot be used to distinguish between different models exhibiting it.

In this context the value of Paper I is clear: it gives a mechanism for how the universality arises that is independent of specific interactions. Further, it gives a simple method for deciding whether a set of interactions will give rise to stripes: is the minimum of the energy spectrum an interior one or not? Using this method makes results such as stripe formation from purely repulsive interactions [23] unsurprising.

2.2 Self-assembly

Self-assembly is a term that gets applied to a wide range of phenomena, sometimes so wide as to make it effectively meaningless. A common way of restricting it is to let it refer to the assembly of pre-existing components, interacting through non-covalent bonds that form reversibly, into larger, more ordered structures. The components can be not only atoms or molecules, but nanostructures and mesoscale particles as well.

A familiar example is protein folding [24]. Here the components are amino acids, interacting with hydrogen bonds (for secondary structure) or any of a host of weaker ones such as salt bridges, disulfide bonds, and hydrophobic effects (for tertiary structure). By randomly reconfiguring, the protein tries out a long sequence of arrangements, tending to stay in those that minimize free energy. It then ends up in the correct structure, though pushing temperatures outside the intended ranges results in errors being made, possibly ending with denaturalization.

Self-assembly is interesting for several reasons. It is a crucial concept for understanding life, as most structure in cells are self-assembled, including membranes from lipids and the aforementioned protein folding [25]. It is also the way many inorganic structures form from molecules, such as molecular crystals [26, 27], liquid crystals [28] and semi-crystalline polymers [29]. Further, it generalizes to larger scales where completely new material properties become available, facilitating for example the design and creation of photonic materials [30]. Lastly, self-assembly provides a method of manufacturing for length scales between the molecular, where chemists can synthesize almost anything, and the microscopic, which constitutes the lower limit for conventional robotics.

The field of self-assembly began within the context of organic chemistry, studied as a molecular process. However, the interactions between molecules are limited by their atom constituents. There is no way to tune the interactions between say a pair of chlorine molecules, and replacing functional groups in larger molecules may destroy the process altogether. Some degree of tunability is possible for large molecules, where van der Waals, electrostatic, and hydrophobic interactions, among others, may depend less drastically on the included atoms. Still, the switch to supramolecular self-assembly has led to new tools becoming available. For larger structures, from nanoscale and up, entirely new forces are important. These include external electromagnetic fields, capillary and entropic forces, as well as electrostatic and magnetic interactions, and have the important feature of being much more tunable. Larger structures are also easier to study; colloidal particles are easily discernible on micrographs and macroscopic objects can be observed with the naked eye.

A key principle of self-assembly is mobility, below the macroscale supplied by Brownian motion. Anything on the mesoscale and smaller will be highly influenced by ever-present thermal noise at room temperature. This random motion has the advantage of being free; no machines are needed to move components. However, it also has the obvious disadvantage of being undirected. This necessitates the reversibility stated in our definition, as otherwise the components would stick randomly and form unordered aggregates (glasses, diffusion limited aggregates) instead of self-assembling into the desired structures. The reversibility is tightly connected to the focus on non-covalent bonds as thermal motion has to be able to break them, giving us a preferred energy scale around the thermal one of $k_B T$.

Self-assembly is possible on the macroscopic scale as well [31]. However, as the free Brownian motion is lacking, some method of propulsion is needed. This can either be externally provided, an analog of Brownian motion achieved by for example shaking the container, or internal, as is the case with robots such as the swarm-bots [32].

Self-assembly, as a field, is rapidly growing with advances being made in both synthesizing components and facilitating the assembly process. One successful strategy is to mimic nature. Repetitive units of block-copolymers can fold into complex three-dimensional structures, reminiscent of protein folding [33]. Similar strategies have been used to self-assemble microelectronic devices [34, 35]. Given the success of such bioinspired methods, why not take it one step further? Indeed, some of the most advanced methods of self-assembly use DNA directly. Its ability to bind to specific strands can be used to build templates for other self-assembly as well as components directly [36–38]. These are just a few examples of the progress in recent years; for further reading see [39–44].

Chapter 3

The spherical model with arbitrary isotropic interactions

Papers I, II, and IV all try to predict what morphologies appear in systems with isotropic interactions. This is done by first formulating a model of the system in terms of some complicated spin or particle model and then showing how that model can be simplified to produce a version of the *spherical spin model.* This model, the topic of the present chapter, turns out to be exactly solvable: we can, given the interactions, say exactly what morphologies appear as its ground state(s). The task then becomes to connect these solutions back to the original model.

This chapter starts with a short introduction to spin models. It then discusses the spherical model at length, highlighting in what senses its different versions are solvable. Following this are some examples of how continuous models are used to approximate discrete ones, a fine old tradition into which our methods fit. The chapter concludes with sections on the predictions of our model, the generalization presented in Paper VI, and the simulations we use to test them.

3.1 Spin models

The original purpose of spin models was to model the ferromagnetic behavior of some metals. All atoms have a magnetic moment referred to as spin. It is a quantum effect but one may classically regard it as the result of the electrons, charged particles, moving in circles around the core, thereby giving rise to a magnetic field. This field is very small and in a normal solid the different atoms have their spin in different directions so no magnetism is observed macroscopically.

In some metals, iron being the most prominent example, the spins can align and give rise to a macroscopic ferromagnetic behavior. The basic observation that a model of magnetism must answer is the following: as temperature rises, the magnetization decreases and then abruptly disappears at the Curie temperature T_C.

To study this behavior Ising [45] used a very simplified model in which classical spins sit on a lattice (similarly to how atoms in a metal sit at specific locations in the crystal) and can only point up or down (instead of any direction in space). These spins are taken to interact only with their nearest neighbors and it costs energy to anti-align with them. The Hamiltonian of the model is thus

$$H = -\sum_{<i,j>} s_i s_j \quad \text{with } s_i \in \{-1, 1\} \tag{3.1}$$

where $< i, j >$ denotes the sum over nearest neighbors.

The ground state of this model is clearly ferromagnetic with all spins aligned, while at high temperatures—with a lot of thermal energy available—the spins point in random directions. The question is when and how the transition between these states occurs. In one dimension the model has zero magnetization for all $T > 0$. Ising threw up his hands and left physics, conjecturing that the model lacked a proper critical temperature in all dimensions. However, years later Onsager solved the two dimensional case in a *tour de force* and showed that it does have a non-zero critical temperature $T_c = 2/\log(1+\sqrt{2})$. It also turns out to capture a lot of the interesting critical behavior of magnets.

In crystals with strong anisotropies, restricting the spins to lie in one direction, this description is adequate. In general in our three dimensional world, one should consider so called Heisenberg spins, which classically lie on the surface of a sphere. If one's system of choice has anisotropies in only one direction, one get the XY-model with spins on a circle. These are all instances of $O(n)$ models, named after their symmetry groups. The spherical model can be regarded as such a model in the limit where n goes to infinity [46].

Spin models were created as models of ferromagnetism but their use has by no means been limited to that [47]. The Ising model in a field can be shown to be equivalent to a lattice gas. Further, by changing the interactions from aligning to anti-aligning, achieved through a change of sign in Eq. (3.1), the models become antiferromagnetic, especially interesting on lattices with triangular geometry as their inherent frustration give rise to complicated behavior. By combining the Ising ferromagnetic Hamiltonian with an anti-aligning long range interaction, the models can be used to study the formation and melting of striped morphologies. In spin glasses, used to model disordered systems, the possibly non-local interactions are taken to be fixed but randomly chosen from some ensemble.

These generalizations are especially interesting to us as the strength of our approach is its ability to handle general interactions, as long as they are isotropic. Thus it does not apply to spin glasses where the interactions vary randomly, but it works surprisingly well for the stripe models, as seen in Paper I. Note that the cases we consider are intermediate between the original Ising model, where the ground state is trivial, and the spin glass models, where determining the ground state is an NP-complete problem. For models with isotropic but non-trivial interactions, the ground states are just complicated enough to be determinable but still interesting.

3.2 The spherical model

The spherical model is a somewhat different spin model, in that the spins take any real value, $s_i \in \mathbb{R}$, with the normalization $\sum_i s_i^2 = N$. This normalization gives an expected squared value of 1 for the spins, comparable to the Ising model, but through one global constraint instead of N local ones. This makes it a great deal easier to work with and the original spherical model, with nearest neighbor interactions, is exactly solvable in the sense that one can calculate the partition function exactly [48]. This means that one can obtain its behavior for any temperature.

We will follow Berlin and Kac [48] in regarding the spherical model as an approximation to the Ising model. As such, it is unphysical in that the normalization effectively couples all spins, thus involving infinite interactions, despite making perfect sense mathematically as a relaxation of a constraint in an optimization problem. However, there is another class of spin models, the n-vector model, with only nearest neighbor interactions of which the spherical model is a limiting case [46]. The spherical model is thus considered a physically acceptable model of critical behavior.

In our work we have considered generalizations of the original model, with arbitrary isotropic interactions. It has real-valued spins s_i situated at positions r_i on some lattice and interacting with potential $V(r)$. The Hamiltonian is then on the form

$$H = \sum_{ij} V_{ij} s_i s_j \tag{3.2}$$

where $V_{ij} = V(|\vec{r}_i - \vec{r}_j|)$ is a matrix describing the interactions.

