Methods, rules and limits of successful self-assembly

A J Williamson
Wadham College
University of Oxford

Supervisor: Dr J P K Doye

Hilary Term 2011

Abstract

Methods, rules and limits of successful self-assembly
A. J. Williamson, Wadham College
Submitted for the degree of Doctor of Philosophy in Physical and Theoretical Chemistry,
Hilary Term 2011.

The self-assembly of structured particles into monodisperse clusters is a challenge on the nano-, micro- and even macro-scale. While biological systems are able to self-assemble with comparative ease, many aspects of this self-assembly are not fully understood. In this thesis, we look at the strategies and rules that can be applied to encourage the formation of monodisperse clusters. Though much of the inspiration is biological in nature, the simulations use a simple minimal patchy particle model and are thus applicable to a wide range of systems. The topics that this thesis addresses include:

Encapsulation: We show how clusters can be used to encapsulate objects and demonstrate that such 'templates' can be used to control the assembly mechanisms and enhance the formation of more complex objects.

Hierarchical self-assembly: We investigate the use of hierarchical mechanisms in enhancing the formation of clusters. We find that, while we are able to extend the ranges where we see successful assembly by using a hierarchical assembly pathway, it does not straightforwardly provide a route to enhance the complexity of structures that can be formed.

Pore formation: We use our simple model to investigate a particular biological example, namely the self-assembly and formation of heptameric α-haemolysin pores, and show that pore insertion is key to rationalising experimental results on this system.

Phase re-entrance: We look at the computation of equilibrium phase diagrams for self-assembling systems, particularly focusing on the possible presence of an unusual liquid-vapour phase re-entrance that has been suggested by dynamical simulations, using a variety of techniques.

Declaration

This thesis is submitted for the degree of Doctor of Philosophy in Physical and Theoretical Chemistry at the University of Oxford. No part of this thesis has been accepted or is currently being submitted for any degree, diploma, certificate or other qualification in this University or elsewhere. This thesis is wholly my own work, except where indicated.

Acknowledgements

I would like to thank the following people and organisations for their support during the writing of this thesis:

- Jessica Andréani Feuillet and Alex Wilber for their scientific discussions and thesis-writing knowledge.
- André for preventing me from throwing the computer out of the window on days when I just didn't get on with LaTeX, and Aleks Reinhardt for proof-reading this mammoth tome.
- Lucy for making sure I took sufficient breaks from working and for showing me that the world continues to spin even when a thesis needs writing.
- Dr Handel, for helping with the proof reading, and Médico for always being ready with a cup of tea and for putting up with me living with him.
- Dr Priest for convincing me to keep going, even when the going got tough.
- Ländy for making sure I understood the true magic of Nandos, and to Loren for helping with proof-reading.
- My brothers, Pip and Jonty, because they've always wanted to be in an acknowledgements section.
- My supervisor, Jon, for always being ready to help.
- The EPSRC for funding.
- My parents for encouraging me to follow my passions and giving me both financial and moral support.

Publications

During my research I have published / been involved with the publishing of several papers [1–7], for which I thank my collaborators and colleagues for their input and time.

To my father, whose courage, strength and honour showed me how to live my life.
May he be at peace with God.

Contents

Appendices

Abbreviations and acronyms

clus	Cluster
cO	Cuboctahedron
CPU	Central Processing Unit
CUDA	Compute Unified Device Architecture
Cuboct	Cuboctahedron
Dd	Dodecahedron
FORTRAN	FORmula TRANslation
GCWL	Grand Canonical Wang-Landau
GE	Gibbs Ensemble
GPGPU	General-Purpose computing on Graphics Processing Units
GPU	Graphics Processing Unit
Hept	Heptamer / Heptamerisation
HS	Hard Sphere
Ih	Icosahedron
MC	(Metropolis) Monte Carlo
MD	Molecular Dynamics
rC	Rhombicuboctahedron
RCP	Random Close-Packed
VMMC	Virtual Move Monte Carlo
WL	Wang-Landau

Glossary

Bacteriocidal	Causes the death of bacteria
Bacteriostatic	Prevents the growth of bacteria
Ergodic	The property that all regions in phase space will be visited given enough time
Gibbs state	A state whose probability distribution is invariant in time
Heteromeric	Made of differing components
Homomeric	Made of two or more of the same component
Neutrophil	A type of white blood cell responsible for defending the body against infection
Number density	The number of particles per unit volume: $\frac{N}{V}$
Monodisperse	Being of the same size
Monte Carlo cycles	Number of Monte Carlo steps divided by the number of particles
Torsional	Relating to a twisting motion along an axis

Symbols

A	Helmholtz free energy
b	Monomer volume exclusion coefficient
B	Cluster volume exclusions coefficient
C_V	The heat capacity at a constant volume
g	Density of a state
G	Gibbs free energy
h	Planck constant
H	Enthalpy
I	Moment of inertia
k	Boltzmann constant
m	Mass of a monomer
N	Total number of patchy particles
N_{A}	Avogadro constant
N_{clus}	Number of clusters
N_{mon}	Number of particles in a monomeric state
p	Pressure
p_{ij}	Probability of forming a link between particle i and j
P	Probability
q	Molecular partition function
q_{ij}	Probability of not forming a link between particle i and j
Q	Ensemble partition function
$\boldsymbol{r}$	Interparticle vector
$\mathcal{R}$	A random number in the interval $[0, 1]$
T	Temperature
T_{agg}	Kinetic temperature at which glassy aggregates form
T_{clus}	Temperature of transition between cluster and monomer gases
T_{Ads}	Temperature of adsorption to a surface
T_{Hept}	Temperature at which a heptamer and a box of seven monomers have the same free energy
T_{Trans}	Temperature at which an inserted heptameric pore and a box of seven monomers have the same free energy
U	Internal energy
U_{ij}	Interaction between particles i and j
v	Volume per particle $\dfrac{V}{N}$
V	Volume
V_{free}	Free volume
W	Probability of selecting a seed particle

α	Probability of attempting a move
β	Inverse temperature $1/kT$
γ	Probability of accepting a move
Δ	A change in
ε	Well depth of a patch-patch interaction
ζ	Patch noise
$\mathit{\Lambda}$	de Broglie wavelength
μ	Chemical potential
ν	Boltzmann probability of a given state
π	Probability of a move being attempted and accepted
ρ	Total number density $\frac{N}{V}$
ρ^*	Reduced number density $\frac{N}{V} \times \sigma^3$
ρ_{mon}	Monomer number density $\frac{N_{\mathrm{mon}}}{V}$
ρ_{clus}	Cluster number density $\frac{N_{\mathrm{clus}}}{V}$
ϕ	Packing fraction
ϕ_{RCP}	Packing fraction of a randomly close-packed system
σ^{PW}	Patch width
σ^{LJ}	Particle length scale (approximate diameter)
σ^{TOR}	Width of the torsional constraint Gaussian (measured in radians)

Chapter 1

Introduction

1.1 Overview

The general aim of this thesis is to explore methods of particle design for self-assembling systems that encourage high yields of correctly formed monodisperse clusters. The bulk of this thesis uses computer simulations to probe both the kinetics and thermodynamics of self-assembly. This chapter introduces the background of novel micro- and nano-technological materials, current research and our interests. As the main inspiration for much of the research we present here has been biological in nature, we also include a simple introduction to the biology of proteins and the self-assembled structures they can form. The next chapter deals with our model system and explains some of the methods used in this thesis in detail.

1.1.1 Aims of the thesis

We use a simple minimal model for particles, which enables us to find general rules of assembly rather than ones specific to a single system. This simple model is explained in section 2.1, while some of the possible applications of such a model are looked at in the following sections of this chapter. The core aim throughout this thesis is to find generally applicable rules for self-assembling systems. The sub-aims are summarised below:

- To explore simple models of biologically interesting systems.
- To find whether templating is an effective strategy to aid the self-assembly of monodisperse shells.
- To examine the effect of hierarchical assembly pathways, particularly whether they lead to an enhancement of self-assembly yields.
- To investigate the phase behaviour of self-assembling systems, particularly in regions of parameter space where self-assembly competes with phase separation.

1.1.2 Structure of the thesis

The thesis has 7 chapters, each of which deals with a different aspect of self-assembling systems. After providing an overview of the background (this chapter), we introduce the model and methods we have used (chapter 2). Then, we look at the use of templating on both the kinetics and the thermodynamics of our simple model (chapter 3). This approach is inspired by the self-assembly of a capsid around a virus's genetic material, and the background to this work is in the current chapter. The next chapter (chapter 4) deals with the idea of hierarchical self-assembly, inspired by the hierarchical nature of assembly for many protein complexes. We then look at a complex biological system more directly: the self-assembly of pores in a membrane (chapter 5) . The model provides a comparison to the experimental data obtained by the group of Dr Mark Wallace [8]. We then study the phase behaviour associated with self-assembly, especially phase re-entrance, using both simulation and theory (chapter 6). The final chapter (chapter 7) summarises the conclusions from our research.

The next section looks at patchy particles and their synthesis. For any reader without a biological background, we then give a short introduction to the biology of proteins, and viruses (section 1.3). Finally, we detail some of the vast body of experimental and simulation work on self-assembling systems that has been carried out by a number of groups worldwide.

1.2 Patchy particles

A prime motivation for the current work is the concerted efforts being made by many groups to develop synthetic methods which produce nano- and micro-particles of colloids with anisotropic interactions [9]. These have the potential for use in a number of areas, including the creation of photonic electronic devices [10]. This potential use would require monodisperse self-assembly of specific structures, which is one of the aims of this thesis. Much of this work is designed to give a set of design rules and thus a framework for experimentalists looking to design cluster-forming patchy colloids.

We first look at the current experimental techniques used in synthesis for background, and to help explain our simulation decisions later in thesis.

1.2.1 Experimental synthesis of "patchy particles"

Synthesising colloids with isotropic interactions occurs daily throughout the world. Milk, for example, is a colloidal suspension of fat droplets. A colloid is defined as a system made up of particles or objects that have at least one dimension in the 1 nm to 1 μm range [11].

The development of synthetic methods to produce colloids that are reliably monodisperse in size has been particularly important [12], and has allowed a wide range of potential areas to be explored, including the study of fundamental physics within the 'colloids as atoms' paradigm.

In order to go beyond the behaviour seen for monodisperse isotropic colloids and nanoparticles, there has been much recent interest in developing methods to produce particles with anisotropic interactions. When these anisotropic interactions involve selective attractions between different parts of the particle surface, the particles are commonly referred to as *patchy particles.* The exploration of the potential behaviour of these particles is the focus of this thesis. Different ways of creating anisotropic colloids have been recently reviewed by Glotzer and Solomon [13]. We will cover and explain the basic ideas behind their *dimensions of anisotropy* before looking at specific experimental examples.

When looking at producing anisotropic particles, Glotzer and Solomon [13] list eight possible variables. We extend this to twelve and then divide these into three categories. It is possible to change the shape, the surface patterns and the clustering of particles formed from smaller component particles [13]. These are pictorially shown in figure 1.1, and explained below:

- **Surface patterns:** This includes the coverage of the patches, the position and shape of patches (called pattern quantisation by Glotzer and Solomon [13]), and the number of patches.

- **Shape:** This covers changing the aspect ratios of the otherwise spherical particles, any branching that may be present on the particles, the chirality of particles, and the number of faces on any shaped particle.

- **Clustering:** When the final particle is made of several bonded components, it is possible for the ordering, relative shapes, positions and number of components to be varied.

Recent developments have allowed a number of these variables to be used when producing anisotropic particles. This has lead to a number of different shapes being produced, including Janus[1] [14] and other hemispherically differentiated particles [15]; branched structures made from both organic molecules and inorganic nano-crystals [16–20]; particles in the shape of platonic solids [21–23]; and many particles of other shapes [24–26]. Often, batch production methods are limited in the production of anisotropic particles, and more recent methods using micro-fluidics have shown promise [26].

---

[1] Janus particles are named after the eponymous Roman god: they are "two-faced" particles with different functionalisation on each face leading to a different type of interaction.