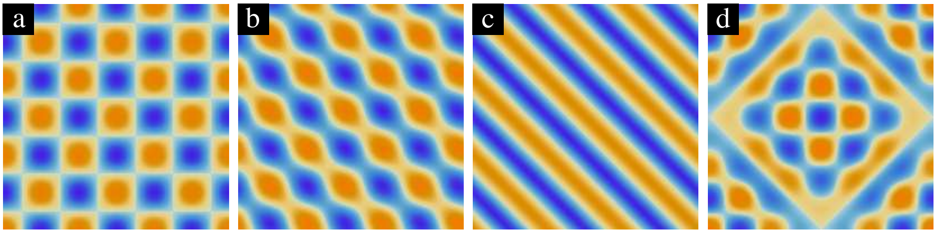

Figure 3.1: *Ground states of the generalized spherical model on the* $\mathbb{Z}^2$*-lattice. Due to the independence of phase, in two dimensions there is a continuous degeneracy, ranging from checkerboard patterns to stripes (a–c). For states with different wavenumbers in the different directions, linear combinations can be used to construct complex morphologies (d). This degeneracy is broken when going to discrete spin models, and only the striped states remains.*

While it is still unclear whether this generalization admits to an exact solution for the full partition function, it turns out to be possible to find its ground state(s) exactly, without making assumptions on the interactions. As we show in Papers I and II, the matrix V_{ij} is diagonalized by the eigenfunctions of the Laplacian.

In Paper I we consider the $\mathbb{Z}^2$-lattice (the square lattice in two dimensions), in which case the eigenfunctions are the Fourier modes. To find the ground state, we thus have to find the Fourier mode with the lowest energy. The easiest way in practice to do this is to take the Fourier transform of the potential, giving the energy spectrum for the given interactions, $E_{\vec{k}} = \mathscr{F}[V(r)]$, and then identifying the lowest point k_{min}. This gives us the ground state of that potential as the Fourier mode with wavenumber k_{min}.

The minimum is in general degenerate; in two dimensions $\vec{k} = (k_1, k_2)$ has the same energy as $(\pm k_1, \pm k_2)$ as well as $(\pm k_2, \pm k_1)$. The Fourier modes of all these are eigenstates, as well as linear combinations of them, illustrated in Figure 3.1. In the continuous limit we also have a degeneracy with respect to any changes to $\vec{k}$ that leave its norm unchanged. This latter degeneracy is a natural result of the rotational invariance of the system. However, taking linear combinations of solutions can lead to some very complicated ground states, resulting in an embarrassment of riches not observed in other spin models. In Paper I we show that these unwanted degeneracies are naturally broken when going to discrete models, leaving only the simplest, striped states.

The above solution goes through in other geometries as well. In Paper II we study solutions to the same Hamiltonian in cylindrical coordinates. This model can be used to study systems with broken translational invariance, where there is a

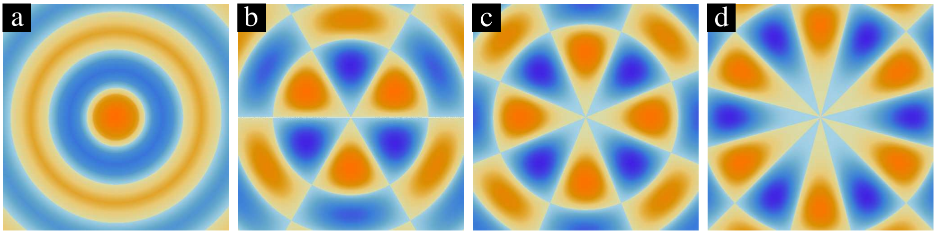

Figure 3.2: *Rotationally invariant ground states of the spherical model, relevant for aggregating systems. All the displayed states (with* $\omega \in \{0,3,4,6\}$ *(a)–(d), respectively) exist and are degenerate regardless of interactions, as well as states with other* ω *and linear combinations of the same. The* ω*-degeneracy persists when going to discrete systems but the relevant linear combinations get fixed.*

natural center of the system. In this geometry the eigenfunctions change from the Fourier harmonics to cylindrical harmonics, *i.e.*, Bessel functions with angular modulations $J_\omega(2\pi kr)\cos(\omega\theta)$. These can be used to describe finite aggregates. Such arise from potentials that in spin systems result in ferromagnetic states, but in particle systems with a finite number of particles the result is a concentration of the particles to a finite region, an aggregate, breaking the translational symmetry. Aggregating potentials can be identified in the spectrum view by their global energy minimum at $k = 0$.

The cylindrical solutions also contain degeneracies, here in ω as the energy only depend on k. Some examples are shown in Figure 3.2. In this case the embarrassment persists; we cannot find reasons to prefer any values of ω above others. However, as shown in Paper II, this degeneracy is real; simulations of discrete models have a weak tendency to select for certain low values of ω but switch between them even for a single set of interactions as the density is varied.

Lastly, in Paper VI we consider solutions on the surfaces of spheres. Here the eigenfunctions of the Laplacian are the spherical harmonics Y_l^m. These are again degenerate, here in m (see Fig 3.3), but for the same reasons as for the $\mathbb{Z}^2$ lattice the degeneracy is broken in discrete models, resulting in striped or spotted states depending on stoichiometry.

3.3 On connecting continuous and discrete models

While the spherical model is interesting in its own right as an exactly solvable model, in our work we are more interested in the discrete Ising spin and related particle models. The goal is then the ground state of such a model and the spher-

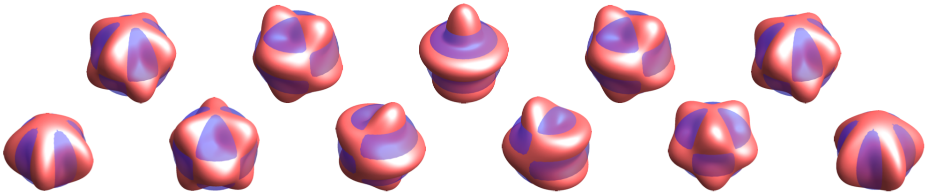

Figure 3.3: *Ground states of the spherical spin model on the surface of a sphere. Shown is the real part of spherical harmonics with* $l = 5$ *and* $m \in \{-5, -4, \dots, 5\}$ *as deviations from a spherical surface.*

ical model a tool for obtaining it. One way of making the connection precise is given in Paper II by considering the Ising spins as continuous variables in the spherical vein with constraints $\sum_i s_i^m = N \; \forall m$. We then get the spherical model by relaxing all but the first two constraints.

This practice—of using continuous models to find solutions to discrete ones—is one that pops up in various places in the literature. It is a way of relaxing constraints, which in its most general form is well-known within optimization theory as penalty methods. There, a variable that in the original problem is restricted to some subset I of $\mathbb{R}$ is allowed to vary outside of its bounds, but with a cost that is slowly raised in the course of the optimization to finally arrive at a solution inside the feasible set [49].

Finding the ground state of the Ising model is basically a combinatorial problem; finding the configuration(s) lowest energy out of the 2^N possible ones. Relaxing a combinatorial problem into a convex one is another instantiation of the general strategy, used for example to recover spike trains [50, 51] and sparse images from a few Fourier modes [52] by minimizing l_1-norm.

For nearest neighbor ferromagnetic interactions the Ising ground state problem is trivial; for a general spin glass with Hamiltonian

$$H = \sum_{ij} J_{ij} s_i s_j,$$

without restrictions on J_{ij}, it is an NP-complete problem in the non-planar case [53]. Using eigenvector analysis in the same vein as ours, Aspvall and Gilbert studied graph coloring [54], a similar problem. Partitioning networks into modules with minimal intra-connectivity is, in the base case of two modules, isomorphic to the spin glass problem and has also been handled using this strategy [55, 56]. The reason such simple methods often work is that NP-completeness is a worst case statement, while in the vast majority of situations the average case is much nicer.

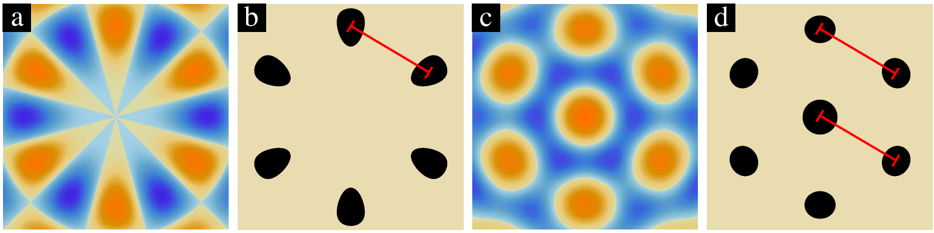

Figure 3.4: *For the aggregating solutions, it is not obvious what base to use and linear combinations of Bessel harmonics may be in some senses simpler than pure ones. For* $\omega = 6$ *(a), a mixing (c) with the corresponding* $\omega = 0$ *state results, after mapping (b and d respectively), in a particle state with a larger fraction of the interactions occurring at the dominating length scale (illustrated in red) due to the extra cluster in the center.*

3.4 Predictions for the discrete models

The solution of the spherical model shows us that the only effect of the interactions is to set the length scale of the morphology, by defining a wavenumber k_{min} through the energy spectrum. All other features are degenerate and has to be determined, if at all possible, from the ways in which our system of interest departs from the model.

So what do we expect to happen to our exact solutions for the spherical model when we go to the arguably more interesting discrete ones? First of all, as alluded to earlier, we expect a great deal of degeneracy to be broken. The continuous model allows for smoothly varying states. Discrete models lack this possibility, introducing sharp interfaces.

For an Ising spin-$1/2$ model the state corresponding to a certain continuous solution is obtained by simply mapping positive states to $+1$ and negative to -1. This will always increase the energy and the increase will be largest at the interfaces, where the departure from the original state is largest. This makes us conclude that the states with minimal interfaces will be the ones observed for each degenerate set. These are the striped ones.

Turning to aggregating particle systems, the situation is more complicated. We still expect simple states to be preferred, but a way of making that notion exact has so far eluded us. Indeed, it is not obvious what constitutes a simple state here as simple surface minimization does not work. Linear combinations of Bessel harmonics are for some ω simpler in the sense that they are to a higher degree dominated by a single distance. Figure 3.4 illustrates the idea.

To explore the results of mapping the spherical states to particle states, we instead turn to numerics. The mapping contains a free parameter c; all states larger than c is mapped to 1 and the rest to 0, with c implicitly defining the density of the resulting particle system. We sweep this and the other parameters of the solutions and compute the energy of the resulting particle states for a large number of random potentials. This allows us to observe which states tend to be favored in general. The result is a morphological alphabet (Figure 3.5): a large but limited set of states with much stronger signals than the rest. This alphabet is our main prediction for aggregating particle systems. Notably, it contains versions of most discrete rotational symmetries, though some are more common than others, indicating that the ω-degeneracy persists.

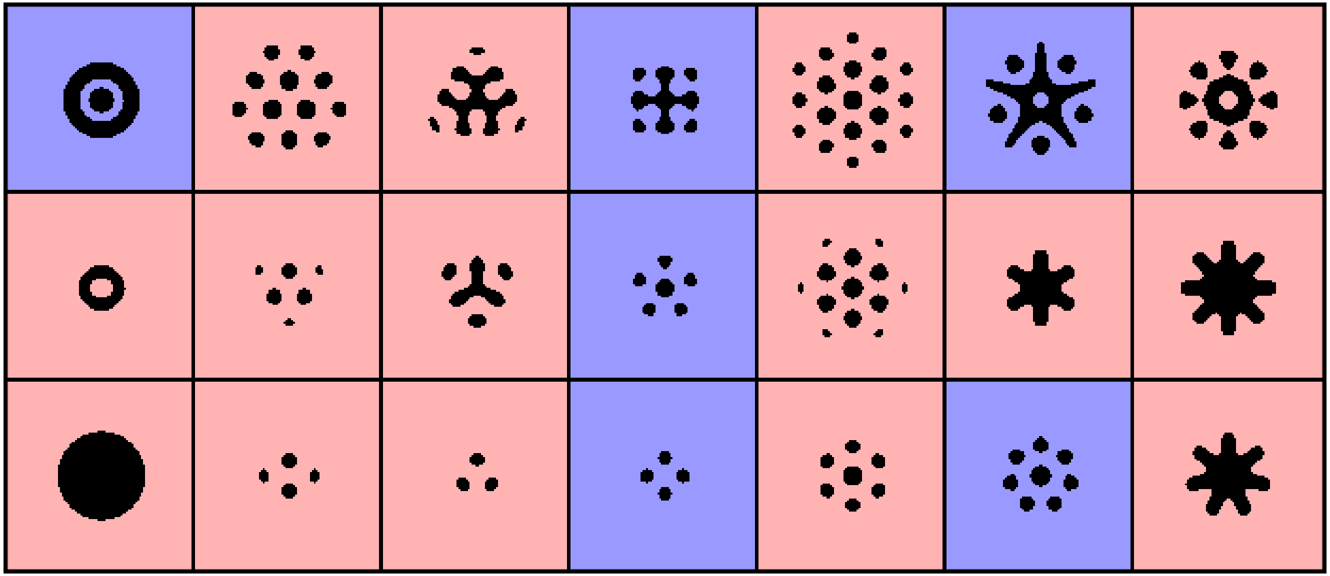

Figure 3.5: *Predicted morphological alphabet for aggregating particle systems, obtained through numerical study of the effect of the mapping on the spherical model. The colors indicate relative strength of signal, with red being predicted to show up more often in particle simulations. The shown alphabet is limited in the rotational symmetry ($\omega \leq 8$) and number of radial oscillations of the morphologies considered. The analysis can be extended but that yields no qualitatively new shapes.*

3.5 Generalizing to multiple particle types

In Paper VI we generalize the spherical model to be able to predict the behavior of systems with multiple particle types. This is a separate question from that of the geometry; in the paper we study spherical geometries but here we will use the language of flat space for simplicity.

The discrete particle system we would like to study has Hamiltonian

$$H = \sum_{\alpha\beta}^{K} \sum_{ij}^{N} \Pi_{i\alpha} V_{ij}^{\alpha\beta} \Pi_{j\beta}, \tag{3.3}$$

where K is the number of particle types. Here we have one interaction potential $V^{\alpha\beta}$ for each pair of particle types α and β, and the state is now described by a matrix $\Pi_{i\alpha}$ that takes the value 1 if site i has a particle of type α, and 0 otherwise. The relaxation into the spherical model then lets Π take any values under only the constraint $\sum_{i\alpha} |\Pi_{i\alpha}|^2 = N$.

The solution of this model goes in two steps. First we follow the same procedure as before, allowing us to diagonalize the interaction matrices $V^{\alpha\beta}$ independently, giving us pair-wise spectra that only take into account the interaction between two particle types at a time. From these we construct a set of small matrices $\Lambda_{\alpha\beta}^{k}$, where k now label the matrices according to which wavenumber they correspond to. These are $K \times K$-matrices and thus simple to diagonalize, for three particle types it is possible to do exactly but in practice it is simpler to do numerically. This diagonalization gives us a new spectrum, with K branches, but also a set of eigenvectors, one for each point in this spectrum (Figure 3.6).

The construction of the ground states of the model involves both the eigenvalues and the eigenvectors. As before, the length scale of the ground state is given by the wavenumber of the lowest point in the spectrum. The eigenvector corresponding to this point gives us weights and phases for the different particle types. For example, if we have three particle types and get a minimum with an eigenvector $(\sqrt{2}, -\sqrt{2}, 0)/2$ the ground state has Fourier modes of particles of type 1 and 2 out of phase and no contribution from type 3. In other words, we construct the columns of Π by taking Fourier modes of wavenumber k_{min} multiplied by the elements of these small eigenvectors.

To map these solutions to predictions for the discrete Hamiltonian is a bit more complicated as well. In essence, this is because each branch of the spectrum only describes how one group of particle types varies in relation to the rest (or no variation in the case of eigenvectors with the same sign on every element). It turns out that to get good predictions, we must look at the minima of several

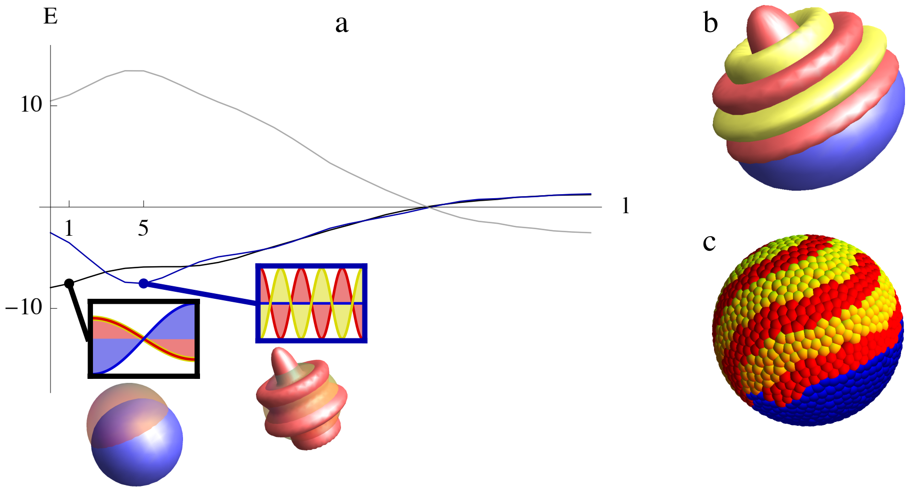

Figure 3.6: *(a) Spectrum of the spherical spin model for three particle types on the surface of a sphere. Minima of the relevant branches are illustrated with spherical harmonics. (b) Combining these spherical harmonics gives a prediction for the particle system. (c) The predictions match the results of Monte Carlo annealing well when done with the same interactions as those used to calculate the spectrum.*

branches. For example, Figure 3.6 shows a spectrum where the global minimum only tells us that the blue particles separate from the rest, and we have to look at the minimum of the next branch to see what happens in the other domain. In this way, the predictions for systems with multiple particle types typically have several length scales, one for each involved minimum.

3.6 Monte Carlo simulations

The main value of the work presented in this thesis lies in its theoretical nature. The field of self-assembly abounds with experimental and simulation based work, but theoretical results are much more far between. The proof of the pudding is in the eating and theoretical work needs to be relevant to be interesting. While the ultimate arbiter of science is experiment, the first testing grounds for theories are provided by computer simulations of less simplified models.

In our work we use simple (Markov chain) Monte Carlo simulations. The basic premise is to start from some initial state and then do random alterations to this state. These alterations are usually local[1] and consist of, for example, spin flips in an Ising model or a move of a single particle in a particle simulation. If a move results in a reduced energy of the system, it is accepted unconditionally. If it instead would increase the energy, it can nevertheless be accepted; it is with a probability that decays exponentially in the energy required, $p \propto e^{-\Delta E/k_B T}$. This simple rule, known as Metropolis–Hastings', ensures that in the long run the states visited are sampled in proportion to their contribution to the thermodynamics of the simulated model. Of course, the long run can get long indeed, especially if your task is to estimate noisy order parameters in a system with meta-stable states.