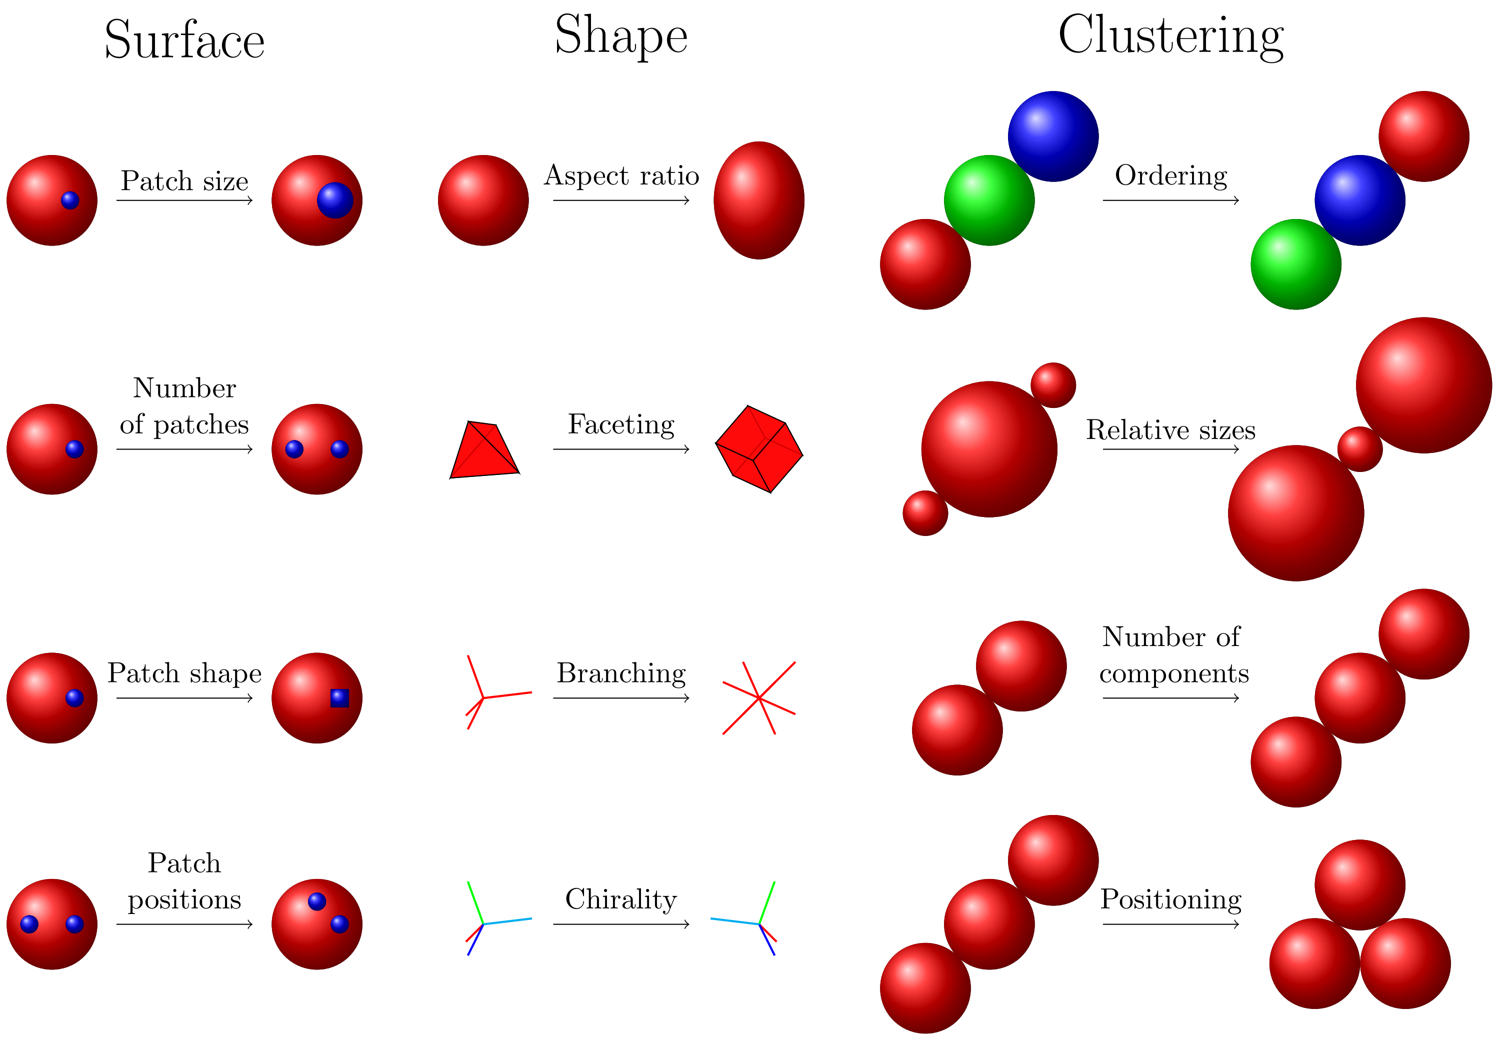

Figure 1.1: Different dimensions of anisotropy that can be varied in a patchy particle.

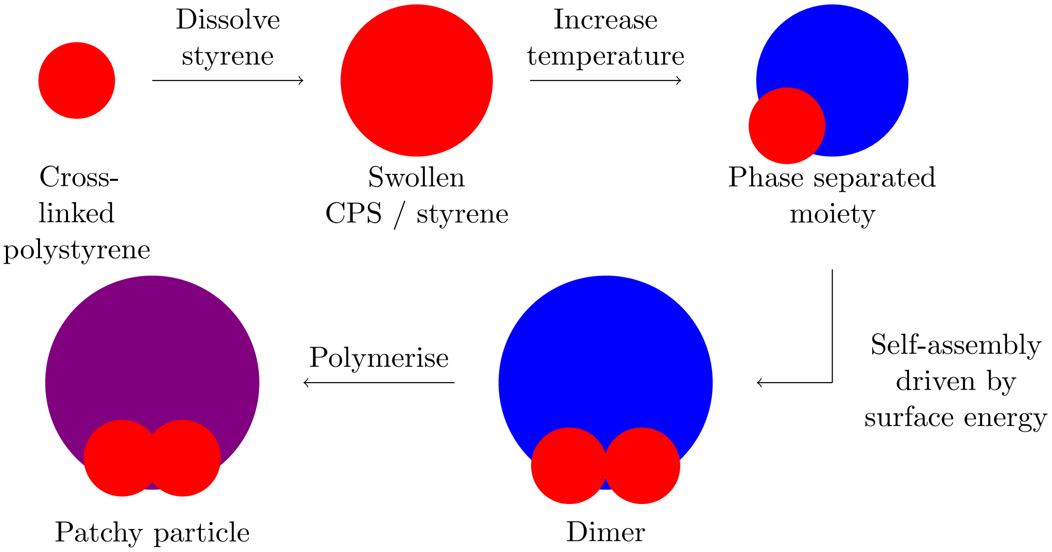

Figure 1.2: A schematic of Kraft's synthetic technique (adapted from [27]).

The reasons that certain inorganic nano-crystals can be made to form a given shape is not fully understood [18], although some techniques have been shown to be very effective in the production of nano-crystals, including limiting the size and shape with templates and using capping agents [22]. Our interest is the self-assembly of nano-crystals/colloids into a given structure, so we do not concentrate much on this area of science here.

Patchy colloids have been experimentally produced with excellent control of the angles between bonds by Kraft *et al* [27]. These are produced using some clever science. Divinylbenzene cross-linked polystyrene spheres are first made, then these are made to swell by dissolving extra styrene monomer inside them. The temperature is then raised, which causes the monomer and cross-linked polymer phases to separate. This phase separation minimises the free energy, and therefore depends on the wetting angle of the cross-linked polystyrene spheres and the amount of monomer dissolved. It causes a protrusion of liquid monomer outside of the seed particle (figure 1.2). If these seed-protrusion moieties are allowed to come together, they again act to minimise free energy, causing contact between the cross-linked spheres (figure 1.2): the angle between these contacts is dependent on the wetting angle and the amount of monomer present. If these are polymerised, then a patchy particle is formed. As the number of seed-protrusion moieties in a particle may vary (though the predominant sizes formed may be controlled to some extent by changing the concentration and time allowed before attempting polymerisation), it is necessary to separate out the assemblies using centrifugation or similar.

The group of David Pine has demonstrated the utility of methods involving oil-in-water emulsions [28–30] in creating patchy colloids. They place colloidal microspheres in toluene so that they swell, and then emulsify the microspheres in water. On removing the toluene they are able to create clusters of solid particles which are stabilised by the surface charge [30]. More recently, they have also demonstrated an approach to the creation of programmable colloids based on the *lock and key* principle of enzyme kinetics [31]. They use a colloidal particle with a spherical cavity (the "lock") and a colloidal sphere able to fit into that cavity (the "key"). The "key" can insert only into a correctly-sized "lock" and this interaction is reversible, and the position of equilibrium may be changed by varying the temperature (which may changes the size of both the "lock" and the "key") [31].

Cross-linking polymer spheres by sulphonation [32] has also been shown to be a valid route to producing patchy colloids of the types we are aiming to simulate. However, all of the approaches mentioned are in the initial phases of development, and an industrial process has still not been found.

1.2.2 Applications of patchy particle clusters

As mentioned previously, colloidal spheres are an incredibly common material. New applications are being found for these substrates, especially in the areas of nanotechnology and optics [11]. One of the most commonly cited potential applications is that of a photonic band gap crystal; such three-dimensional photonic band gap materials were originally postulated in the late 1980s [33, 34].

A photonic band gap material is one which does not permit the transmission of certain wavelengths of light [34]. It is normally produced from materials whose dielectric constant varies throughout the material. They have many possible uses in photonic computer applications [10]. Theoretical calculations have shown that a diamond-like crystal made from colloidal spheres would create a band gap in the visible spectrum [35]. The creation of such a structure has been shown possible in simulations of patchy particles [36].

The idea of a photonic band gap seems strange at first, but is no more complicated than the idea of electronic band gap materials not permitting electrons of certain energies from existing (which is the underpinning of much of the semi-conductor technology applications we see today). It is, in fact, for this reason that there is such a great deal of interest in the production of photonic band gap materials. Photonic band gap materials have possible uses as logic gates for photons [37]. This would allow computers that use light in their circuits over electrons, and give a large potential increase in computing power (the speed of light is orders of magnitude faster than that of electrons in a circuit). In fact, as light is often used to transmit data between two computers (by the use of optical fibres), photonic technology in routers and gateways would prevent the need for slow optical to electronic conversion and the associated loss of data signalling speed [10].

Other uses

There are many other possible uses for colloidal crystals, including their use in catalysis. Nano-crystals have a surface on which many reactions can take place, and their small size means that they have a large surface area to volume ratio, making them cost-effective in chemical reactors. The strength of certain nano-crystals is that they possess the properties both of the functional solid state, as well as the advantage of being able to be placed in solution [38]. Solar energy conversion has been postulated as a possible use for crystals formed of copper(I) oxide colloidal clusters [39]. Even data storage is seen as a possible application [22]. In fact, the versatility and scale of colloidal crystals and patchy particles is incredibly wide, owing to the large range of sizes and materials that can be used [9, 38].

1.3 Biology: a quick primer

This section deals with the basic ideas behind protein and virus biology. The explanations are brief and are designed only to give the reader an understanding of why we made certain choices in our simulations and what inspired some of our research in this thesis.

1.3.1 An introduction to the biology of proteins

Proteins are ubiquitous in biology. They can act as catalysts, structural building blocks, sensors and transporters, to name but a few of their uses. Much of our research has been inspired by the behaviour of biological systems: it has been argued that a simple patchy particle model can in fact be used as a minimal model for protein assembly [40], and has in fact already been used for a variety of tasks, including explaining the evolution of protein complexes [41], the formation of fibrils [42] and even protein crystallisation [43].

Let us take a look at some of the properties of proteins in order to explain the background behind many of the inspirations for this thesis and the choices made in simulations.

Structure

Proteins are created from the association of one or more polypeptide chains, with the possible addition of other rings or groups. Polypeptides are polymers made from amino acids. These monomers contain both a carboxylic acid group and an amine group. The ordering of amino acids is controlled by the genetic code of an organism (generally DNA).

Biologists typically use four levels of structure when referring to proteins [44, 45]:

- **Primary:** The actual ordering of amino acids in a polypeptide chain

- **Secondary:** The initial folding of the polypeptide chains into α-helices and β-pleated sheets
- **Tertiary:** The folding of the polypeptide chains into more complex shapes
- **Quaternary:** The final structure composed of one or more polypeptide chains and any non-polypeptide parts (for example, a haem group in haemoglobin)

Our research is mainly inspired by the aggregation of proteins or polypeptides with each other to form a final quaternary structure. This can occur between identical polypeptides (forming a *homomeric* complex) or polypeptides with different primary structure (forming a *heteromeric* complex). The interface between polypeptides is often well-defined, and can be treated as a *patch* [46]. In our model, we have effectively assumed that monomeric proteins fold into their tertiary ground state on a faster timescale than that of self-assembly into a quaternary structure. While this is true for most protein complexes, we are ignoring those in which partially-folded proteins may form aggregates [47], and proteins with no stable monomeric state for which the first assembly step is the dimerisation of two unfolded polypeptide chains [47, 48].

Intra- and inter-protein bonding

Bonding within a protein can take many different forms. A basic explanation of some of the common intra- and inter-peptide bonds found in biology is given below [49, 50]:

- **Hydrogen bonding:** This occurs when there is a significant difference in electronegativity between hydrogen and another atom (normally nitrogen or oxygen). The uneven sharing of electrons leads to charge separation. This charge separation causes electrostatic attractions.
- **Charged interactions:** At physiological pH, certain side groups are charged. These charged groups can interact through space in a coulombic way. It is important to note that often, the permittivity relevant to the interactions is not that of the bulk solvent, as the interaction may occur through other parts of the protein chain, or in areas where the solvent is organised.
- **Covalent bonds:** As well as the amide bonds forming the backbone of the protein, side groups are able to bond to each other. The most common is that of two cysteine residues forming a bond between two sulphur atoms. This is known as a *disulphide bridge* and is a major part of keratin structure.

- **Hydrophobic interactions:** As some side groups are non-polar, they do not dissolve in water and thus often tend to come together. This is a complex effect with both entropic and enthalpic factors [51].

Internal degrees of freedom

The actual structure of a protein depends on a number of environmental factors as well as the identity of the amino acids (primary structure). A change in pH, for example, can cause the protonation or deprotonation of an amino acid's side chain, and this thus can change the charge density and the stability of different conformations [52]. Changes in connectivity may also occur when the protein binds to a specific substrate; for example, in the catalytic reaction of co-enzyme B_{12}, the bond between the enzyme and the cobalt of the colbalamin ring breaks homolytically, changing the conformation of the protein [53]. Changes are also common when proteins with a similar primary structure form a complex cluster together, especially in virus capsid coat assembly.

Co-operativity

If two or more polypeptides are in close proximity, then changes in internal degrees of freedom in one polypeptide can influence the energy landscape of another polypeptide's internal degrees of freedom. This phenomenon is known as *co-operativity.* A classical example of this is the bonding of oxygen to haemoglobin [44]. Haemoglobin is a tetramer of iron-bound proteins, and the binding of oxygen to one causes the structure of the final protein to shift in such a way as to permit an oxygen molecule binding to any of the remaining free haem-sites to be much easier than the initial binding [54]. We look into the idea of both internal degrees of freedom and co-operativity in chapter 5.

1.3.2 Introduction to viruses

Viruses are not technically "alive" in the standard biological sense of the word [55]. They are unable to reproduce without a host, do not respire and are unable to move save for Brownian and other motion of the solvent in which they are carried. Despite this, they are an excellent example of biological self-assembly in action. Viruses are a very successful form of self-replicating machinery and are common throughout the world: they have been found in such diverse environments as the depths of the sea [56] and even boiling geysers of Yellowstone Park [57].

Viral genomes are of the order of tens of thousands of nucleotides; *e.g.* the turnip yellow mosaic virus has 6 318 base pairs [58]). In comparison, the human genome has 2.91 billion

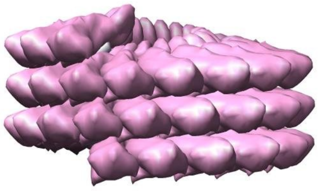

Figure 1.3: Rod shaped virus (section of a capsid from a tobacco mosaic virus [72]).

base pairs (in the euchromatic portion (gene rich nuclear DNA)) containing an estimated 26 383 genes [59]. The smallest bacterial genome on the other hand is that of *Mycoplasma genitalium*, with 580 000 base pairs and the ability to express 480 gene products [60]. This demonstrates that the number of nucleotides, and thus genes, is incredibly small in viruses when compared to other biotic systems. However, there are some large viruses, and some small prokaryotes; for example, the *Tipula* iridescent virus and the *Vaccinia* virus have approximately 1.5×10^8 u of DNA, which is greater than that of some bacterial genomes [61, 62]. Nevertheless, we will not be concentrating on these in this thesis.

Structure

Viruses are incredibly small and simple. Their diameter is of the order of 100 nm [63]. This compares with the average bacterial cell being on the micron length scale [64], and a human adipose cell, which has a diameter of approximately 100 µm [65]. Viruses consist of a nucleic acid (which may be RNA or DNA, depending on the type of virus [66]) encapsulated in a protein coat [67]. Some viruses also have a membrane envelope associated with them (created from the host's cell membrane), which has glycoproteins protruding from it, to allow it to detect binding to a host cell [68].

As mentioned previously, the small size and thus compact genome of some viruses means that they are only able to contain enough information to encode a few proteins [62, 69], and this often limits their possible structures and complexity. Despite the limitation of genetic material storage to code for proteins, viruses come in many shapes and sizes [70, 71] and are able to self-assemble into monodisperse shapes under a variety of conditions. Viruses are generally classified into two structural families: rod-shaped viruses (for example, the tobacco mosaic virus in figure 1.3 [72]) and spherical viruses (figure 1.4). The most common family of viruses is the spherical viruses [70]. These form protein capsids with icosahedral symmetry with genetic material encapsulated within. These are the inspiration for our work on encapsulation (chapter 3).

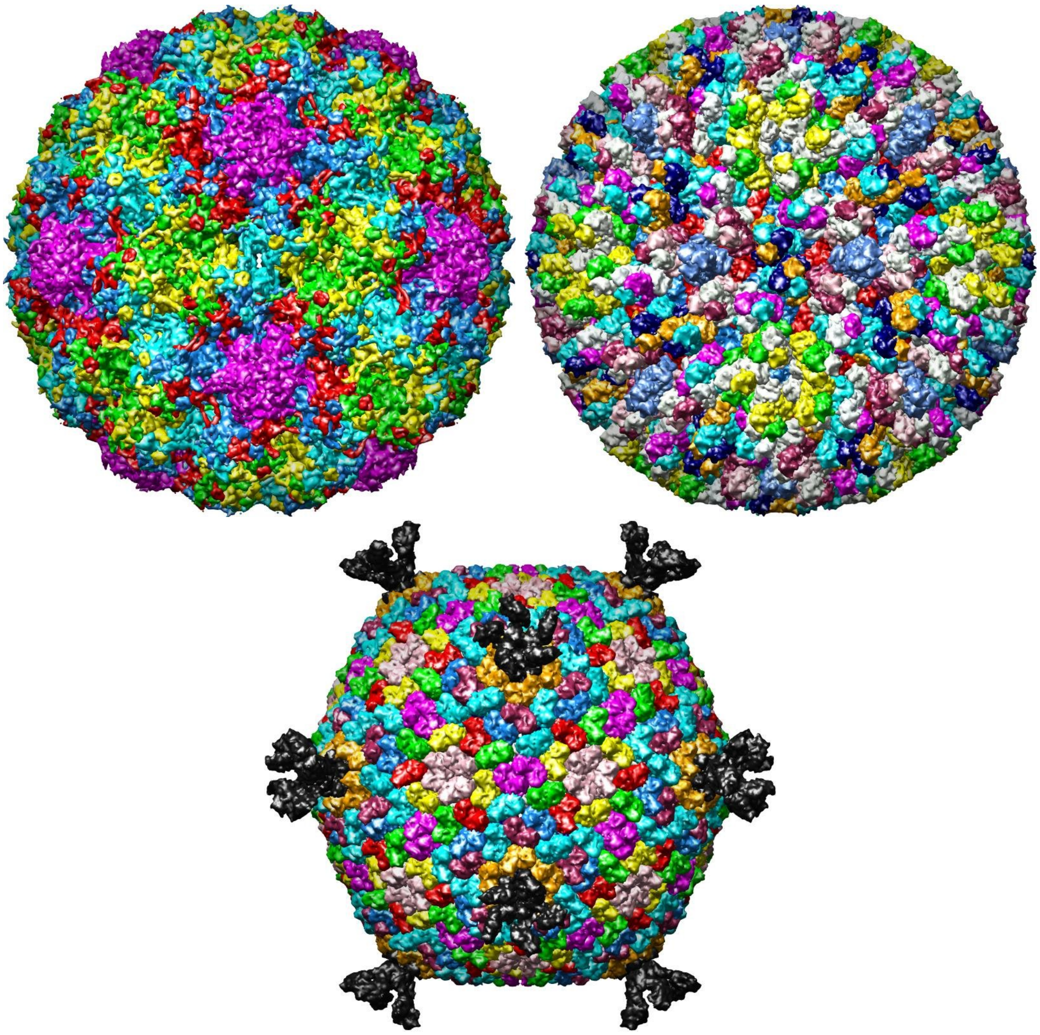

Figure 1.4: Different icosahedral viruses reproduced from [70] (top left: bovine papillomavirus type 1 outer capsid, $T = 1$, made of 60 subunits; top right: bluetongue virus core, $T = 13$, made of 780 subunits; bottom: entire structure of the shell of the lipid-containing bacteriophage PM2, pseudo-$T = 21$, made of 1260 subunits in the core). Note that all of the viruses have different numbers of proteins and connectivity, but have the same overall symmetry.

As figure 1.4 shows, it is possible to form an icosahedral capsid from a range of different numbers of protein monomers. Unsurprisingly, only certain values are allowed. The permitted numbers are best thought of as different ways in which the triangular faces can be made. They are known as *T-numbers.* For any virus with a T-number greater than 1, the proteins which make up each face must occupy a slightly different environment; this is termed *quasi-equivalence* [62]. A virus capsid contains $60 \times T$ monomers, and T may range from 1 (*e.g.* satellite tobacco mosaic virus [73]) to 219 (*e.g. Phaeocystis pouchetii* virus, PpV01 [74]).

Assembly

The assembly of the capsid varies greatly from virus to virus. Some capsids assemble using the virus's own genetic material as a template [75], while others form the capsid separately, then pump the genetic material into the capsid [76]. Section 1.4.1 has a summary of experimental evidence for the self-assembly of viruses, and the references there are strongly recommended to anyone interested in understanding the complexities of some viral assembly routes.

Life-cycle

Viruses have two life-cycles, known as the lytic and the lysogenic life-cycles. The lytic life-cycle involves the virus hijacking the host cell's organelles to make many copies of itself and then causing the host cell to explode and release the viral particles. The lysogenic cycle involves the virus inserting its genetic material into the DNA of the host, where it can remain in a dormant state known as a *prophage*. This means that every time the cell copies itself, the virus's genetic code is also copied. Viruses may undergo a lysogenic or lytic life-cycle as a response to their environment [77], and can easily switch between the two when the environment changes.

Why do we care?

Understanding how viruses function is interesting for a number of reasons. There is the obvious fact that if we understand the formation of viruses, then we are more able to develop anti-viral drugs [78] and indeed recently *in vitro* experimental screening techniques have been published for interfering with this process [79, 80].

Empty viral capsids are also used in some vaccines [81]. To the immune system, these appear far more similar to an active virus than an isolated viral capsid protein (they are said to be more *immunogenic*), and thus are more effective in vaccines [82, 83]. Viruses also have potential uses as cargo carriers in a wide range of scenarios [84–87]. Other uses of viruses and virus-like particles include applications in immunoassays [88], and even as tools for finding out the connectivity of the central nervous system [89].

1.4 Experimental work on self-assembly

There is a large body of work on self-assembly. We concentrate mainly on the formation of mono-disperse structures, of which viral assembly is the archetypal example. Much research is available on the final products of many self-assembly routes. However, there is less experimental data looking at the actual route of self-assembly, and those which do exist tend to concentrate on a single aspect of assembly (as the actual details are complex). A small number, however, do look at the overall process of assembly. The important and interesting areas of research are summarised below.

1.4.1 Viruses

Virus assembly is one of the archetypal examples of monodisperse self-assembly, and is one of the most well-studied. This interest is both due to the fundamental interest in understanding how the assembly actually occurs, but also as these are systems of great pharmaceutical interest (viruses are responsible for many deaths worldwide [90] and very few successful anti-viral drugs exist, as viral resistance develops rapidly [91]) and are easy to characterise (they generally crystallise well, which came as a surprise to early researchers [92]). Studies have shown that some virus capsids are able to form reversibly *in vitro* in the absence of a nucleic acid [93]; however, most also require the presence of their genomic material (presumably because they have evolved to encapsulate their genome). Other experiments on viruses have shown that the capsid proteins of a virus can interact with its nucleic acid, thus changing the structure of the RNA and then the pathway of self-assembly [94]. In fact, the formation pathway changes in the presence of RNA [95].

Viral self-assembly of certain viruses is templated either by their genetic material [94] or by scaffold proteins [96], and may occur hierarchically [97]: two topics which we examine in detail in this thesis.

1.4.2 Other biological systems

Many other biological systems also self-assemble. For example, the creation of a new human life involves a significant amount of self-assembly, both on the intra- and inter-cellular scales. The formation of many protein complexes occurs by self-assembly [98], even of complex molecular machinery [99, 100]. A good illustration of the complexity often associated with such biological self-assembly is the animation of the assembly of the bacterial flagellum motor producted by the group of Keichi Namba [101].

Although many protein complexes exist and assemble in solution, there are also interesting examples of self-assembly of proteins on membranes. Furthermore, their presence on a membrane potentially allows much more direct visualization of the assembly mechanism using single-molecule techniques. For example, in chapter 5, we have attempted to use our simple model to examine the formation of heptameric pores of α-haemolysin, a system for which there is particularly interesting recent experimental results [8].

In the cell DNA is mainly found in its iconic double helical form, which is based on the ability of complementary sequences of DNA to very reliably self-assemble. However, this ability to self-assemble is increasingly being used by scientists to create an array of nano-structures in a way that is easy to program through the DNA sequence. The most interesting examples of this recent work include the formation of a tetrahedron from four DNA strands by the Turberfield group in Oxford [102], and the self-assembly of DNA star motifs into a variety of polyhedra [103–106].

1.4.3 Synthetic systems

The use of directional chemical bonds to form co-ordinated clusters has been well characterised since the development of transition metal chemistry [107]. Recent developments have meant that actual building blocks can be designed to allow a "Lego" of different base components. These can be made to assemble into certain high-symmetry clusters by varying the choice of conditions and precursors in the mixture [108]. The ideas learnt through self-assembly of chemically bonded species are often applied when designing self-assembling systems, even to other length scales; this is termed *molecule-mimetic chemistry* [16]. Much of the research looks at the possible outcomes and uses of the assembly pathway, rather than looking at the particulars of the pathway itself.

A more macroscopic and rather interesting example of self-assembly is the so-called *washing machine experiment* of Burnley and Cox [109]. They used magnetic interactions to create a system able to self-assemble into a dodecahedral cluster by use of a templating mechanism analogous to viral assembly. Their system was placed into a drum and agitated, and correct assembly of dodecahedra was seen in certain circumstances. The final clusters were of the order of a few centimetres in diameter. Their results showed that globular, internal templates (in this case, a bundle of wire wool) increased the size of the clusters formed, but did not help the overall yield, whereas a flexible, external template did increase the yield of the target structure. We explore the idea of templating further in chapter 3.

1.5 Simulation work on self-assembly

As mentioned previously, the difficulties of creating perfect patchy colloids particularly in sufficient quantities and with sufficiently well-defined interactions has meant that the majority of detailed work in this field is in the area of simulation.

1.5.1 Within our group and our direct collaborators

Our group is interested in the simulation of patchy particles, in the gas phase (or solution) and their crystallisation.

Patchy particles

Previous research within our group has looked at a number of areas, including the evolution of protein complexes [4, 110], the formation of clusters with the use of torsional constraints [111], the differences in stabilities of certain cluster shapes [112], and looking into the formation of protein fibres [42].

We have also published a number of papers which form part of the research in this thesis [6, 7].

Crystallisation

More recently, a member of the group has begun to look at the phenomenon of nucleation in patchy particles using a novel technique known as forward flux sampling [113, 114].

1.5.2 Other groups

There is a large body of work available on the simulation of self-assembling systems. The simulation results may be obtained directly by simulating the particles' change in position (using molecular dynamics or Monte Carlo, termed *direct simulation*), or by simulating the changes in population by the use of kinetic rate equations (termed *indirect simulation*).

Direct simulation

Research in other groups has covered a wide range of areas, including the assembly of viral particles. Simulations have mainly concentrated on empty capsids [115–126], with some simulation work on filled capsids also becoming more common recently [127–130]. Some all-atom molecular dynamics simulations have even been carried out for $T = 1$ viruses [73].

Francesco Sciortino's group has also researched much into the thermodynamics and kinetics of patchy particles, including the assembly of binary mixtures of colloidal particles [131], the importance of fluid-crystal coexistence [132], phase diagrams and the possibility of phase re-entrance [133]. The group's articles are strongly recommended to anyone wishing to get a clear general overview of the subject, especially the recent review article [134].

Indirect simulation and theoretical treatment

Other theoreticians have also looked into the assembly of both empty [119, 120, 135–139] and filled [140–142] viral capsids from this perspective. They generally make an attempt to fit their results closely to the experimental data available on viruses. We are more interested in general rules of assembly, so we do not attempt this.

Some recent papers have looked at the use of statistical mechanical models to find the phase diagrams of patchy particles [143], and we have used a similar approach in chapter 6.

Chapter 2

Methods and model

2.1 The patchy particle model

Our model, used throughout this project, consists of spherical particles with patches on the surface, of a certain Gaussian width, σ^{PW}, with optional torsional specificity. The interaction energy, U, of the patchy particles along the inter-particle vector, $\boldsymbol{r}$, is of a Lennard-Jones type, modified in the attractive region only by an angular factor U^{ANG} (equation 2.1). Only the two patches which minimise the angle θ between the inter-particle vector, $\boldsymbol{r}$, and the vector connecting the centre of the particle to the patch on the particle's surface, $\boldsymbol{A}$, are used to calculate interactions; or if different patch types (A, B) are possible, the pair which minimises U_{ij} is used. The parameters of the potential are the location of the patches, the size of the particle, σ^{LJ}, and the strength of their interactions. Subscripts i, j refer to specific particles; ϕ, χ refer to particle types; and A, B refer to the particle patch types. The full potential is

$$U_{ij,\,AB,\,\phi\chi}(\boldsymbol{r}_{ij}) = \begin{cases} U^{\mathrm{LJ}}_{AB,\,\phi\chi} & |\boldsymbol{r}_{ij}| < \sigma^{\mathrm{LJ}}_{\phi\chi} \\ U^{\mathrm{LJ}}_{AB,\,\phi\chi} \times U^{\mathrm{ANG}}_{ij,\,AB} & |\boldsymbol{r}_{ij}| \geq \sigma^{\mathrm{LJ}}_{\phi\chi} \end{cases} \tag{2.1}$$

$$U^{\mathrm{LJ}}_{AB,\,\phi\chi}(\boldsymbol{r}_{ij}) = 4\,\varepsilon_{AB}\left[\left(\frac{\sigma^{\mathrm{LJ}}_{\phi\chi}}{|\boldsymbol{r}_{ij}|}\right)^{12} - \left(\frac{\sigma^{\mathrm{LJ}}_{\phi\chi}}{|\boldsymbol{r}_{ij}|}\right)^{6}\right] \tag{2.2}$$

$$U^{\mathrm{ANG}}_{ij,\,AB}(\boldsymbol{r}_{ij}) = G_{ij,\,A}(\boldsymbol{r}_{ij})G_{ji,\,B}(\boldsymbol{r}_{ji}) \tag{2.3}$$

$$G_{ij,\,A}(\boldsymbol{r}_{ij}) = \exp\left(-\frac{\theta^2_{A,\,\boldsymbol{r}_{ij}}}{2(\sigma^{\mathrm{PW}}_{A})^2}\right) \tag{2.4}$$

Note that only the particle possesses a volume; the patches do not have any true volume, though they are often diagrammatically shown as such for clarity in this thesis. For particles of differing sizes, the arithmetic mean is used in the interactions, $\sigma_{\phi\chi}^{\mathrm{LJ}} = (\sigma_{\phi\phi}^{\mathrm{LJ}} + \sigma_{\chi\chi}^{\mathrm{LJ}})/2$.