Thankfully, we do not have to suffer this as we mostly are interested in ground states. On the other hand, these have their own issues. Real proof that a certain state is the ground state of your potential of choice is generally not available and you have to settle for the "beyond reasonable doubt" kind. This is usually obtained by multiple runs of simulated annealing where the above Monte Carlo algorithm is run while slowly lowering the temperature. This can be combined with comparisons to energy calculations for one or more hypothesized perfect patterns or lattice structures, as applicable.

In the papers we use these methods to test the predictions of the spherical model. We find remarkably good agreement, considering the simplicity of the model. In general, the length scale determined by the energy spectrum predicts the length scales of the simulated morphologies well. We have not made extensive numerical tests, but in Paper II we show the results for a typical potential.[2] The measured k agrees to within a couple of percent with its predicted value, which is within the uncertainty of the measurement. However, the errors do have a bias, suggesting more analysis might be fruitful.

For the stripes similar results hold, with the caveat that the lattice spacing has to be scaled relative to the potential so that the resulting stripes do not span too many lattice sites laterally. With increasing stripe width the underestimation increases and the asymptotic scaling is wrong. For example, a scaling analysis of the potential used in [20] shows our predicted scaling to be linear while the correct one is exponential.

Regarding the other features of the morphology, we do indeed observe stripes, and only stripes, in the spin models. This is also consistent with the existing literature, for the different specific interactions considered. For the aggregating particle systems the predictions also play out well, with almost total overlap between

---

[1] Though global moves, such as flipping large clusters of spins or moving whole collections of particles, can often greatly speed up convergence.

[2] The potential is chosen for its diversity of morphologies, not for goodness of fit.

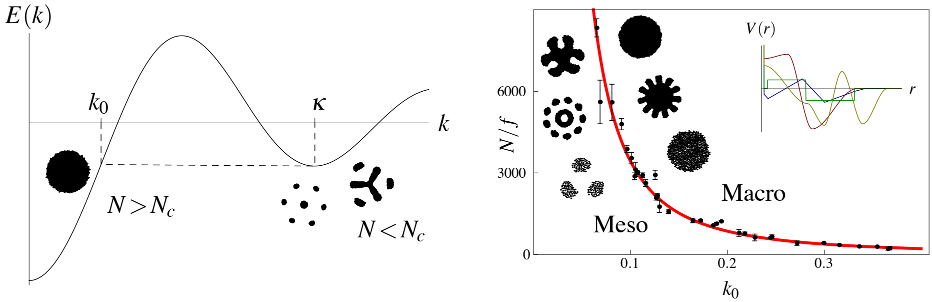

Figure 3.7: *Predicting the transition from disks at the macroscale to more complex morphologies at the mesoscale. (Left) Aggregating potentials have their global minimum at* $k = 0$*, so at large particle numbers their preferred morphology is a disk. However, if the particle number decreases below some critical point* N_c *there is no longer enough particles to occupy such a state and more complex morphologies appear. This point can be predicted as the number corresponding to a* $k = k_0$ *such that* $E(k_0) = E(\kappa)$*, where* κ *is the location of the interior minimum of the energy spectrum. (Right) Results of particle simulations to determine* N_c *(adjusted for packing fraction* f*) for potentials with various* k_0*. The red line is the parameter free prediction. (Inset) Examples of randomly generated potentials used.*

the predicted alphabet and the observed morphologies.

Paper II presents an additional result, a phase diagram showing the particle numbers at which the complex morphologies appear, as opposed to the simpler disks common at large particle numbers. Figure 3.7 demonstrates the prediction and its fit with our data. This serves as another demonstration of the strength of the spectral view for these systems. It also constitutes an implicit test of the length scale predictions.

The predicted ω-degeneracy plays out in a surprising way. It would not be convincing if all we could say was that we cannot find a feature of the potentials that predicts ω. However, fixing the potential and simulating systems with different particle numbers shows a switching back and forth between different values of ω. This shows that the breaking of the degeneracy is a complicated matter, one we cannot hope to disentangle with such a simple model.

In Paper VI we construct a simplified model of alkanethiol-on-gold systems. These are experimental systems involving gold nano-particles coated with hydrocarbon chains that form patterns due to competition between phase separating and mixing interactions. This model has many transitions between different patterns, and these turn out to be reflected in the energy spectra allowing us to predict them.

Interestingly enough there are several ways this happens. The simplest is that the minima moves and thus changes wavenumber. There are also times when there is a change in which branch has the global minimum. In one example this gives a change from a state dominated by a blue domain with red-yellow stripes on the other side to a state dominated by red spots with a blue-yellow phase-separation between the spots (see Figure 4b and c in Paper VI).

Lastly, the sign structure of the eigenvector of a branch minimum involved in the solution can change. Most dramatically the global minimum can change into an unphysical state which does not describe variation at all, with all elements of the eigenvector having the same sign. We then have to use another minimum to get enough variations to predict a full state. While one might suspect that this is an artifact of the method and expect that the predictions would be unreliable in such cases, the simulations show that the predictions are correct.

Chapter 4

Designing interactions for self-assembly

Broken symmetry is one of the central concepts of modern physics and a fruitful one for understanding the appearance of non-trivial morphologies in systems with isotropic initial conditions. In the previous chapters we have seen many examples of this, in the appearance of a macroscopic magnetization in the Ising model, in the growth of Fourier modes in pattern formation, and the breaking of the degeneracy of the ground state of the spherical model.

However, as Anderson put it, "[a]t some point we have to stop talking about decreasing symmetry and start calling it increasing complexity" [57]. This chapter crosses this line. We now turn from the prediction of patterns that typically appear in different systems toward the design of very special systems with much more complex morphology. We will look at two different approaches to design, one explorative that uses the prediction methods of the previous chapter and one direct that in its purest form gives a potential fulfilling the desiderata without parameter tuning.

The task of design is in essence a search problem: find a description which we can follow to construct a desired structure. The point of directed self-assembly is to move as much as possible of this construction into the design, simplifying the former at the cost of the latter. Design methods, as any methods for solving search problems, exist on a scale. On one end there is the absence of method, trying solutions purely at random. At the other are ideal methods that require no intelligent input and minimal calculations. In between we have a sliding scale of heuristics, semi-automated methods, and simulation based approaches.

In the next section we will look at a method that combines heuristics, some hand tuning, and the prediction methods of the previous chapter. The heuristics give us desired patterns on spheres, which if we can achieve them in colloidal systems would cause self-assembly into target structures. These patterns could self-assemble themselves if their constituents have interactions whose spectra have certain properties. We then use our prediction method to allow us to tune parameters in our model until the desired spectral properties are achieved.

We then look at a direct design method for lattice structures. Given the target structure, it specifies an energy spectrum and inverts it to produce an isotropic potential whose ground state is that target. This is a much faster method than the previous one, but the power comes at the price of experimental realizability: the potentials produced are much too complicated to implement in the experimental systems we have today and probably in the foreseeable future. We thus spend the last section discussing methods for simplifying the designed potentials.

4.1 Designing self-assembly of and with patchy colloids

Patchy colloids are colloidal particles (10 nm-1 μm sized) with patches that give them directional interactions. This makes them resemble atoms, leading to talk about creating a colloidal chemistry, but with much higher theoretical freedom in designing the interactions; in principle a higher diversity in form should be possible, and we should for example be able to turn interactions on and off by changing properties such as salinity of the surrounding solution. However, while many methods for fabricating patchy colloids exist, it is still hard to do so in sufficient quantities for self-assembly, and more complex patch arrangements are just beginning to come within reach. Thus so far only relatively simple structures, using cleverly designed simple patchy colloids, have been self-assembled.

As an example, consider the self-assembling Kagome lattice design by Granick and coworkers [58]. By creating caps of gold with glancing angle deposition on latex particles and then coating them with alkanethiols, they create particles with an attractive patch on each pole. The scheme is chosen to create caps of a particular size. This causes the favored arrangement of the particle to be triplets with inward facing caps, and, as this configuration is characteristic of the Kagome lattice, causes self-assembly into this two-dimensional structure after sedimentation. Note that the naive patch arrangement for Kagome assembly would be one with a patch in the direction of each neighbor.

A proposed method for creating patches on colloids is to self-assemble the desired pattern itself. This has been done for simple striped patterns in alkanethiol-

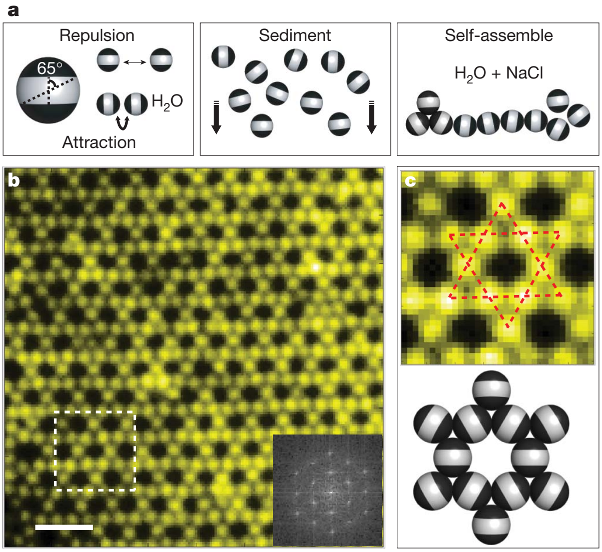

Figure 4.1: *Self-assembly of a colloidal Kagome lattice. Latex particles with hydrophobic patches form triplets that arrange into the Kagome structure under sedimentation. Copyright NPG. Reproduced with permission from [58].*

coated gold nanoparticles, and in Paper VI we used a model of this system as our test case for our prediction method. In Paper VII we use this method to design systems exhibiting hierarchical self-assembly, where the first the building blocks self-assemble and then those building blocks assemble into the desired macro structure.