2.1.1 Torsional constraints

Additionally to the energy term imposed by our model on the relative positions of patchy particles, previous research within the group [41, 42, 111] has used a *torsional* constraint. This constraint is multiplicative with the previous angular constraint. Torsional constraints are common to proteins, which have multiple bonding sites leading to preferred binding orientations. While much of our work is inspired by biological systems, we have attempted to look mainly at ways of getting patchy particles to form in the absence of torsional constraints, as the current generation of experimentally formed patchy colloids is not able to form torsional interactions [13, 29].

Our particle interaction is modified as follows:

$$U_{ij,\,AB,\,\phi\chi}(\boldsymbol{r}_{ij}) = \begin{cases} U_{AB,\,\phi\chi}^{\mathrm{LJ}} & |\boldsymbol{r}_{ij}| < \sigma_{\phi\chi}^{\mathrm{LJ}} \\ U_{AB,\,\phi\chi}^{\mathrm{LJ}} \times U_{ij,\,AB}^{\mathrm{ANG}} \times U_{ij,\,AB}^{\mathrm{TOR}} & |\boldsymbol{r}_{ij}| \geq \sigma_{\phi\chi}^{\mathrm{LJ}} \end{cases} \tag{2.5}$$

$$U_{ij,\,AB}^{\mathrm{TOR}}(\boldsymbol{r}_{ij}) = H_{ij,\,A}(\boldsymbol{r}_{ij}) H_{ji,\,B}(\boldsymbol{r}_{ji}) \tag{2.6}$$

$$H_{ij,\,A}(\boldsymbol{r}_{ij}) = \exp\left(-\frac{(\varphi_{A,\boldsymbol{r}_{ij}} - \varphi_{\mathrm{ref}})^2}{2(\sigma_A^{\mathrm{TOR}})^2}\right) \tag{2.7}$$

The torsional angle, φ, is best explained diagrammatically (see figure 2.1). The angle of interest (φ) is the angle between the projection of the two torsional vectors of the particles ($\boldsymbol{T}$) onto a plane perpendicular to the interparticle vector ($\boldsymbol{r}$).

For torsional vectors $\boldsymbol{T}_{A,B}$ on each particle with an interparticle vector $\boldsymbol{r}$, we find the torsional angle (φ) using a number of simple vector calculations. First, we find the unit vector of the cross products between the interparticle vector ($\boldsymbol{r}$) and the two torsion vectors ($\boldsymbol{T}_A$ and $\boldsymbol{T}_B$), then find the angle between them. In the event that our reference torsion angle, φ_{ref}, is zero, we have the torsional angle. However, if we wish to have the energy minimum offset (for example, if we are forming a helix), then it it necessary to further use

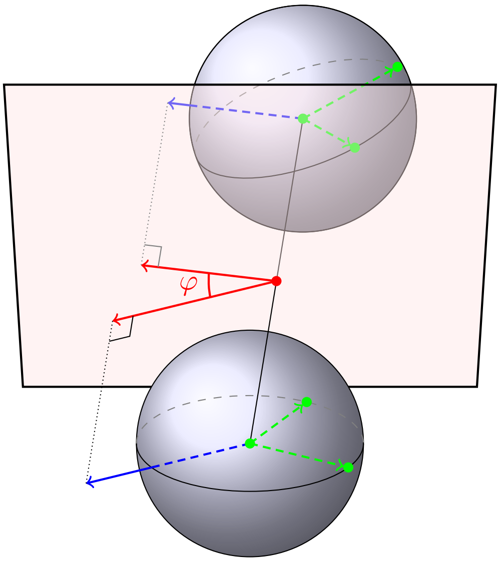

Figure 2.1: Torsional angle calculation (adapted from [41]).

vector mathematics to find the direction of the angle.

$$\boldsymbol{a} = \frac{\boldsymbol{r} \times \boldsymbol{T}_A}{|\boldsymbol{r} \times \boldsymbol{T}_A|} \tag{2.8}$$

$$\boldsymbol{b} = \frac{\boldsymbol{r} \times \boldsymbol{T}_B}{|\boldsymbol{r} \times \boldsymbol{T}_B|} \tag{2.9}$$

$$\chi = \arccos(\boldsymbol{a} \cdot \boldsymbol{b}) \tag{2.10}$$

$$\boldsymbol{c} = \boldsymbol{a} \times \boldsymbol{b} \tag{2.11}$$

$$g = \boldsymbol{c} \cdot \boldsymbol{r} \tag{2.12}$$

$$\varphi = \begin{cases} -\chi & g < 0 \\ \chi & g \geq 0 \end{cases} \tag{2.13}$$

This approach is one of many possible methods for finding the torsional angle. An advantage of this approach is that it is computationally relatively fast to evaluate. The torsional vectors are normally defined as the average vector of all the patch vectors, though the program allows several other possibilities, all of which were necessary for modelling evolution of complexes [41].

2.1.2 Internal degrees of freedom

In order to simulate internal rearrangements in particles, or the binding of an initiator, we allow the system to be set up with different possible patch configurations for a particle. A new move type was created that allowed particles to change their ‘internal’ structure configurations (which includes the number of patches, their position, their types, and even the interaction terms and sizes of the particles). We term these moves *flipping moves.* This is used and explained in more detail in chapter 5.

2.1.3 Modifying the range of interactions

The initial work has been carried out using a Lennard-Jones potential, as this is a common model for interactions throughout the chemical and biological world. This has been extended in a limited number of cases to a Lennard-Jones-esque potential, allowing the range of the interaction to be varied using a single parameter, $\digamma$, while keeping the general shape the same (see equation 2.14).

$$U^{\mathrm{LJE}}_{AB,\phi\chi}(\boldsymbol{r}_{ij}) = 4\varepsilon_{AB}\left[\left(\frac{\sigma^{\mathrm{LJE}}_{\phi\chi}}{|\boldsymbol{r}_{ij}|}\right)^{2\digamma} - \left(\frac{\sigma^{\mathrm{LJE}}_{\phi\chi}}{|\boldsymbol{r}_{ij}|}\right)^{\digamma}\right] \tag{2.14}$$

The effect of changing the range of interactions is explored in appendix B.

2.2 Simulation techniques

2.2.1 Monte Carlo

A Metropolis Monte Carlo (MC) simulation [144, 145] in the canonical (constant N, V, T) ensemble was used for the majority of dynamic simulations of the model system. The basic program used moves that were single particle translations (randomly distributed within a sphere) and rotations (using a quaternion representation and adding a perturbation), as well as cluster moves according to the virtual move MC algorithm of Whitelam and Geissler [146].

What is a Monte Carlo simulation?

Monte Carlo simulations are a common and useful tool not only in computational chemistry, but also found in such varied fields as physics and finance [147]. Monte Carlo simulations involve attempting random moves through configuration space, and the choice of whether to accept or reject them is based on how the system's properties change. In particle simulations, it is normal for the energy to be used as the property which helps to determine whether a move is accepted or rejected. Monte Carlo simulations are generally used to find thermodynamic averages of properties [145], but they can be adapted to give information on phase diagrams of systems [148, 149], chemical potential and pressure of a system [150], and even on the free energies as a function of cluster sizes [151]. The next sections explain some of the mathematics and methods which underpin this thesis.

Using Monte Carlo to simulate dynamics

Generally, Monte Carlo simulations are used to find thermodynamic equilibrium data in molecular simulations. In our coarse-grained models, we want to be able to reproduce the diffusive dynamics expected of colloids, proteins and nanoparticles in solution. Monte Carlo simulations using small local moves naturally produce the relaxation of these non-equilibrium systems (to a time proportionality constant), and are thus ideal for our purposes [152, 153]. The more cumbersome molecular dynamics methods require significantly more computer time, as the calculations of forces, the derivatives of the potential energy, are required [153], and are thus computationally expensive [154].

Requirements of a Monte Carlo simulation

In order for a Monte Carlo simulation to give reliable statistics, it is necessary for a number of conditions to be met. In this derivation, we label the possible states using the Greek letters phi (ϕ) and chi (χ), the probability of occupying them is labelled ν, and the probability of a transition is known as π. The first of the requirements is that, at equilibrium, the overall probability of leaving one state must be the same as returning to it (this is simply the definition of a dynamic equilibrium) [144] (equation 2.15). This is known as *balance*.

$$\sum_i \nu(\phi) \times \pi(\phi \to \chi_i) = \sum_i \nu(\chi_i) \times \pi(\chi_i \to \phi) \tag{2.15}$$

Keeping to this condition in its broadest form would involve knowledge of many states. Thus, Monte Carlo algorithms usually also ensure that, at equilibrium, the probability of a forward step leaving one state for another is exactly the same as the probability of returning from that new state back again to the original state, a condition known as *detailed balance*. This is a sufficient condition, but not a necessary one [155]: it meets the requirement of balance, but is actually far easier to work with. It is shown mathematically in equation 2.16.

$$\nu(\phi) \times \pi(\phi \to \chi) = \nu(\chi) \times \pi(\chi \to \phi) \tag{2.16}$$

This mathematical statement of the detailed balance condition allows us to decide on our acceptance criteria for a given move. The transition probability can be written as the product of the probability of attempting a given move $\alpha(\phi \to \chi)$ and the probability that we accept the move, $\gamma(\phi \to \chi)$.

$$\pi(\phi \to \chi) = \alpha(\phi \to \chi) \times \gamma(\phi \to \chi) \tag{2.17}$$

We keep the probability of attempting any forward move to be the same as its reverse counterpart (*i.e.* we set $\alpha(\phi \to \chi) = \alpha(\chi \to \phi)$), in accordance with the original Metropolis algorithm [145, 156]. This allows us to rewrite equation 2.16 in terms of the acceptance probability, and then to find our acceptance probabilities in terms of the probability of being in a given state.

$$\nu(\phi) \times \gamma(\phi \to \chi) \times \alpha(\phi \to \chi) = \nu(\chi) \times \gamma(\chi \to \phi) \times \alpha(\chi \to \phi) \tag{2.18}$$

$$\frac{\alpha(\phi \to \chi)}{\alpha(\chi \to \phi)} = 1 \tag{2.19}$$

$$\Rightarrow \frac{\gamma(\phi \to \chi)}{\gamma(\chi \to \phi)} = \frac{\nu(\chi)}{\nu(\phi)} \tag{2.20}$$

For a system in the canonical ensemble, the probability of being in a given state is proportional to the canonical Boltzmann factor $\exp(-\beta U)$ [157], so we know that our acceptance probability must therefore be

$$\frac{\gamma(\phi \to \chi)}{\gamma(\chi \to \phi)} = \frac{\nu(\chi)}{\nu(\phi)} = \exp(-\beta(U_\chi - U_\phi)). \tag{2.21}$$

Many possibilities are available that would ensure this ratio is kept. We set that any step which decreases the energy of the system should be accepted, and any other state depends on the relative Boltzmann probability [158]; *i.e.*

$$\gamma(\phi \to \chi) = \min\left[1,\, e^{-\beta(U_\chi - U_\phi)}\right]. \tag{2.22}$$

In other ensembles, this may change, as the probability of being in a given state contains other terms [159]. These are explored in later sections.

The next requirement is that every point of the fitness landscape must be accessible by the simulation, and there must exist a pathway between any two points. This is known as *ergodicity*. For Monte Carlo runs to be valid, there must be an equivalence between time and ensemble averages, and thus ergodicity is required [160]. In simulations where there are bottlenecks which do not allow sampling between two states, special methods are required [150].

Algorithm details

We begin by considering Monte Carlo simulations in the canonical ensemble. Other ensembles and the way in which they change the algorithm are explored later. As mentioned previously, a *move* may be a translation or a rotation of a single particle. In order to ensure that detailed balance is followed, we must first require that after any forward move, the reverse must be possible. This requires that we select our particle to be moved randomly at each step, and do not sample across particles systematically. It also means that our maximum move sizes should remain constant.

Generating displacement moves Our program permits displacement moves to be selected from both within a cube and a sphere, with a set maximum cutoff, ξ. We choose our random move (Δx, Δy, Δz) using a function which generates random and arbitrary numbers ($\mathcal{R}$) in the interval $[0,1]$ as follows:

$$\Delta x = 2\xi \times (\mathcal{R}_1 - 0.5) \tag{2.23}$$

$$\Delta y = 2\xi \times (\mathcal{R}_2 - 0.5) \tag{2.24}$$

$$\Delta z = 2\xi \times (\mathcal{R}_3 - 0.5) \tag{2.25}$$

When selecting a random and arbitrary displacement within a sphere, we tried several approaches. The first involved the use of quaternions (see section 2.2.1), and the second (our chosen method) was simply to generate all of our moves within a cube, and then ignore any move vectors that had a length greater than our maximum step size ξ. When testing these methods for speed, it was found that the latter method was actually fastest (on average, approximately two random moves need to be generated each time in order to obtain one within the sphere). Most simulations in this thesis use translational moves that are randomly chosen from within a sphere; the effect of using moves distributed within a cube is shown in appendix B.1.

The cluster move algorithm [146] that we also used is discussed in section 2.2.3.

Boundary effects

When simulating a small system, the effect of the edges rapidly becomes relevant. This is because the proportion of particles near the edge compared to those in the bulk is unrepresentative in small systems compared to large systems (and we are ultimately interested in the behaviour of large systems). In order to maximise the efficiency of our simulations, periodic boundary conditions were used. This periodic boundary means that the system is treated as an infinite array of boxes with the unit cell equal to our simulation box. When a particle moves off one edge, it simply reappears on the opposite side.