We assume that if we can create a patch structure within the alkanethiol-on-gold model we can then make one or more of the particle types attractive. The design process thus consists of choosing a patch structure that causes assembly into the target structure, finding a way to implement it on the surface of a sphere using spherical harmonics, and then tuning the parameters in the alkanethiol-on-gold model until a spectrum with the correct branch minima is achieved.

For a simple example, consider the task of assembling a diamond lattice. A sufficient patch structure would be one with four patches equally separated on the sphere, analog to the bonds of the carbon atom in a real diamond. As the simplest pattern exhibited by the $l = 3$ spherical harmonic is such a pattern (given a limited stoichiometry to result in patches rather than stripes), we just need to find parameters of the model that gives an energy spectrum with a minimum at this point. By choosing the stoichiometry we can then make sure we get patches of a size that gives reasonable flexibility in the bond formation (Figure 2a in Paper VII).

More complex structures requires more complicated constructions. For example, the 90° separation between patches needed to assemble cubes does correspond to the $l = 4$ spherical harmonic, but this would give a cubic lattice. If we want cubes, we need to make sure no patches form on the backs of the colloids. Here the solution involves three particle types and three minima in the spectrum, the extra type forming a domain on the back to restrict the area available to the rest (Figure 5a in Paper VII).

4.2 A direct design method for lattice self-assembly

There exists several schemes for designing isotropic potentials that have specific lattice structures as their ground state [59, 60]. These typically take some family of potentials, an initial guess based on the different nearest neighbor distances of the target lattice, and proceed to optimize the interactions according to some criterion to try and achieve the correct structure. Through such schemes several isotropic potentials with low-coordinated lattice structures have been reported, *e.g.*, square, honeycomb, simple cubic and diamond lattices [61–63]. One of the strengths of this approach is that other properties can be optimized for as well, such as the target being ground state over as large a density range as possible [64]. These methods, however, tend to be somewhat unwieldy, requiring Monte Carlo simulations of the designed system for each step in the optimization process. Indeed, all the lattices mentioned above are fairly simple compared to what is possible.

The key property of lattices is periodicity. Thus it is natural to study them in reciprocal space. Indeed, many first courses in solid state physics begins with a thorough primer on various reciprocal space concepts.[1] Despite this, methods for designing interaction potentials $V(r)$ for lattice self-assembly have so far worked in real space.

We focus instead on designing an energy spectrum $E(k)$ that is then, through a simple Fourier transform, turned into a potential. If we describe a configuration as a sum of delta functions $\rho(\mathbf{r}) = \sum_i \delta(\mathbf{r} - \mathbf{r}_i)$, we can write the energy as

$$E = \int d\mathbf{r}d\mathbf{r}'V(|\mathbf{r} - \mathbf{r}'|)\rho(\mathbf{r})\rho(\mathbf{r}') = \int dk\widehat{V}(k)|\hat{\rho}(k)|^2 \tag{4.1}$$

---

[1] Whether this is a good order for teaching the actual physics is less clear; for a counterexample see the excellent Ashcroft/Mermin [65].

where $\hat{\rho}(k)$ is related to the structure factor of the configuration and $\widehat{V}(k)$ is the energy spectrum, that is the Fourier transform of the potential.

The key idea is to match local minima of $E(k)$ to the magnitudes of the dominating wavevectors of the lattice's power spectrum. From Eq. (4.1) we see that if we construct $\widehat{V}(k)$ such that it is strictly positive everywhere except at the points where $\hat{\rho}(k)$ of the target lattice is non-zero and zero at such points, the energy of the target lattice will be zero and any structure which has contributions from other k will have a non-zero energy. This means that the target will be a ground state, and places strong restrictions on what any competing ground state could look like.

For complex lattices there generally are such competing lattices. For example, the structure factor of the Kagome lattice studied in Paper III (Figure 1) has the same support as that of the closely packed lattice, meaning that this first construction normally would cause collapse into the latter, rather than self-assembly into the target Kagome lattice. However, even though the support may be the same different lattices have different magnitudes of the structure factor. Thus we complete the method with a small perturbation at one or more points, letting the spectrum become negative at the points where the structure factor is maximal for the target. This turns out to solve the problem.

Note that from our construction we can prove that the target is indeed a ground state. This is interesting in light of the fact that it is normally very hard to determine rigorously the ground state of a given potential as the differences in energies between competing structures are very small compared to the total energies involved. Indeed, the exact nature of the ground state of even the simple Lennard–Jones potential is not rigorously known [66]. Somewhat counterintuitively, with our method in the toolbox it is easier to construct a potential with a specific lattice as a ground state than to identify the ground states of even very simple potentials.

4.3 Simplifying designed potentials

Although matching wavevectors leads to methods fulfilling the requirement of producing potentials that cause self-assembly of desired lattices, naive versions horribly fail another reasonable requirement: we want the produced potentials to be simple enough to possibly realize in an experiment. While it is hard to say what are the exact capabilities today and where they will be in the near future, it is part of due diligence as a designer of self-assembly potentials to simplify them at least as much as is easily achievable.

There are two obvious directions from which to approach this problem, the experimental and the theoretical. Starting from a specific experimental setting, we might try and identify the interactions we have at our disposal, look at their spectra and try to combine them into one that fulfills the prescriptions. As combining spectra is a linear operation, optimizing in such a setting is perfectly feasible and the main problem lies in identifying the interactions, as well as assuring that they are sufficiently varied for us to reach our target.

The alternative is to begin on the theoretical side. Starting with the complicated designed potential, we might try to simplify it into something more feasible. Without reference to a specific experimental system it is hard to determine what is a complex potential, but in general we might expect potentials with fewer features and shorter range to be easier to implement. We thus focus on smoothing and screening potentials.

With screening we mean a transformation $V(r) \to \xi(r)V(r)$ such that $\xi(r)$ is some positive decaying function. We can see that such a transformation correspond to a convolution of the spectrum $\mathscr{F}[\xi(r)V(r)] = \int d\mathbf{k}'\hat{\xi}(|\mathbf{k}-\mathbf{k}'|)\widehat{V}(|\mathbf{k}|)$. The uncertainty relation for Fourier transforms then lets us place constraints on the simultaneous localization of ξ and $\hat{\xi}$:

$$\mathrm{Var}[\xi^2]\,\mathrm{Var}[\hat{\xi}^2] \geq \frac{1}{4\pi^2}.$$

$\mathrm{Var}[\xi^2]$ is here the amount of screening performed by ξ and $\mathrm{Var}[\hat{\xi}^2]$ should be interpreted as the loss of distinction the same screening costs. In other words, this places restrictions on the amount of screening we can achieve without destroying the properties of the potential in reciprocal space that we designed for in the first place. As the lower bound is achieved for Gaussians, we thus propose to use the equivalent of a Gaussian screening, namely heat kernel smoothing in reciprocal space: $\widehat{V}_\tau(k) = \int d\mathbf{k}' e^{-|\mathbf{k}-\mathbf{k}'|^2/\tau}\widehat{V}(|\mathbf{k}'|)$, where τ describes the degree of smoothing. This was for screening in real space; an equivalent derivation leads us to use Gaussian screening in reciprocal space to achieve smoother potentials. Figure 2 of Paper V gives nice demonstrations of the simplifications possible.

Chapter 5

Paper overview

Reading a body of work is much easier with a roadmap. This chapter gives an overview of the papers contained in this thesis, with an emphasis on the connections and similarities between the papers.

In Papers I, II, and IV we study morphologies that arise from isotropic potentials in various geometries. They all build on the same idea, that we can formulate the problem of predicting the morphology as a question about the low-energy states for some spin model. This spin model can then be approximated with one with relaxed constraints, giving the spherical spin model. We show that we can solve for the ground state of the spherical model for arbitrary potentials. The different papers then explore the implications in various geometries and, in Paper VI, for several particle types, and show that the approximation works very well for predicting the morphologies of the original systems.

Papers III-V concern the design of potentials to cause self-assembly into specific morphologies. Paper III describes and Papers IV and V extends a method for designing potentials for targeted assembly of crystal structures (lattices). The method works in reciprocal space, utilizing the fact that lattices are most simply described there.

Paper I: Universality of Striped Morphologies

Paper I tries to explain the ubiquitousness of stripe morphologies. In a sense, the reason is simple: a stripe is simply the least complicated way in which translational invariance can be broken and thus typical for systems where this occurs.

On the other hand, seeing how this occurs in an exactly solvable model is instructive; it forces us to make this vague reasoning exact and demonstrates subtleties.

The main theoretical result of the first paper is thus the exact solution for the ground state of the spherical model. This result is something we stumbled over while studying spin glasses and the paper was originally intended to be a short detour, a detour which grew into this thesis.