Cutoff distances

In order to prevent the energy calculations being carried out over an infinitely sized box, we introduce a *cutoff* distance and shift the potential so that at this cutoff, there is no

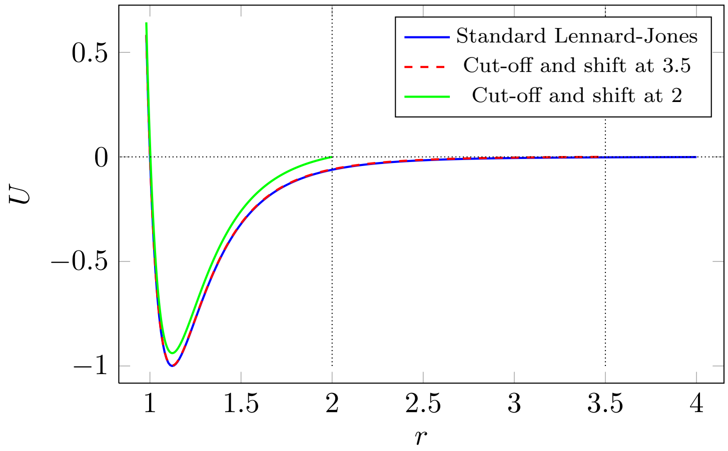

Figure 2.2: Comparison of different cutoff distances.

interaction between particles. This is known as 'cutoff and shift'. The potential becomes

$$U^{\text{cutoff and shift}}_{ij,\,AB,\,\phi\chi}(\boldsymbol{r}_{ij}) = \begin{cases} U^{\text{LJ, cutoff}}_{AB,\,\phi\chi} & |\boldsymbol{r}_{ij}| < \sigma^{\text{LJ,cutoff}}_{\phi\chi} \\ U^{\text{LJ, cutoff}}_{AB,\,\phi\chi} \times U^{\text{ANG}}_{ij,\,AB} & |\boldsymbol{r}^{\text{cutoff}}_{ij}| \geq |\boldsymbol{r}_{ij}| \geq \sigma^{\text{LJ,cutoff}}_{\phi\chi} \\ 0 & |\boldsymbol{r}_{ij}| \geq |\boldsymbol{r}^{\text{cutoff}}| \end{cases} \tag{2.26}$$

where $\sigma^{\text{LJ,cutoff}}_{\phi\chi}$ is the shifted Lennard-Jones radius, and

$$U^{\text{LJ, cutoff}}_{AB,\,\phi\chi}(\boldsymbol{r}_{ij}) = U^{\text{LJ}}_{ij,\,AB,\,\phi\chi}(\boldsymbol{r}_{ij}) - U^{\text{LJ}}_{ij,\,AB,\,\phi\chi}(\boldsymbol{r}^{\text{cutoff}}). \tag{2.27}$$

As figure 2.2 shows, when we use a short cutoff distance of $2\,\sigma^{\text{LJ}}$, we change the shape of the potential, though not significantly. However, when using larger cutoff distances, there is little difference. Alternative methods for the cutoff include σ matching and depth matching, but, for the purposes of our simple model system, there is no need to increase the complexity of the interaction potential. In all our dynamic simulations, we used a cutoff of at least $3.5\,\sigma$, which changes the well depth by $0.2\,\%$. A smaller cutoff of $2.5\,\sigma^{\text{LJ}}$ was used for runs involving the Gibbs ensemble (section 2.2.7) and one of $2\,\sigma^{\text{LJ}}$ was used for certain Wang-Landau runs (we consider the effect this has in appendix B.4).

Quaternions

In order to represent the current rotational position of a molecule, three parameters are required. Normally, these are the three Euler angles ψ, θ, and ϕ. There are several conventions for the naming and definition of these angles; we use the following [161] (for a particle with internal axes a, b, c situated at a certain rotation in the simulation box's axes x, y, z):

ψ: The angle between the x-axis and the vector formed from $\hat{\boldsymbol{z}} \times \hat{\boldsymbol{c}}$

θ: The angle between the particle's axis $\boldsymbol{c}$ and the simulation box's axis $\boldsymbol{z}$.

ϕ: The angle between the vector formed from $\hat{\boldsymbol{z}} \times \hat{\boldsymbol{c}}$ and the $\boldsymbol{a}$ axis.

Using Euler angles in a simulation requires the calculation of trigonometric functions. These are slow to calculate on a computer (compared to multiplication) [162], and also suffer from some numerical inaccuracy.

In order to minimise the need to calculate computationally expensive trigonometric functions, we used quaternions [163]. A quaternion is a four dimensional vector which is restricted to being on the surface of a four dimensional sphere (thus giving 3 degrees of freedom). The standard nomenclature of a quaternion is a four dimensional column matrix [164] (equation 2.28), or a complex 2-dimensional matrix [165] (equation 2.29). Note that the multiplication rules are more complicated than simple matrix multiplication for the column form (see equation 2.34).

$$q = \begin{pmatrix} w \\ x \\ y \\ z \end{pmatrix} \tag{2.28}$$

$$q = \begin{pmatrix} w + \mathrm{i}x & y + \mathrm{i}z \\ -y + \mathrm{i}z & w - \mathrm{i}x \end{pmatrix} \tag{2.29}$$

The components can be related to the Euler angles as follows [144] (adapted from [113]):

$$w = \cos\left(\frac{\theta}{2}\right)\cos\left(\frac{\phi+\psi}{2}\right) \tag{2.30}$$

$$x = \sin\left(\frac{\theta}{2}\right)\cos\left(\frac{\phi-\psi}{2}\right) \tag{2.31}$$

$$y = \sin\left(\frac{\theta}{2}\right)\sin\left(\frac{\phi-\psi}{2}\right) \tag{2.32}$$

$$z = \cos\left(\frac{\theta}{2}\right)\sin\left(\frac{\phi+\psi}{2}\right) \tag{2.33}$$

When multiplying two quaternion vectors, the following is used [165] (adapted from [113]):

$$q_a \times q_b \equiv \begin{pmatrix} a_w \\ a_x \\ a_y \\ a_z \end{pmatrix} \times \begin{pmatrix} b_w \\ b_x \\ b_y \\ b_z \end{pmatrix} = \begin{pmatrix} a_w b_w - a_x b_x - a_y b_y - a_z b_z \\ a_w b_x + a_x b_w + a_y b_z - a_z b_y \\ a_w b_y - a_x b_z + a_y b_w + a_z b_x \\ a_w b_z + a_x b_y - a_y b_x + a_z b_w \end{pmatrix} \tag{2.34}$$

This is important when rotating a quaternion. The rotation matrix necessary to rotate by a given quaternion is as follows [165] (adapted from [113]):

$$\boldsymbol{R} = \begin{pmatrix} w^2 + x^2 - y^2 - z^2 & 2xy - 2wz & 2wy + 2xz \\ 2wz + 2xy & w^2 - x^2 + y^2 - z^2 & 2yz - 2wx \\ 2xz - 2wy & 2wx + 2yz & w^2 - x^2 - y^2 + z^2 \end{pmatrix} \tag{2.35}$$

In order to generate a random quaternion, the procedure according to Vesely was used [163]:

1. Generate a pair of random numbers a_1 and a_2 in the range $[-1, 1]$, and repeat until $S_1 = a_1^2 + a_2^2 < 1$ (this is actually a computationally faster method than attempting to find the allowable values after generating the first random number).
2. Generate another pair of random numbers a_3 and a_4 similarly, with $S_2 = a_3^2 + a_4^2 < 1$.
3. The quaternion components are given by $w = a_1$, $x = a_2$, $y = a_3 \times \sqrt{\frac{1 - S_1}{S_2}}$ and $z = a_4 \times \sqrt{\frac{1 - S_1}{S_2}}$.

This procedure has been shown to be approximately five times faster than the equivalent generation of a random, arbitrary Euler rotation [163]. When combined with the additional increase in computational efficiency given by the need to calculate fewer trigonometric functions, it is clear that the use of quaternions is sensible for all particle simulation applications [166].

Random numbers, seeds and initial configurations

Initially, all particles are placed on an orthorhombic lattice (the size of each side of the cuboid is chosen such that the maximum amount of space is filled). The particles are then simulated without attractions (but with the repulsive region of the interactions only) for 1000 cycles. This ensures that they are randomly placed in the box. In order to generate random numbers, we used a third party's code, which used the method reviewed by Marsaglia *et al.* [167]. The initial random seeds were either generated as a function of the instantaneous free memory on a node and the time on the system clock, or taken from a list of random seeds (ensuring each simulation in a group of runs had a different random seed).

2.2.2 Ensembles used

When performing simulations, one must choose an intensive and two extensive factors to keep constant. In statistical mechanics, the selection of these variables is known as the *ensemble*. In order for an ensemble to be useful in a Monte Carlo simulation, the underlying probability distribution must be time invariant (a *Gibbs state*). It must also be possible to get between any two points on the probability surface in a finite number of steps (the system must be *ergodic*).

Our program (MCP) can run in three different simple ensemble modes, each one useful to find out different information about a system.

Canonical ensemble

This ensemble is the one we use the most. Here, the number of particles, the volume and the temperature are kept constant. This is the equivalent to having a closed system of constant volume which can exchange energy with a heat bath; for example, a closed reaction vessel. As explained in section 2.2.1, an acceptance probability for a given move which maintains detailed balance is given in equation 2.21.

Grand canonical ensemble

In this ensemble, the system is not closed and particles may be exchanged with an external "reservoir" which is kept at constant chemical potential. An example of a physical phenomenon equivalent to this ensemble is a surface in contact with a large reservoir of gas molecules in the atmosphere (which acts as the infinite reservoir). The Boltzmann probability (ν) of being in a state of given chemical potential (μ), volume (V) and temperature (T) is [144]

$$\nu(\mu, N, V, T, U) \propto \frac{\exp(\beta\mu N)V^N}{\Lambda^{3N}N!}\exp(-\beta U). \tag{2.36}$$

When we make a displacement or rotation move, our acceptance probability remains unchanged compared to equation 2.21 (as all of the other factors cancel). However, we additionally have acceptance probabilities for the addition and removal of a particle that are given by

$$\gamma(N_{\text{old}} \to N_{\text{new}}) = \min\left[1, \frac{\exp(-\beta(U_{\text{new}} - U_{\text{old}})) \times V^{N_{\text{new}}-N_{\text{old}}} \times \exp(\beta\mu(N_{\text{new}} - N_{\text{old}}))}{\Lambda^{3(N_{\text{new}}-N_{\text{old}})}\dfrac{N_{\text{new}}!}{N_{\text{old}}!}}\right]. \tag{2.37}$$

These can be simplified depending on whether particles are added or removed:

$$\gamma(N \to N+1) = \min\left[1,\ \exp\left(-\beta\Delta U + \ln V\right) \times \frac{\exp\left(+\beta\mu\right)}{(N+1)\lambda^3}\right] \tag{2.38}$$

$$\gamma(N \to N-1) = \min\left[1,\ \exp\left(-\beta\Delta U - \ln V\right) \times \exp\left(-\beta\mu\right) \times N\lambda^3\right] \tag{2.39}$$

As most of our systems are at low density, we use a simple insertion technique where we randomly attempt to place a particle with random orientation and position on addition. Other insertion methods are also possible for more complex systems [168], but they are not used here.

Isothermal-isobaric ensemble

This is the isobaric form of the canonical ensemble [169]. Rather than the volume remaining constant, we fix the pressure. The system is able to change the size of the box; moves attempting this are known as *volume moves.* A physical analogue to the simulation would be a reaction taking place inside a cubic balloon (that had no elastic effects from the balloon's skin). The program can change the box keeping all of the edges the same length (cubic box), or allow the three dimensions to vary independently (orthorhombic box).

The probability of being in a given state in this ensemble is:

$$\nu(p, N, V, T, U) \propto V^N \exp(-\beta p V)\exp(-\beta U) \tag{2.40}$$

Displacement and rotation moves have the same acceptance criterion as the canonical ensemble (equation 2.21). However, for volume moves, there are two possibilities.

We can sample either the logarithm of the volume ($\ln(V)$) or the volume directly (V). Generally, we used sampling in $\ln(V)$, as this means that the size of volume move steps is varied with the size of the box, and thus gives a better acceptance ratio. When we change the value of $\ln(V)$ by a randomly chosen constant, this is the equivalent to multiplying the volume by a randomly chosen amount.

When generating the move, it is important to remember that we have chosen to set the probability of generating a move to be equal for any forward or reverse move. Therefore, for both our sampling schemes, we select the move size to be a random amount in the interval $[-\xi_{\mathrm{vol}}, \xi_{\mathrm{vol}}]$, added to either the volume or the logarithm of the volume.

The acceptance of a volume move depends on our sampling. When we aim to sample the volume V, we get the acceptance rule [170]

$$\gamma(V_{\mathrm{old}} \to V_{\mathrm{new}}) = \min\left\{1,\ \exp\left[-\beta\left(\Delta U + p\Delta V\right) + N\ln\left(\frac{V_{\mathrm{new}}}{V_{\mathrm{old}}}\right)\right]\right\}, \tag{2.41}$$

while sampling in $\ln(V)$ yields the acceptance rule [144, 171]

$$\gamma(V_{\text{old}} \to V_{\text{new}}) = \min\left\{1, \exp\left[-\beta\left(\Delta U + p\Delta V\right) + (N+1)\ln\left(\frac{V_{\text{new}}}{V_{\text{old}}}\right)\right]\right\}. \tag{2.42}$$

2.2.3 Cluster moves

Although single-particle move Monte Carlo algorithms approximate diffusive dynamics, the relative diffusion rates of different-sized clusters is incorrect, with larger clusters moving too slowly. This deficiency would be a particular problem for simulating hierarchical assembly (chapter 4), where assembly requires intermediates to diffuse sufficiently rapidly. In order to correct this problem, we used a cluster move system as created by Whitelam and Geissler [146, 172] that is specifically designed to give more realistic diffusion rates for clusters of different sizes. The algorithm not only complies with the Monte Carlo requirements for balance, but goes two steps further and follows super-detailed balance [173]. We derive and explain it below for the canonical ensemble.

Dynamic cluster formation

The algorithm of Whitelam and Geissler [146, 172] aims to form clusters using a *virtual move* Monte Carlo algorithm. It is applicable to pairwise-interacting particles in its simplest form (described below). It works by attempting to move a single particle, and seeing the effect this movement has on the interaction energies of individual pairs of particles (within an interaction cutoff) compared with moving them together. The algorithm is as follows:

1. Select a seed particle i uniformly and randomly and assign a *move map* (translation or rotation).

2. Attempt a *virtual move* from state ϕ to state χ.

3. Link particles j and i with probability p_{ij} according to

$$p_{ij}(\phi \to \chi) = I_{ij}^{\phi} \max\left[0,\, 1 - \exp\left(\beta U_{\text{Cluster}}^{ij} - \beta U_{\text{Individual}}^{ij}\right)\right], \tag{2.43}$$

   where I_{ij}^{ϕ} is a binary variable for the interaction between molecules i and j (*i.e.* it is 1 if they are within the interaction cutoff distance and 0 otherwise), U_{Cluster}^{ij} is the energy of interaction between i and j if the particles move in concert (with the same move map – this is equal to their interaction energy if they remain still in the case that they are in an isotropic field), and $U_{\text{Individual}}^{ij}$ is their energy of interaction if they do not (*i.e.* i moves while j stays still).