The basis for the result is that, for any isotropic interaction, the spherical model is diagonalized by the Fourier modes if periodic boundary conditions are used. This can be proved in several ways, e.g. by using the translational invariance or considering the interaction matrix as a circulant matrix [67]. In this paper we prove that any pair of such interaction matrices commute and note that the Laplacian is such a matrix, with the Fourier modes as its eigenvectors. For a different perspective, in Paper II we note that any isotropic interaction matrix can be built up by linear combinations of different powers of the Laplacian, thus showing they have eigenvectors coinciding with the generating Laplace operator.

The ground state of the spherical model with a given interaction potential is then given by the Fourier modes corresponding to the minimum of the energy spectrum of the same potential. The spectrum can be simply calculated as the Fourier transform of the potential and only depends on the magnitude of the wavevector. This latter fact leads to degeneracy of the ground state with checkerboards, stripes, and linear combinations of the two as well as more complicated states all having the same energy.

In Paper I we then connect this result to the behavior of Ising spin models with general isotropic interactions (in two dimensions for concreteness). The mapping from the continuous spins of the spherical model to the discrete Ising spins is nonlinear and we show that this resolve the above mentioned embarrassment of riches and singles out striped states as the ground states. Monte Carlo simulations show that this prediction is indeed correct for any potential we test which has an interior minimum in the energy spectrum. If the minimum lies at $k = 0$, we effectively have an infinite stripe width, corresponding to a ferromagnetic state.

The location of the minimum predicts the width of the stripes and the simulations bears out this prediction. One point worth noting, not central to and thus not present in the paper due to space limitations, is that the prediction works best when the stripe width is not too large compared to the lattice spacing. This can be understood by comparing the shapes of a sine wave (the solution in the spherical model) and a step function (the shape of the stripes in the Ising spin models). One should thus be careful in trying to predict the scaling of the stripe width with respect to parameters in the potential if the predictions grow too much.

Paper II: Novel Self-Assembled Morphologies from Isotropic Interactions

In this paper we turn our attention to a subtlety of the solution in Paper I: the eigenfunctions of the Laplacian, which we showed determines the ground state of the models of interest, depend on the boundary conditions. With periodic boundary conditions comes the translationally invariant Fourier modes; when we instead consider functions decaying to zero at infinity the translational invariance is spontaneously broken and we get localized solutions of Bessel or spherical harmonics type.

Paper II explores the implications of this result for aggregating systems in two dimensions; that is, systems of particles that for low temperatures tend to gather into a cluster. The theory turns out to work very well. Looking at the energy spectrum of a potential we can predict whether it will aggregate or not, the length scale of the aggregate as well as a transition from a disk at high particle numbers to more complex Bessel morphologies at lower particle numbers. These morphologies are well described by the localized eigenfunctions of the Laplacian, on the form $J_\omega(2\pi kr)\cos(\omega\theta)$ with $k \in \mathbb{R}$, $\omega \in \{0, 2, 3, ...\}$.

One simple sufficient condition for an aggregating potential is that its spectrum has a global minimum at $k = 0$. This would mean a ferromagnetic ground state in a spin model and for a finite number of particles it means that they cluster into a disk of as small radius as possible. However, if the spectrum has an interior local minimum as well, more interesting things can happen. As low k means large disks, the actual spectrum accessible to the system is limited from below by the particle number. Thus, for small particle numbers the $k = 0$ minimum is inaccessible and the interior minimum κ governs the behavior instead. In the intermediate regime when the particles are too few for a large disk but too many for a κ-sized one, instead Bessel functions of higher order arise as they allow more particles for a given κ.

This line of reasoning allows us to predict at what number of particles a given potential will exhibit these more complex morphologies from its energy spectrum. It may also be the key to understanding why these morphologies have not been previously observed, neither in experiments nor in simulations, to the best of our knowledge. Most of the potentials we study in Paper II are medium range; they have non-trivial structure on length scales considerably larger than the size of the particles. This is not common in systems usually studied, but is probably achievable in colloidal science today.

Paper III: Designing Isotropic Interactions for Self-Assembly of Complex Lattices

In the first two papers we focused on the energy spectra corresponding to different interaction potentials as they determine the behavior of the spherical model. This focus on spectra lead us to the realization that designing potentials for self-assembly of lattice structures, while not connected to the spherical model, should be done in momentum space as the translational invariance is most easily represented there. Work in this direction lead to Paper III, which presents such a spectral method.

The main ideas of the method is simple. First a target lattice structure is chosen and its Fourier spectrum is identified. One then designs an energy spectrum with minima at the points corresponding to maxima of the lattice Fourier spectrum. Lastly the energy spectrum is transformed into a potential by the inverse Fourier transform. This potential can then be used in, for example, Monte Carlo simulations to self-assemble the desired lattice structure. This allows us to design potentials that cause self-assembly into the Kagome lattice and a truncated square tiling.

There are of course a lot of subtleties swept under the rug in this description. Most important is the design of the spectrum: the mentioned constraints leaves a lot of degrees of freedom for this design. To achieve a spectrum with a corresponding potential that is not too complicated is highly non-trivial. In Paper III we do some *ad hoc* simplifications while a more systematic method had to wait for Paper V, but the seeds in the form of the realization that smoothing and screening are the correct operations is here.

Paper IV: Chiral Surfaces Self-Assembling from Isotropic Interactions

The goal of self-assembly is to not only construct form, but form with specific function. In Paper IV we use the design method of Paper III to design interaction that cause self-assembly into lattices with chiral structure.

A chiral structure is one which exists in two forms which are related through a mirror symmetry, but not through rotations. Our hands are one example, and has given the term its name (from $\chi\varepsilon\iota\rho$, Greek for hand).

The self-assembly of chiral lattices is interesting from a physics perspective as an example of spontaneous symmetry breaking of mirror symmetry. However,

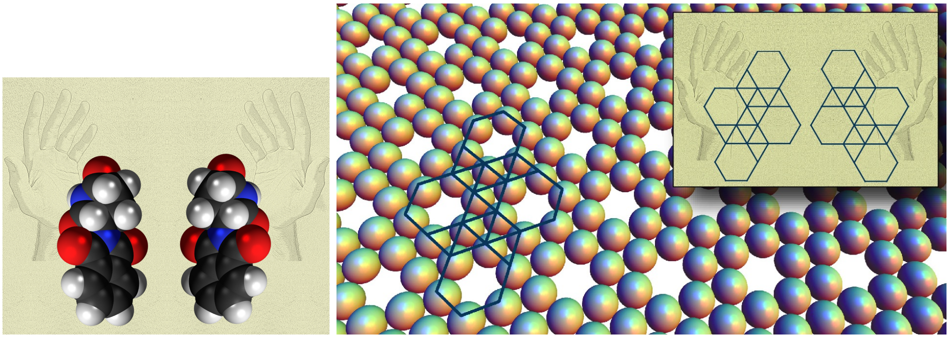

Figure 5.1: *(Left) Thalidomide, a chiral molecule. (Right) The snub hexagonal lattice, a two-dimensional lattice with chiral super-cells.*

the practical importance is found within chemistry, where methods for producing a specific enantiomer (specific "handedness") is vital for some pharmaceutical compounds. The most well-known example is thalidomide (see Figure 5.1) which has one enantiomer which is a sedative and one that causes birth defects if taken by the mother during pregnancy.

In Paper IV we design two chiral lattices, one simple composed of scalene triangles and one more complex (Figure 5.1), where the chirality comes from the mismatch between the symmetry of basis and lattice. We also study the phase diagram for the latter, varying the temperature and the strength of a perturbation parameter used in the construction. It turns out to be fairly complex with several different phases. We hypothesize that this is a general feature of systems designed with our method, as it selects for a set of features with different strengths, and varying parameters will change which of those features are displayed.

Paper V: Using the Uncertainty Principle to Design Simple Interactions for Targeted Self-Assembly

The results of Papers III and IV are interesting from a theoretical point of view, showing the existence of isotropic potentials that cause self-assembly into specific lattices. Some of these models are also worthy of study in their own right. For example, the assembly process of the second chiral lattice of Paper IV could give insight into the formation of chiral domains, a practical problem for any fabrication of chiral crystal structures.

However, the potentials constructed in Papers III and IV are too complicated to implement in experiment. In Paper V we therefore turn our attention to the problem of simplifying designed potentials.

To be completely certain what is complicated to implement in a given experimental setting requires full knowledge of said setting. However, some general features of potentials can be expected to be problematic in generic cases, such as potentials with many extremal points and long ranges. We thus focus on screening and smoothing as methods to simplify the designed potentials.

The main observation of Paper V is that the uncertainty relation places bounds on the simultaneous localization of the potential in real and reciprocal space. If we want to shorten the interaction range in real space, we cannot avoid a smoothing in reciprocal space. As our design method selects features in reciprocal space, we want to minimize this effect. We show that this prescribes the usage of Gaussian screening and smoothing.

We apply these results to the Kagome lattice studied in Paper III. There we did some unsystematic simplifications and managed to reduce the complexity from the naive implementation of the design method. However, we show that with the Gaussian screening and smoothing we can simplify the potential much further without losing its ability to self-assembly the target lattice.We design simple potentials that cause self-assembly into diamond lattices, interesting for its optical band-gap when implemented in colloids, and a disordered structure with suppressed diffraction in a given frequency range.

Paper VI: Predicting Self-Assembled Patterns on Spheres with Multicomponent Coatings

In this paper we generalize the approach of Papers I and II to spherical geometries[1] and, more importantly, multiple particle types. We also take a step closer towards experiment, working with a potential that is modeled on the alkanethiol-on-gold system.