4. Decide whether to reject or accept the move.

Our detailed balance condition requires that

$$\nu(\phi) \times \pi(\phi \rightarrow \chi) = \nu(\chi) \times \pi(\chi \rightarrow \phi). \tag{2.44}$$

Using the same symbols as before, where ν is the probabilty of being in a given state, the probability of a transition is known as π and is equal to the product of the probability of attempting a given move $\alpha(\phi \rightarrow \chi)$ and the probability we accept the move, $\gamma(\phi \rightarrow \chi)$.

This leads to equation 2.45 (which is similar, but not identical to equation 2.21, as the probability of generating a move is not the same for all moves).

$$\frac{\gamma(\phi \rightarrow \chi)}{\gamma(\chi \rightarrow \phi)} = \frac{\nu(\chi)}{\nu(\phi)} \times \frac{\alpha(\chi \rightarrow \phi)}{\alpha(\phi \rightarrow \chi)} \tag{2.45}$$

We can write our probability of generating a move, α, using equation 2.43. We can write this as a sum over all probabilities of the possible ways we could build our cluster (R). Each way of building a cluster is simply the product of the probabilities of generating a certain seed particle $(W_{\mathrm{seed}}(\phi))$, not making bonds outside of the cluster ($q_{ij} = 1 - p_{ij}$, multiplied over all unformed bonds, X), and of making bonds in the cluster (p_{ij}), multiplied over all formed bonds, Y. This gives us equation 2.46.

$$\alpha(\phi \rightarrow \chi) = \sum_{R} \left(W_{\mathrm{seed}}(\phi) \times \prod_{X}^{R} q_{ij}(\phi \rightarrow \chi) \times \prod_{Y}^{R} p_{ij}(\phi \rightarrow \chi) \right) \tag{2.46}$$

If we had to keep track of all of the different ways in which a cluster could be formed, there would be a lot of calculations needed. Instead, we introduce the idea of super-detailed balance [173]. In detailed balance, we made sure that the probabilities of a forward and reverse move between any two states maintained equilibrium. In super-detailed balance, we go one step further and set the probabilities of completing a forward or reverse move *in a certain way* such that we maintain equilibrium. This obeys detailed balance, but simplifies keeping track of many different routes by which a cluster can be formed or moved. Our new requirement for super-detailed balance can we written in terms of probabilities of generating or accepting a move in a certain way, R, which we write as $\alpha(\phi \rightarrow \chi|R)$ and $\gamma(\phi \rightarrow \chi|R)$, giving us

$$\frac{\gamma(\phi \rightarrow \chi|R)}{\gamma(\chi \rightarrow \phi|R)} = \frac{\nu(\chi)}{\nu(\phi)} \times \frac{\alpha(\chi \rightarrow \phi|R)}{\alpha(\phi \rightarrow \chi|R)}. \tag{2.47}$$

This leads us to:

$$\alpha(\phi \to \chi|R) = W_{\text{seed}}(\phi) \times \prod_X^R q_{ij}(\phi \to \chi) \times \prod_Y^R p_{ij}(\phi \to \chi). \tag{2.48}$$

If we have canonical Boltzmann weights, and select our seed particle arbitrarily and randomly (so that $W_{\text{seed}}(\phi) = W_{\text{seed}}(\chi)$), then we are able to write equation 2.47 as:

$$\frac{\gamma(\phi \to \chi|R)}{\gamma(\chi \to \phi|R)} = \exp^{-\beta(U_\chi - U_\phi)} \times \frac{\prod_X^R q_{ij}(\chi \to \phi) \times \prod_Y^R p_{ij}(\chi \to \phi)}{\prod_X^R q_{ij}(\phi \to \chi) \times \prod_Y^R p_{ij}(\phi \to \chi))} \tag{2.49}$$

We make two further modifications to the probabilities before giving our final acceptance criterion. As long as we do so symmetrically, we do not affect our super-detailed balance. The algorithm in its current form will preferentially move larger clusters (as the probability of picking a particle in a large cluster as the seed is higher than picking one that is not). This will not mimic diffusive dynamics (which is our aim in our Monte Carlo simulations). To deal with this, we select a random cutoff size, n_c, before each move and reject any move which attempts to move a cluster larger than this. This adds a sharp cutoff function $\Theta(n - n_\text{c})$ into our acceptance ratios ($\Theta(n - n_\text{c})$ is 1 if $n \leq n_\text{c}$ and zero otherwise). We select n_c by taking the integer floor of the inverse of a random floating point number $\mathcal{R}$ in the range 0 to 1, *i.e.* $n_\text{c} = \lfloor \frac{1}{\mathcal{R}} \rfloor$. This perfectly eliminates the unfair bias (equation 2.50).

$$\begin{aligned} P(\text{generation of cluster with size } n) &\propto n \times \int_{\mathcal{R}=0}^{\mathcal{R}=\frac{1}{n}} \mathrm{d}\mathcal{R} \\ &\propto 1 \;\; \forall n \geq 1 \end{aligned} \tag{2.50}$$

We also need to deal with some hydrodynamic issues. The equation we have currently derived allows clusters to move the same distances, regardless of their cross-sectional area. This does not mimic diffusive dynamics. This is easily rectified by applying Stokes's laws (equation 2.51, [146]) to the acceptance rate of the move, and by noting that the effective radius can be found using equation 2.52 ($\hat{\boldsymbol{n}}$ is the unit movement vector (direction of translation or axis of rotation), $\boldsymbol{r}_\text{c}$ is our centre of mass position vector (for translations) or moment of inertia vector (for rotations), $\boldsymbol{r}_i$ is the position of each particle, and the average occurs over all N particles in the cluster). D_t is the diffusion coefficient for translational moves and D_r is the diffusion coefficient used in rotational moves. R_0 is the size of the smallest monomeric particle.

$$\begin{aligned} D_{\mathrm{t}}(R) &= \frac{\mathit{\Gamma}}{3R} = \frac{D_{\mathrm{t}}(R_0) \times R_0}{R} \\ D_{\mathrm{r}}(R) &= \frac{\mathit{\Gamma}}{8R^3} = D_{\mathrm{r}}(R_0) \left(\frac{R_0}{R} \right)^3 \end{aligned} \tag{2.51}$$

$$(R - R_0)^2 = \sum_i \frac{|(\boldsymbol{r}_i - \boldsymbol{r}_{\mathrm{c}}) \times \hat{\boldsymbol{n}}|^2}{N} \tag{2.52}$$

We generally set $D_r(R_0) = D_t(R_0) = 1$, as only the relative rates matter.

This gives us our final probability of acceptance as being:

$$\begin{aligned} \gamma(\phi \to \chi | R) = \mathit{\Theta}(n - n_{\mathrm{c}}) \times D_{\mathrm{t/r}}(R) \times \\ \min \left[1, \exp\left(-\beta(U_\chi - U_\phi)\right) \times \frac{\prod_X^R q_{ij}(\chi \to \phi) \times \prod_Y^R p_{ij}(\chi \to \phi)}{\prod_X^R q_{ij}(\phi \to \chi) \times \prod_Y^R p_{ij}(\phi \to \chi)} \right] \end{aligned} \tag{2.53}$$

Early rejection is possible whenever a proposed cluster is greater than the cutoff, n_{c}, thus saving calculations for certain moves.

The algorithm is very effective at avoiding kinetic traps that single particle moves often fall into, and we have found settings that ensure approximate diffusive dynamics of a wide range of sizes of clusters. We have a probability of a rotation move being attempted set to 0.4, the maximum size for a translation set to $0.3\,\sigma^{\mathrm{LJ}}$, and the maximum size of a quaternion that is added to induce a rotation set to 0.3.

2.2.4 Umbrella sampling

Umbrella sampling is the standard methodology used to glean information about the underlying thermodynamics of a system [144, 174]. Every Monte Carlo step's acceptance probability is *biased* by a factor which is a function of an observable of the system. This *order parameter* may be, for example, the size of the largest liquid cluster [175], the size of a growing cluster [112] or the number of triangles in all the clusters.

The technique is used to find the probabilities of being in each state in a way which overcomes the limitations of certain states being unlikely to be visited. The mathematics is long, but does not involve any complicated concepts [176], and is explained below.

Mathematical derivation

The probability of measuring the value of any observable, $\mathit{\Omega}$, from all the possible configurations of a system, $\mathit{\Gamma}$, can be written as

$$P(\mathit{\Omega}) = \int P(\mathit{\Gamma}, \mathit{\Omega}) \, \mathrm{d}\mathit{\Gamma} \tag{2.54}$$

to give the canonical partition function.

The expectation of an observable can therefore be written as:

$$\begin{aligned} \langle \mathit{\Omega} \rangle &= \int \mathit{\Omega} P(\mathit{\Omega}) \, \mathrm{d}\mathit{\Omega} && (2.55) \\ &= \iint \mathit{\Omega} P(\mathit{\Gamma}, \mathit{\Omega}) \, \mathrm{d}\mathit{\Gamma} \, \mathrm{d}\mathit{\Omega} && (2.56) \end{aligned}$$

In a standard canonical Monte Carlo simulation, the probability of a given state is simply the Boltzmann probability, so we can rewrite the probability using the Kronecker delta function acting on the value of the observable of a given state ($\omega(\mathit{\Gamma})$) and our desired observable ($\mathit{\Omega}$).

$$\begin{aligned} P(\mathit{\Gamma}, \mathit{\Omega}) &= P(\mathit{\Gamma}) \delta_{\mathit{\Omega}, \omega(\mathit{\Gamma})} && (2.57) \\ P(\mathit{\Gamma}) &= \frac{\exp(-\beta U(\mathit{\Gamma}))}{\int \exp(-\beta U(S)) \, \mathrm{d}S} && (2.58) \\ \Rightarrow P(\mathit{\Gamma}, \mathit{\Omega}) &= \delta_{\mathit{\Omega}, \omega(\mathit{\Gamma})} \frac{\exp(-\beta U(\mathit{\Gamma}))}{\int \exp(-\beta U(S)) \, \mathrm{d}S} && (2.59) \end{aligned}$$

Thus our overall probability of a given value of an observable can be written using the Dirac delta function as follows:

$$\begin{aligned} P(\mathit{\Omega}) &= \int P(\mathit{\Gamma}) \delta(\mathit{\Omega} - \omega(\mathit{\Gamma})) \, \mathrm{d}\mathit{\Gamma} && (2.60) \\ &= \frac{\int \exp(-\beta U(\mathit{\Gamma})) \delta(\mathit{\Omega} - \omega(\mathit{\Gamma})) \, \mathrm{d}\mathit{\Gamma}}{\int \exp(-\beta U(S)) \, \mathrm{d}S} && (2.61) \end{aligned}$$

Our expectation value is therefore:

$$\begin{aligned} \langle \mathit{\Omega} \rangle &= \int \mathit{\Omega} \int P(\mathit{\Gamma}) \delta(\mathit{\Omega} - \omega(\mathit{\Gamma})) \, \mathrm{d}\mathit{\Gamma} \, \mathrm{d}\mathit{\Omega} && (2.62) \\ &= \int \mathit{\Omega} \frac{\int \exp(-\beta U(\mathit{\Gamma})) \delta(\mathit{\Omega} - \omega(\mathit{\Gamma})) \, \mathrm{d}\mathit{\Gamma}}{\int \exp(-\beta U(S)) \, \mathrm{d}S} \, \mathrm{d}\mathit{\Omega} && (2.63) \end{aligned}$$

In umbrella sampling, we are aiming to find out how we have to bias the states which have a certain value of an observable (the order parameter) in order to ensure that the probability of visiting all possible values of the observable within a certain range is equal.

Thence we can find the underlying probability of the states as a function of the order parameter and the expectation value of the observable.

In order to get uniform sampling over our order parameter, our aim is to bias our underlying probability distribution by a factor (l) as follows:

$$l_{\text{bias, ideal}}(\Omega) = \frac{1}{P(\Omega)} \tag{2.64}$$

This would give us new, non-Boltzmann, probabilities of certain states, as follows:

$$p_{\text{US}}(\Gamma, \Omega) = \frac{l_{\text{bias,ideal}}(\Omega) \times \exp(-\beta U(\Gamma)) \times \delta_{\Omega,\omega(\Gamma)}}{\int l_{\text{bias,ideal}}(\Omega(S)) \times \exp(-\beta U(S))\,\mathrm{d}S} \tag{2.65}$$

Unfortunately, we do not know the probabilities *a priori*, and often cannot find them from direct simulation. This is especially relevant when the values of $P(\Omega)$ are very low, as we may not visit any states with certain values of the order parameter without some initial biasing. We need to estimate our parameter and update it until we find a good estimate of the ideal biasing factor.

In our program, we bias the distribution by a parameter $\exp(-\beta x(\Omega))$. There are many other possibilities for this factor; this choice just means we are able to combine the energy and biasing factor when calculating acceptance. We change the probability of a given state:

$$p_{\text{US}}(\Gamma, \Omega) = \frac{\exp(-\beta x(\Omega)) \times \exp(-\beta U(\Gamma)) \times \delta_{\Omega,\omega(\Gamma)}}{\int \exp(-\beta x(\omega(S))) \times \exp(-\beta U(S))\,\mathrm{d}S} \tag{2.66}$$

Our move acceptance criterion thus becomes:

$$\gamma(\phi \to \chi) = \min\left[1,\, \mathrm{e}^{-\beta(x(\Omega_\chi)-x(\Omega_\phi))} \times \mathrm{e}^{-\beta(U_\chi - U\phi)}\right] \tag{2.67}$$

$$\gamma(\phi \to \chi) = \min\left[1,\, \mathrm{e}^{-\beta(U_\chi - U\phi + x(\Omega_\chi)-x(\Omega_\phi))}\right] \tag{2.68}$$

After each run, we can estimate the current probability of obtaining certain values of our order parameter ($p_{\text{obtained}}(\Omega)$), and thus change the biasing function so that our next run gives a smoother sampling. Generally, we are only interested in the relative values of the order parameter; therefore, we can use the number of times we visit each state ($C(\Omega) = p_{\text{obtained}}(\Omega) \times N_{\text{steps}}$) rather than the actual probability. For this derivation, $l_{\text{bias}}(\Omega) = \exp(-\beta x(\Omega))$. The relevant concept is that the more we visit a given state relative to the other states, the smaller the biasing factor should become.

$$l_{\text{bias, new}}(\mathit{\Omega}) = \frac{l_{\text{bias, old}}(\mathit{\Omega})}{C(\mathit{\Omega})} \tag{2.69}$$

$$\log(l_{\text{bias, new}}(\mathit{\Omega})) = \log(l_{\text{bias, old}}(\mathit{\Omega})) - \log(C(\mathit{\Omega})) \tag{2.70}$$

$$\frac{\log(l_{\text{bias, new}}(\mathit{\Omega}))}{\beta} = \frac{\log(l_{\text{bias, old}}(\mathit{\Omega}))}{\beta} - \frac{\log(C(\mathit{\Omega}))}{\beta} \tag{2.71}$$

$$-\frac{\log(l_{\text{bias, new}}(\mathit{\Omega}))}{\beta} = -\frac{\log(l_{\text{bias, old}}(\mathit{\Omega}))}{\beta} + \frac{\log(C(\mathit{\Omega}))}{\beta} \tag{2.72}$$

Our biasing factor update is therefore:

$$x_{\text{new}}(\mathit{\Omega}) = x_{\text{old}}(\mathit{\Omega}) + \frac{\log(C(\mathit{\Omega}))}{\beta} \tag{2.73}$$

Generally, we limit the maximum value of the change of the biasing factor at each step to ensure that we do not suffer from hysteresis (and so that the change in bias on unvisited states is not huge). We also shift the biasing factors so that we do not suffer from any numerical errors by ending up with incredibly small factors. We then run another simulation using our new factors and repeat until we find that we have an equal probability of visiting each value of the order parameter we are interested in. At this point, we have the correct biasing probability and can thus calculate the underlying probability of a given value of order parameter simply by remembering that:

$$l_{\text{bias, ideal}}(\mathit{\Omega}) = \frac{1}{P(\mathit{\Omega})} \tag{2.74}$$

$$l_{\text{bias, ideal}}(\mathit{\Omega}) = \exp(-\beta x_{\text{ideal}}(\mathit{\Omega})) \tag{2.75}$$

$$\Rightarrow P(\mathit{\Omega}) = \exp(\beta x_{\text{ideal}}(\mathit{\Omega})) \tag{2.76}$$

It is important in our umbrella sampling that we get not only even sampling, but also fast exchange between states. This is demonstrated in figures 2.3. Once we get fast exchange between states, we are able to update our biasing factor without needing to limit the change in the biasing factor to a maximum.

This is similar to the Wang-Landau algorithm (section 2.2.5), but, rather than updating the bias at each step, we use the number of visits after each run to estimate the bias needed, then repeat the simulation with new biasing probabilities until we get fast regular exchange between the states of interest.

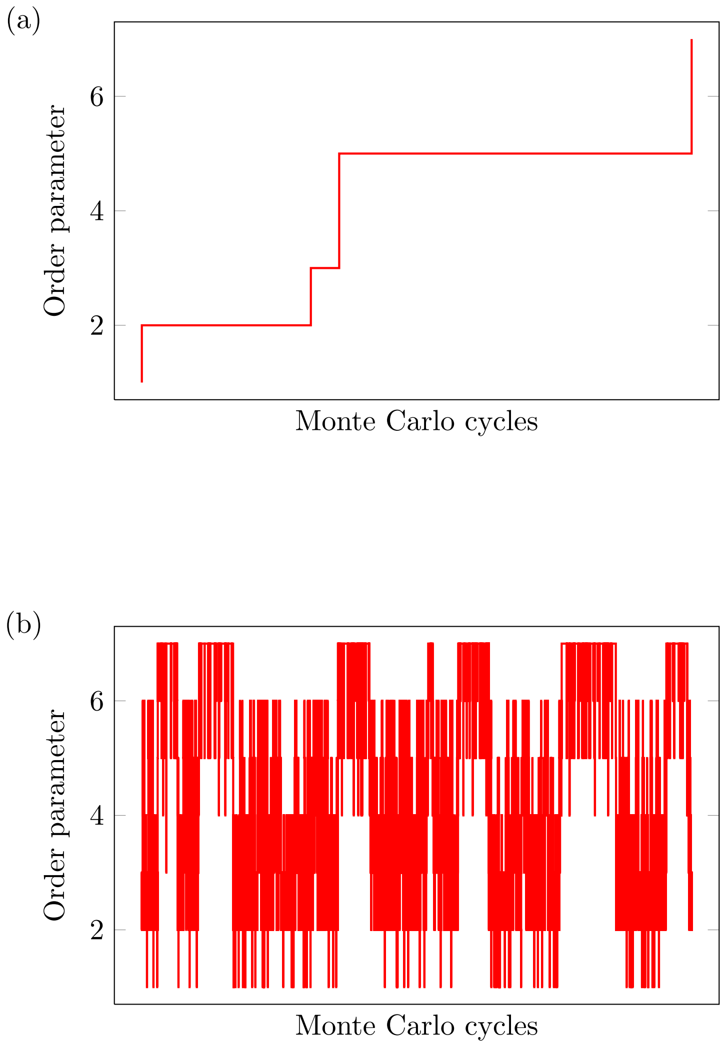

Figure 2.3: Changes of the order parameter as a function of Monte Carlo time in a regime of (a) inadequate exchange and (b) adequate exchange.

2.2.5 The Wang-Landau algorithm

Thermodynamic data have been collected using the Wang-Landau algorithm [148]. The Wang-Landau method uses a histogram approach to estimating the density of states of a given system using a probability-biased random walk. This estimate of the density of states can then be used to calculate simple thermodynamic properties, such as the heat capacity (see equation 2.78). We initially look at the algorithm in its simplest form, the canonical ensemble, before expanding to the more complicated mathematics required for the grand canonical ensemble. It is worth noting that the Wang-Landau algorithm does not follow detailed balance, though in the limit of a long run (and thus small convergence factor (f value)), it can be thought of as approximately doing so.

The stages of the Wang-Landau algorithm are as follows:

1. Set up a visit histogram as a function of energy, $H(U)$, in the range of interest, and an initial estimate of the density of states as a function of energy, $g(U)$ (in the original paper, the initial estimate is always flat, but it has been found in this system that giving an approximate initial profile greatly speeds convergence). Then select an initial convergence factor f. We used $\log(f) = 1$. Within our program, logarithmic units are used for the density of states and the convergence factor, $G(U) = \log(g(U))$, $F = \log(f)$, and F is initialised to 1 ($f = \mathrm{e}$).

2. Take a Monte Carlo type random walk with the acceptance probability γ, according to equation 2.77 for steps between state ϕ and χ.

$$\gamma(\phi \to \chi) = \min\left[1, \frac{g_\phi}{g_\chi}\right] \tag{2.77}$$

3. If a move is accepted, then increment the visit histogram $H(U_\chi) = H(U_\chi) + 1$, and multiply the density of states estimate $g(U_\chi)$ by the convergence factor f (or, equivalently, set $G(U_\chi) = G(U_\chi) + F$). Otherwise, increment the original state's visit histogram and density of states estimate similarly ($H(U_\phi) = H(U_\phi) + 1$ and $G(U_\phi) = G(U_\phi) + F$).

4. Steps 2 and 3 are repeated until the visit histogram H is 'flat' (there are several definitions of flat possible, see below). The visit histogram is then reset and the value of f decreased as $f = \sqrt{f}$.

5. Steps 2 to 4 are repeated until f decreases to below a threshold close to 1. Some recent work [177] suggested that the value of f should decrease as a simple inverse function of time if it ever fell below the value of $\frac{1}{t}$, where t is the number of Monte Carlo cycles that have occurred. This was attempted for this system, with only a minor improvement in efficiency in a few cases (most simulations did not fall below the trigger criterion of $f < \frac{1}{t}$).

6. The heat capacity can then be calculated using equation 2.78 or 2.79. Experimentally, equation 2.79 is less affected by noisy results, but it takes approximately three times longer to run than post-processing based on equation 2.78. Generally, equation 2.79 is used for final results, but when checking the progress of a run, equation 2.78 is

used.

$$\begin{aligned} Q(T) &= \int g(U) \times \mathrm{e}^{(-\beta U)} \mathrm{d}U \\ U(T) - U(0) &= -\frac{\partial \log(Q(T))}{\partial \beta} \\ C_V(T) &= \left.\frac{\partial U(T)}{\partial T}\right|_V \end{aligned} \tag{2.78}$$

$$\begin{aligned} \langle U(T) \rangle &= \frac{\int U \times g(U) \times \mathrm{e}^{(-\beta U)} \mathrm{d}U}{Q} \\ \langle U^2(T) \rangle &= \frac{\int U^2 \times g(U) \times \mathrm{e}^{(-\beta U)} \mathrm{d}U}{Q} \\ C_V(T) &= \left.\frac{\partial \langle U \rangle}{\partial T}\right|_V \\ &= \frac{\langle U^2 \rangle - \langle U \rangle^2}{kT^2} \end{aligned} \tag{2.79}$$

The actual calculations assume that the integral and sum are approximately equal (equation 2.80). This is felt to be a valid approximation for a large number of bins and at high temperature (see appendix B.5.2).

$$\int g(U) \times \mathrm{e}^{-\beta U} \mathrm{d}U \approx \sum_i g(U_i) \times (U_{i+1} - U_i) \times \mathrm{e}^{-\beta U_i} \tag{2.80}$$

The Wang-Landau method was further modified with the use of an automated bin adaptation algorithm, where the width of each bin was not fixed for the run, but was rather allowed to adapt to ensure more sampling in areas which carried more information about the density of states. We allowed the bins to change so that each bin contained an equal density of information about the system (the bins at energies where there is a higher density of states were smaller). This gave results that contained less noise, but, unfortunately, also decreased the speed with which a run could converge on the density of states (by orders of magnitude).

When defining whether the visit histogram was flat, we tried several different definitions. The first was that the number of visits in each bin should be within a certain proportion of the average visits multiplied by the current F-value. This ensured that the criterion became stricter later in the run. The next possibility was just to ensure that each bin was visited a certain minimum number of times. This gave a faster decrease of F-value, but the data appeared noisier. The final (and most effective) definition was to ensure that the number of visits in each bin were within a certain proportion of the mean number of visits (generally, 80 % was used), regardless of the current F-value.

It was found that the visit histogram count should be reset if it was not flat for a long time (approximately every 10^9 MC cycles). This occurred because of random poor sampling at the initial stages and if the reset did not occur, then it took a very long time for the visit histogram to become flat, and hysteresis was a risk. We later implemented a mode toggle that allowed only the most recent visits to be taken into account by keeping track of the visits at each update and having a rolling average. This gave better sampling under certain conditions, though, in order to ensure fast convergence, more effort had to be put into keeping track of over how long the averaging should occur for different systems.

We also looked at the effect of changing the way the f-value was updated. Normally, after the visit histogram is flat, the f-value becomes the square root of its previous value (F becomes half its previous size). However, it is possible to change the rate at which we decrease this. By setting a slower decrease of the f-value, the definition of a flat histogram could be made much looser, thus giving an decrease in the time taken for the program to converge. In general, it was found that decreasing the F-value by a factor of 0.5 was actually the most effective method to use in conjunction with the histogram flatness definition that just required all bin visit-densities to be within 80 % of the mean, but that with the F-value only being decreased by a factor of 0.9, a method that just required 1000 visits per bin could be used instead. The overall run time was similar for simple systems.

Proof of principle

We first tested our implementation of the Wang-Landau algorithm on a system of patchy particles able to form tetramers. We used this because theoretical heat capacity data had already been calculated for this system as part of another project [178], assuming the gases of clusters to be ideal and any vibrations to be harmonic (a similar set of approximations to those which we use in section 6.6).

The system consisted of 4 patchy particles, each with two orthogonal patches (A and B), where the A patches only attracted A patches and B patches similarly attracted only B patches (figure 2.4). The behaviour of this system was studied as a function of the relative strength of these interactions. Figure 2.5 shows that the theoretical and Wang-Landau algorithm results are clearly commensurate, thus validating the method within this model system. As the Wang-Landau algorithm was only run with 4 particles present, it also shows that it is not necessary to have a large system to find the thermodynamics. This ensures that we are only sampling a small energy space, and our chapter on encapsulation (chapter 3) uses single complete clusters and estimates from this the thermodynamic behaviour of larger systems. We review this further in appendix B.3.1.

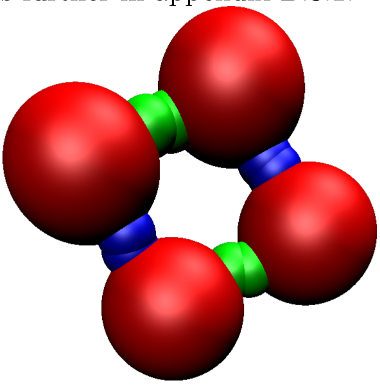

Figure 2.4: The cluster used in the test system (a square made of two dimers: we refer to this as an *AABB square* or dimer of dimers).