Until now we have only considered systems with a single particle type (in the case of the striped patterns in binary space filling systems we can regard one of the types to be vacuum without loss of generality). Thus we have only handled a single potential, but here we need to be able to do more as with n particle types, there are $\binom{n}{2}$ unique potentials.

---

[1] Note and take care to avoid the possibility of confusing the word spherical when applied to the geometry and in the name of the spherical spin model.

It is still possible to diagonalize the Hamiltonian for the spherical spin model with multiple particle types. The diagonalization is mathematically the same as for the solution for the sound waves in atomic chains with multiple atoms in the basis. In both calculations a number of branches in the spectrum (phonon spectrum for the sound waves) appear as a result of the differences between the particle types breaking a degeneracy.

The structure of the diagonalization is the same regardless of geometry; as in Paper II, the spherical geometry enters through the eigenfunction, here as spherical harmonics. The ground state of the spherical spin model can thus be determined as being linear combinations of spherical harmonics with orbital number l determined by the minimum of the energy spectrum.

While the diagonalization of the spherical model is a fairly simple extension of the previous work, the mapping to discrete particle systems is a bit more involved. It turns out that each of the branches in the spectrum describes a partition of the particle types into two groups and how these groups vary relative to each other. For example, the spectrum corresponding to the system shown in Figure 3.6 has one branch whose minimum describes the phase separation of the blue particles and one that describes the stripe pattern between the red and yellow particles.

To test the predictions of our method, we formulated a simplified model of the alkanethiol-on-gold systems studied by Stellaci's and Glotzer's groups, among others. These consist of spherical gold nanoparticles covered with several types of hydrocarbon chains anchored with sulfur heads. The interactions of the chains cause them to form patterns in the shape of various combinations of stripes and spots. Our model captures most of this behavior (as shown in the supplementary material of Paper VI), and we thus use it as a test case for our method. The method turns out to describe both the morphologies and transitions between them in Monte Carlo simulations of this alkanethiol model.

Paper VII: A Design Path for Hierarchical Self-Assembly of Patchy Colloids

In the last paper we tie together the prediction methods with the idea of designing self-assembling systems. This is a direct continuation of Paper VI in that we use its form of the prediction method and our model system builds on the alkanethiol-on-gold model formulated there. In this paper we show that if we have a model for the interactions between the constituents of a coating on a sphere and we can make a given type attractive, we can hierarchically self-assemble a range of structures.

The design proceeds as follows. First a structure is chosen; in the paper we do diamond and cubic lattices, spherical vesicles of different sizes, strings and sheets, as well as isolated cubes. Then a patch arrangement that should cause assembly into these is chosen by positing attractive patches in positions corresponding to the neighbors in the finished structure. Next the arrangement has to be achieved in spherical harmonics; this can be guided by expanding the desired pattern with the spherical harmonics as the basis. Lastly, parameters for the alkanethiol-on-gold model that give a spectrum with minima at the correct harmonics have to be found and a stoichiometry chosen. This last step involves a bit of hand tuning, which is helped by a fast implementation of the prediction method as the spectrum can be recalculated almost in real time. The design is then tested by doing Monte Carlo simulations, first of the alkanethiol-on-gold method to check that the correct patch arrangement appears and then of a collection of colloids with those patches to determine the final structure.

Compared to the direct design method of Papers III-V, the setting proposed in Paper VII is much closer to experimental realizability. While Paper V goes a long way towards simplifying the designed potentials, there is still missing a model of a specific experimental system and the work of matching parameters of the latter with the design. The alkanethiol-on-gold system might not be what is eventually used, but the design path presented here would work with any system with a reasonable range of patterns, as the stripes and spots are universal in the sense of Paper I.

Chapter 6

Summary and outlook

In this thesis we have seen how exactly solvable models can be applied to various problems in self-assembly. By generalizing the exactly solvable spherical spin model to arbitrary isotropic interactions we have been able to predict morphologies appearing in a range of systems. By designing energy spectra and inverting them we have been able to construct interactions whose ground states are exactly some desired lattice.

A lesson from the use of exact solutions in statistical mechanics is that "quantities and functions, introduced primarily as mathematical devices, almost invariably acquire a fundamental physical meaning." [68] For example, the Lagrange multiplier β that occurs in thermodynamical calculations turns out to be nothing less than the (inverse) temperature, $\beta = 1/k_B T$. In this thesis we have seen this in the solution to the spherical model in cylindrical geometries, where we saw no way to lift the ω-degeneracy and this turned out to persist in the particle simulations, showing that the degeneracy is indeed fundamental. Relatedly, in applying the solution to the multiple particle types generalization of the spherical spin model we saw that the branch structure reflected the morphology of the particle systems, and that when small quantitative changes caused qualitative changes of this structure that corresponded to phase transitions in the particle system.

Fruitful studies always point out new directions for research, and to whatever extent the reverse holds this thesis would seem a success. Some possibilities that immediately suggest themselves follow.

The solutions to the various generalizations of the spherical spin model presented here all concern ground states only. For the original, nearest neighbor version of the model an exact solution for the complete partition function is known.

Such a solution might well be obtainable for some or all of our generalizations as well, possibly with restrictions on the interactions.

As alluded to in Chapter 4, while the design method of Papers III-V have already given results of theoretical interest, there is still the question of whether it is possible to use them to design experimental systems directly. Here identification of a system for which effective interactions is known and sufficiently rich is needed, or construction of such an effective description. Further, exploration of the idea of using optimization in reciprocal space to restrict the method to a given collection of potentials is needed from the theory side.

The prediction methods of Papers I, II, and VI tell us about what morphologies typically appear in different types of systems with isotropic interactions, while Papers III-V show one way to design specific interactions to achieve desired morphologies. This suggests a third line of inquiry, regarding the range of possible structures. Are there morphologies such that no isotropic potentials can be designed to have them as (unique or degenerate) ground states? This question is intimately connected to the question of degeneracy of the pair correlation function; if two structures has the same pair correlation function $g(r)$, then they will have the same energy for all isotropic potentials and one would need anisotropies (or possibly entropic factors at non-zero temperature) to select for either of them. The unstated assumption in large parts of the literature seems to be that these degeneracies are rare and special cases, but this is far from proven and an important open question. Classes of degenerate processes [69] and structures [70, 71] are known but whether they present important limitations is less clear; practical results of image reconstruction methods suggests most structures are fully specified at least by the correlation function $g(\mathbf{r})$ (up to rotations and reflections), see for example [72]. Ideally we would want to have probabilistic statements of the type possible for reconstruction from sparse Fourier signals [52].

The final question, the challenge this thesis presents to the self-assembly community, is whether other exact methods and models can be adapted to our problems. Perhaps results from spin glass theory have bearing on questions regarding formation and solution of defects in systems undergoing self-assembly. Can methods from computer science be used to prove possibilities and limitations? The nascent field of algorithmic self-assembly suggests so [73].

All in all, theoretical understanding of self-assembly is in its infancy and much remains to be done.

Bibliography

1. Baxter, R. J. *Exactly Solved Models in Statistical Mechanics* (Academic Press, London, UK, 1982).
2. Axilrod, B. M. & Teller, E. Interaction of the van der Waals Type Between Three Atoms. *Journal of Chemical Physics* **11,** 299 (1943).
3. Likos, C. N. Effective interactions in soft condensed matter physics. *Physics Reports* **348,** 267–439 (2001).
4. Sullivan, L. H. The tall office building artistically considered. *Lippincott's Monthly Magazine* **57,** 403–09 (1896).
5. Ball, P. *The Self-Made Tapestry: Pattern Formation in Nature* (Oxford University Press, 2001).
6. Thompson, D. W. *On Growth and Form* (Cambridge University Press, 1917).
7. Cross, M. C. & Hohenberg, P. C. Pattern formation outside of equilibrium. *Reviews of Modern Physics* **65,** 851–1112 (1993).
8. Vanderbauwhede, A. *Local Bifurcation and Symmetry* Habilitation Thesis (Rijkuniversiteit Gent, 1980).
9. Cicogna, G. Symmetry breakdown from bifurcation. *Lettere Al Nuovo Cimento* **31,** 600–602 (1981).
10. Hoyle, R. *Pattern Formation: An Introduction to Methods* (Cambridge University Press, 2006).
11. Keller, S. L. & McConnell, H. M. Stripe Phases in Lipid Monolayers near a Miscibility Critical Point. *Physical Review Letters* **82,** 1602–1605 (1999).

12. Kern, K., Niehus, H., Schatz, A., Zeppenfeld, P., Goerge, J. & Comsa, G. Long-Range Spatial Self-Organization in the Adsorbate-Induced Restructuring of Surfaces: Cu100-(2×1)O. *Physical Review Letters* **67,** 855–858 (1991).

13. Rosensweig, R. E., Zahn, M. & Shumovich, R. Labyrinthine instability in magnetic and dielectric fluids. *Journal of Magnetism and Magnetic Materials* **39,** 127–132 (1983).

14. Seul, M. & Wolfe, R. Evolution of disorder in magnetic stripe domains. I. Transverse instabilities and disclination unbinding in lamellar patterns. *Physical Review A* **46,** 7519–7533 (1992).

15. Tranquada, J. M., Sternlieb, B. J., Axe, J. D., Nakamura, Y. & Uchida, S. Evidence for stripe correlations of spins and holes in copper oxide superconductors. *Nature* **375,** 561–563 (1995).

16. Orenstein, J. & Millis, A. Advances in the Physics of High-Temperature Superconductivity. *Science* **288,** 468–474 (2000).

17. Garel, T. & Doniach, S. Phase transitions with spontaneous modulation – the dipolar Ising ferromagnet. *Physical Review B* **26,** 325–329 (1982).

18. Löw, U., Emery, V. J., Fabricius, K. & Kivelson, S. A. Study of an Ising Model with Competing Long- and Short-Range Interactions. *Physical Review Letters* **72,** 1918–1921 (1994).

19. Grousson, M., Tarjus, G. & Viot, P. Phase diagram of an Ising model with long-range frustrating interactions: A theoretical analysis. *Physical Review E* **62,** 7781–7792 (2000).

20. Stoycheva, A. D. & Singer, S. J. Stripe Melting in a Two-Dimensional System with Competing Interactions. *Physical Review Letters* **84,** 4657–4660 (2000).

21. Giuliani, A., Lebowitz, J. L. & Lieb, E. H. Ising models with long-range antiferromagnetic and short-range ferromagnetic interactions. *Physical Review B* **74,** 064420 (2006).

22. Zaanen, J. Current ideas on the origin of stripes. *Journal of Physics and Chemistry of Solids* **59,** 1769–1773 (1998).

23. Malescio, G. & Pellicane, G. Stripe phases from isotropic repulsive interactions. *Nature Materials* **2,** 97–100 (2003).

24. Grantcharova, V., Alm, E. J., Baker, D. & Horwich, A. L. Mechanisms of protein folding. *Current Opinion in Structural Biology* **11,** 70–82 (2001).

25. Albert, B., Johnson, A., Lewis, J., Raff, M., Roberts, K. & Watson, J. D. *Molecular Biology of the Cell* 4th ed. (Garland, New York, 2002).

26. Schwiebert, K. E., Chin, D. N., MacDonald, J. C. & Whitesides, G. M. Engineering the Solid State with 2-Benzimidazolones. *Journal of the American Chemical Society* **118,** 4018–4029 (1996).

27. Desiraju, G. *Crystal Engineering: the Design of Organic Solids* (Elsevier, New York, 1989).

28. Schmidt-Mende, L., Fechtenkötter, A., Müllen, K., Moons, E., Friend, R. H. & MacKenzie, J. D. Self-Organized Discotic Liquid Crystals for High-Efficiency Organic Photovoltaics. *Science* **293,** 1119 –1122 (2001).

29. De Rosa, C., Park, C., Thomas, E. L. & Lotz, B. Microdomain patterns from directional eutectic solidification and epitaxy. *Nature* **405,** 433–437 (2000).

30. Maldovan, M. & Thomas, E. L. Diamond-structured photonic crystals. *Nature Materials* **3,** 593–600 (2004).

31. Gross, R. & Dorigo, M. Self-Assembly at the Macroscopic Scale. *Proceedings of the IEEE* **96,** 1490–1508 (2008).

32. Gross, R., Bonani, M., Mondada, F. & Dorigo, M. Autonomous Self-Assembly in Swarm-Bots. *IEEE Transactions on Robotics* **22,** 1115–1130 (2006).

33. Thomas, E. L. The ABCs of Self-Assembly. *Science* **286,** 1307 (1999).

34. Gracias, D. H., Tien, J., Breen, T. L., Hsu, C. & Whitesides, G. M. Forming Electrical Networks in Three Dimensions by Self-Assembly. *Science* **289,** 1170 –1172 (2000).

35. Boncheva, M., Gracias, D. H., Jacobs, H. O. & Whitesides, G. M. Biomimetic self-assembly of a functional asymmetrical electronic device. *Proceedings of the National Academy of Sciences* **99,** 4937 –4940 (2002).

36. Seeman, N. C. DNA nanotechnology: novel DNA constructions. *Annual Review of Biophysics and Biomolecular Structure* **27.** PMID: 9646868, 225–248 (1998).

37. Seeman, N. C. At the Crossroads of Chemistry, Biology, and Materials: Structural DNA Nanotechnology. *Chemistry & Biology* **10,** 1151–1159 (2003).

38. Fleishman, S. J., Whitehead, T. A., Ekiert, D. C., Dreyfus, C., Corn, J. E., Strauch, E., Wilson, I. A. & Baker, D. Computational Design of Proteins Targeting the Conserved Stem Region of Influenza Hemagglutinin. *Science* **332,** 816 –821 (2011).

39. Lehn, J. Toward Self-Organization and Complex Matter. *Science* **295,** 2400 –2403 (2002).
40. Whitesides, G. M. & Boncheva, M. Beyond molecules: Self-assembly of mesoscopic and macroscopic components. *Proceedings of the National Academy of Sciences* **99,** 4769 –4774 (2002).
41. Whitesides, G. M. & Grzybowski, B. Self-Assembly at All Scales. *Science* **295,** 2418 –2421 (2002).
42. Zhang, S. Fabrication of novel biomaterials through molecular self-assembly. *Nature Biotechnology* **21,** 1171–1178 (2003).
43. Zhang & Glotzer, S. C. Self-Assembly of Patchy Particles. *Nano Letters* **4,** 1407–1413 (2004).
44. Barth, J. V., Costantini, G. & Kern, K. Engineering atomic and molecular nanostructures at surfaces. *Nature* **437,** 671–679 (2005).
45. Ising, E. Beitrag zur Theorie des Ferromagnetismus. *Zeitschrift für Physik* **31,** 253–258 (1925).
46. Stanley, H. E. Spherical Model as the Limit of Infinite Spin Dimensionality. *Physical Review* **176,** 718 (1968).
47. Brush, S. G. History of the Lenz-Ising Model. *Reviews of Modern Physics* **39,** 883 (1967).
48. Berlin, T. H. & Kac, M. The Spherical Model of a Ferromagnet. *Physical Review* **86,** 821 (1952).
49. Andreasson, N., Evgrafov, A. & Patriksson, M. *An Introduction to Continuous Optimization* (Studentlitteratur, Lund, Sweden, 2005).
50. Santosa, F. & Symes, W. W. Linear Inversion of Band-Limited Reflection Seismograms. *SIAM Journal on Scientific and Statistical Computing* **7,** 1307 (1986).
51. Chen, S. S., Donoho, D. L. & Saunders, M. A. Atomic Decomposition by Basis Pursuit. *SIAM Review* **43,** 129 (2001).
52. Candes, E., Romberg, J. & Tao, T. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory* **52,** 489–509 (2006).
53. Barahona, F. On the computational complexity of Ising spin glass models. *Journal of Physics A: Mathematical and General* **15,** 3241–3253 (1982).

54. Aspvall, B. & Gilbert, J. Graph coloring using eigenvalue decomposition. *SIAM Journal on Algebraic and Discrete Methods* **5,** 526–538 (1984).
55. Fiedler, M. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal* **23,** 298–305 (1973).
56. Newman, M. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences* **103,** 8577–8582 (2006).
57. Anderson, P. W. More Is Different. *Science* **177,** 393–396 (1972).
58. Chen, Q., Bae, S. C. & Granick, S. Directed self-assembly of a colloidal kagome lattice. *Nature* **469,** 381–384 (2011).
59. Torquato, S. Inverse optimization techniques for targeted self-assembly. *Soft Matter* **5,** 1157 (2009).
60. Cohn, H. & Kumar, A. Algorithmic design of self-assembling structures. *Proceedings of the National Academy of Sciences* **106,** 9570 –9575 (2009).
61. Rechtsman, M., Stillinger, F. H. & Torquato, S. Optimized Interactions for Targeted Self-Assembly: Application to a Honeycomb Lattice. *Physical Review Letters* **95,** 1–4 (2005).
62. Rechtsman, M. C., Stillinger, F. H. & Torquato, S. Self-assembly of the simple cubic lattice with an isotropic potential. *Physical Review E* **74,** 021404 (2006).
63. Rechtsman, M., Stillinger, F. & Torquato, S. Synthetic diamond and wurtzite structures self-assemble with isotropic pair interactions. *Physical Review E* (2007).
64. Jain, A., Errington, J. R. & Truskett, T. M. Inverse design of simple pairwise interactions with low-coordinated 3D lattice ground states. en. *Soft Matter* **9,** 3866–3870 (2013).
65. Ashcroft, N. W. & Mermin, N. D. *Solid State Physics* (Brooks Cole, Belmont, USA, 1976).
66. Schachinger, W., Addis, B., Bomze, I. M. & Schoen, F. New results for molecular formation under pairwise potential minimization. *Computational Optimization and Applications* **38,** 329–349 (2007).
67. Davis, P. J. *Circulant Matrices* (Chelsea Pub Co, Providence, RI., 1994).
68. Schrödinger, E. *Statistical Thermodynamics* English (Cambridge University Press, 1944).

69. Baddeley, A. J. & Silverman, B. W. A Cautionary Example on the Use of Second-Order Methods for Analyzing Point Patterns. *Biometrics* **40,** 1089–1093 (1984).
70. Jiao, Y, Stillinger, F. H. & Torquato, S. Geometrical ambiguity of pair statistics: Point configurations. *Physical Review E* **81,** 011105 (2010).
71. Jiao, Y., Stillinger, F. H. & Torquato, S. Geometrical ambiguity of pair statistics. II. Heterogeneous media. *Physical Review E* **82,** 011106 (2010).
72. Rozman, M. G. & Utz, M. Uniqueness of Reconstruction of Multiphase Morphologies from Two-Point Correlation Functions. *Physical Review Letters* **89,** 135501 (2002).
73. Doty, D. Theory of algorithmic self-assembly. *Communications of the ACM* **55,** 78 (2012).