2.2.6 The Wang-Landau algorithm in the grand canonical ensemble

The Wang-Landau algorithm is actually applicable to many different ensembles. The original implementation used in the program was in the canonical ensemble, sampling over energy space. It is possible to sample over more than one variable, though this is slow for complicated systems. We decided to follow the example of Ganzenmüller and Camp [149], and use the Wang-Landau algorithm in the grand canonical ensemble, sampling over the number of particles. This attempts to find a biasing probability for the number of particles and continues until the probability of visiting each density is approximately equal. The biasing probability allows the relative values of the canonical partition function, $Q(N, V, T)$, to be found. These are related to the grand canonical partition function $Q(\mu, V, T)$ by

$$Q(\mu, V, T) = \sum_{N=0}^{\infty} Q(N, V, T) \times \exp(\beta\mu N) \tag{2.81}$$

and can be used to calculate the grand canonical probability distribution for the number of particles,

$$P(N, V, T) = \frac{Q(N, V, T) \times \exp(\beta\mu N)}{Q(\mu, V, T)}. \tag{2.82}$$

Unlike the Wang-Landau algorithm in the canonical ensemble, the biasing probability requires some derivation and explanation. For a given configuration of particles, $\mathit{\Gamma}$, if we wish to have a uniform sampling of the number of particles, N, then we can define our biasing probability, $p_{\text{bias}}(\mathit{\Gamma}, N)$, as follows.

$$p_{\text{bias}}(\mathit{\Gamma}, N) = \frac{P(\mathit{\Gamma}, N)}{P(N)} \tag{2.83}$$

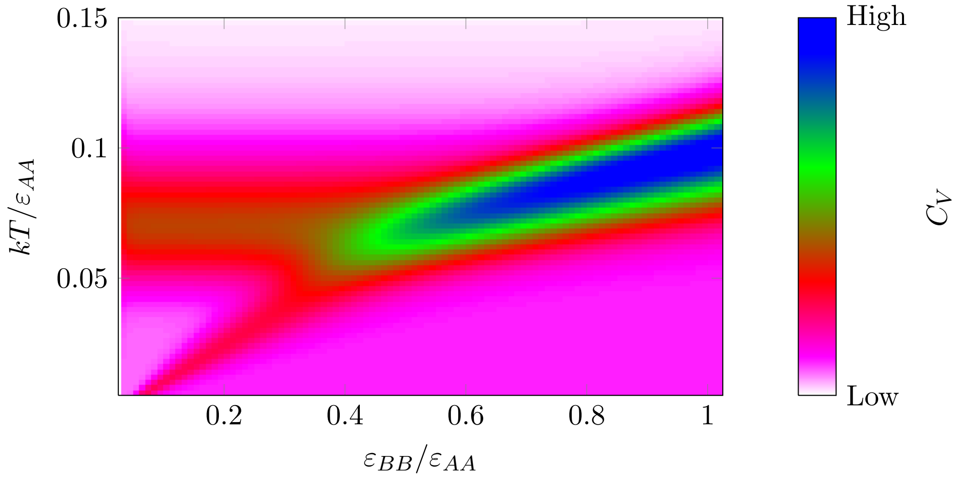

(a) Theoretical

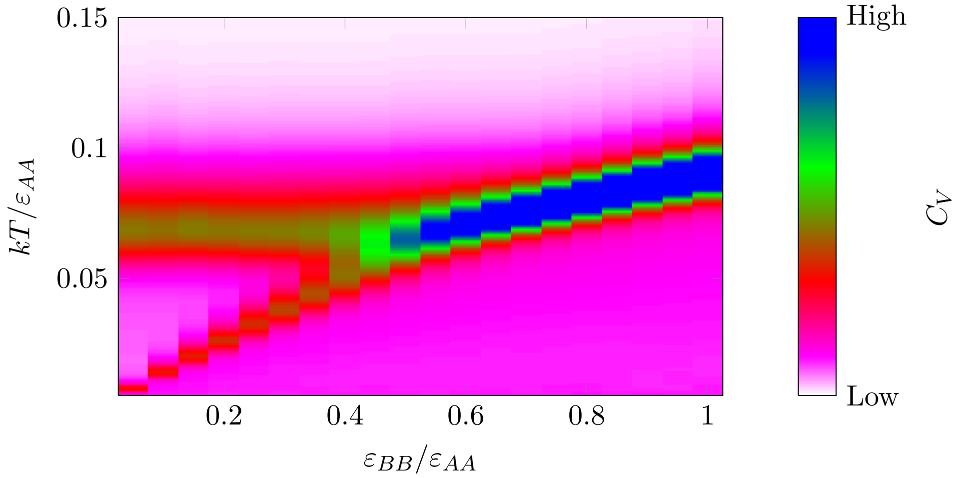

(b) Wang-Landau

Figure 2.5: Comparison of the theoretical (from [178]) and Wang-Landau derived heat capacities of the tetrameric system.

When we move between two states, ϕ and χ, without any bias for generating the moves, then our standard Monte Carlo detailed balance condition for our acceptance probability, γ, continues to hold. Namely:

$$\begin{aligned} p_{\text{bias}}(\mathit{\Gamma}_\phi, N_\phi)\gamma(\phi \to \chi) &= p_{\text{bias}}(\mathit{\Gamma}_\chi, N_\chi)\gamma(\chi \to \phi) \\ \implies \gamma(\phi \to \chi) &= \min\left[1, \frac{p_{\text{bias}}(\mathit{\Gamma}_\chi, N_\chi)}{p_{\text{bias}}(\mathit{\Gamma}_\phi, N_\phi)}\right] \end{aligned} \tag{2.84}$$

We know from statistical mechanics that the Boltzmann probability, $P(\mathit{\Gamma}, N)$, is:

$$P(\mathit{\Gamma}, N) = \frac{V^N \exp(\beta\mu N - \beta U(\mathit{\Gamma}))}{N!\mathit{\Lambda}^{3N} Q(\mu, V, T)} \tag{2.85}$$

The probability of a certain number of particles, $P(N)$, can be written as:

$$\begin{aligned} P(N) &= \int P(\mathit{\Gamma}, N)(d)\mathit{\Gamma} \\ &= \frac{Q(N,V,T) \times \exp(\beta\mu N)}{Q(\mu,V,T)} \end{aligned} \tag{2.86}$$

Thus we find from equation 2.83 that our biasing probability can be written as:

$$p_{\text{bias}}(\mathit{\Gamma}, N) = \frac{V^N \exp(-\beta U(\mathit{\Gamma}))}{(N!\mathit{\Lambda}^{3N} Q(N,V,T))} \tag{2.87}$$

Thus our acceptance probability is:

$$\gamma(\phi \to \chi) = \min\left[1, \frac{Q(N_\phi,V,T)}{Q(N_\chi,V,T)} \times \frac{V^{N_\chi} N_\phi! \mathit{\Lambda}^{3N_\phi}}{V^{N_\phi} N_\chi! \mathit{\Lambda}^{3N_\chi}} \times \frac{\exp(-\beta U(\mathit{\Gamma}_\chi))}{\exp(-\beta U(\mathit{\Gamma}_\phi))}\right] \tag{2.88}$$

We estimate $Q(N,V,T)$ during our simulation using a similar method to the canonical Wang-Landau method described in the previous section.

Sampling over the number of particles has significant advantages over continuous energy distribution, as sampling of a discrete variable is significantly easier to deal with than for a continuous variable and is also faster. Using a single simulation to find all data in large systems was slow; thus, we attempted to split the density landscape into smaller, overlapping sections: a technique known as *windowing*. This technique does not prevent ergodicity as it may do when windowing over energy variables. As the algorithm only allows us to estimate the relative partition functions ($Q(N,V,T)$), we need to scale our estimates when we combine the windowed system. This is strongly affected by any noise in the system, and any effects of the edge of the windowed range are thus amplified. For our clustering systems, this was not found to be significantly faster than our non-windowed runs overall (which had a maximum of approximately 500 particles), and required more computer core time overall (since many runs were needed), though it may be useful for very large systems.

Derivation

Once we have approximate values for the partition functions $Q(N,V,T)$, we are able to calculate the average number densities of the system at different chemical potentials. We use equation 2.82 for the probability and

$$\langle\rho\rangle = \frac{\langle N\rangle}{V} = \frac{1}{V}\sum_{N=0}^{\infty} P(N) \times N. \tag{2.89}$$

From the dependence of $P(N, V, T)$ and $\langle\rho\rangle$ on the chemical potential, it is possible to locate the coexistence points for liquid-vapour coexistence. At the chemical potential corresponding to coexistence, $\langle\rho\rangle$ will jump between that for liquid and vapour, and the centre of the transition can be defined as the point at which the two peaks in $P(N, V, T)$ corresponding to liquid and vapour states have equal area. In practice, this is achieved by a numerical search. We use a quadrature search technique, where the solution is found by converging on the correct answer iteratively. We first choose two starting limits for the chemical potential (one in the low density phase and one in the high density phase), then we test whether the mean density of the chemical potential between them is nearer the higher or lower density phase. We then repeat the search using the central chemical potential and either the original upper or lower limit (keeping the expected density of the two points in different phases), until the answer is found to the relevant degree of accuracy. Once the chemical potential at which there is an equal probability of the system being in either phase is located, we can find the density of the two coexisting phases.

The grand canonical Wang-Landau algorithm may experience some issues based on the size of the simulation box owing to clustering and the number of particles simulated. We explain the root of these and their effect in appendix B.3.

Proof of principle

The method was first tested on a Lennard-Jones system to ensure that it recovered the previously established phase diagram [179]. As can be seen from figure 2.6, our implementation of the algorithm gives results which agree with the previously established results from a similar length cutoff ($2.5\,\sigma$ was used in this case to allow comparison with the Gibbs ensemble data) when compared to both Gibbs ensemble simulation data ([180] as cited in [179]), and to the empirical modified Benedict-Webb-Rubin equation [179] (which is not expected to be accurate near the critical point for cutoffs less that $3\,\sigma$). For simple systems, the grand canonical Wang-Landau method is efficient at computing thermodynamic data, and our results are in agreement with other calculations.

2.2.7 Gibbs ensemble method

An alternative method to compute liquid-vapour phase behaviour is the Gibbs ensemble technique, proposed by Panagiotopoulos [181]. The Gibbs ensemble method aims to simulate two coexisting phases that are in thermodynamic, but not physical contact, allowing thermodynamic information about both to be obtained simultaneously. This method removes the issues with forming an interface between two phases (which changes

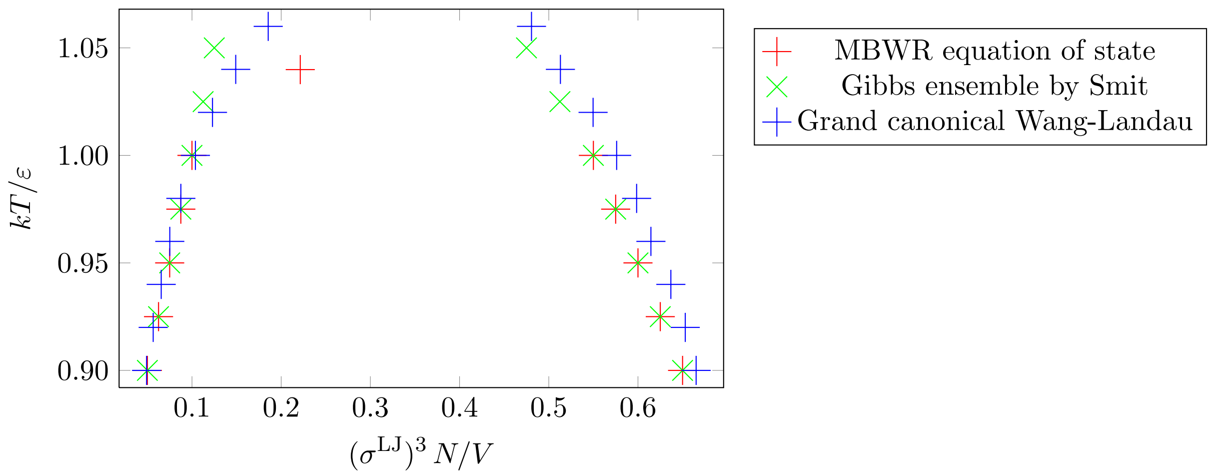

Figure 2.6: A plot of density against temperature for the Lennard-Jones system (MBWR equation from [179], Gibbs ensemble results from [180] as cited in [179]).

the free energy significantly for small simulations), by simulating two boxes which allow exchange of particles and volume between each other.

In Gibbs ensemble simulations, there are three classes of move possible. The first is the standard Monte Carlo displacement or rotation. The next is to remove a particle from one box and place it randomly in the other. The final move class is to change the volume of the two individual boxes, while conserving the overall volume of the system.

The displacement or rotation of a particle occurs using the same acceptance criterion as for a canonical ensemble normal move, and is explained in section 2.2.1. Virtual move Monte Carlo (section 2.2.3) can also be used if required. The displacement and rotation moves in each box occur independently of each other.

Particle exchange moves occur at the same time in both boxes. It can be thought of as two grand canonical Monte Carlo exchange moves occurring simultaneously (see section 2.2.2). First, a particle is selected for removal from one box (here, we use ϕ to signify the box which loses a particle), and the energy change when the particle is removed is calculated (ΔE_{ϕ}). The next part of the move is to add a particle to box χ and calculate the change in energy (ΔE_{χ}). The overall probability of the move being accepted is the product of the probabilities of the two components being accepted. Note that terms involving μ cancel, and thus we do not need to specify a chemical potential. From equation

2.37, we are able to derive the following probabilities:

$$P_{\text{add}} = \exp\left(-\beta \Delta U_\chi - \ln \frac{N_\chi + 1}{V_\chi}\right) \times \frac{\exp\left(+\beta\mu\right)}{\lambda^3} \tag{2.90}$$

$$P_{\text{rem}} = \exp\left(-\beta \Delta U_\phi - \ln \frac{V_\phi}{N_\phi}\right) \times \exp\left(-\beta\mu\right) \times \lambda^3 \tag{2.91}$$

$$\gamma_{\text{exchange}} = P_{\text{add}} \times P_{\text{rem}} \tag{2.92}$$

$$\gamma_{\text{exchange, accept}} = \min\left[1, \exp\left(-\beta\Delta U_\phi - \beta\Delta U_\chi - \ln \frac{V_\phi (N_\chi + 1)}{V_\chi N_\phi}\right)\right] \tag{2.93}$$

$$\tag{2.94}$$

Volume moves can be thought of in terms of the isothermal-isobaric ensemble (section 2.2.2). Here, we conserve the total volume of both systems. This is again a correlated move. We need to look at the combined probability of changing the volume of each system. Again, we have two systems: ϕ, which decreases in volume, and χ, which increases in volume, both by the same amount, ΔV. As in section 2.2.2, there are two possible ways of making a volume move, by performing a random walk in volume space or the natural logarithm of volume space. The former gives the following probabilities:

$$P_{\text{grow}} = \exp\left[-\beta\left(\Delta U_\chi + p\Delta V\right) + N_\chi \ln\left(\frac{V_\chi + \Delta V}{V_\chi}\right)\right] \tag{2.95}$$

$$P_{\text{shrink}} = \exp\left[-\beta\left(\Delta U_\phi - p\Delta V\right) + N_\phi \ln\left(\frac{V_\phi - \Delta V}{V_\phi}\right)\right] \tag{2.96}$$

$$P_{\text{vol move}} = P_{\text{grow}} \times P_{\text{shrink}} \tag{2.97}$$

$$\gamma_{\text{vol move}} = \min\left[1, \exp\left(-\beta\Delta U_\phi - \beta\Delta U_\chi + N_\chi \ln \frac{V_\chi + \Delta V}{V_\chi} + N_\phi \ln \frac{V_\phi - \Delta V}{V_\phi}\right)\right] \tag{2.98}$$

The random walk in the natural-logarithm-of-volume space, on the other hand, gives the following:

$$P_{\text{grow}} = \exp\left[-\beta\left(\Delta U_\chi + p\Delta V\right) + (N_\chi + 1) \ln\left(\frac{V_\chi + \Delta V}{V_\chi}\right)\right] \tag{2.99}$$

$$P_{\text{shrink}} = \exp\left[-\beta\left(\Delta U_\phi - p\Delta V\right) + (N_\phi + 1) \ln\left(\frac{V_\phi - \Delta V}{V_\phi}\right)\right] \tag{2.100}$$

$$P_{\text{vol move}} = P_{\text{grow}} \times P_{\text{shrink}} \tag{2.101}$$

$$\gamma_{\text{vol move}} = \min\left[1, \exp\left(-\beta\Delta U_\phi - \beta\Delta U_\chi + (N_\chi + 1) \ln \frac{V_\chi + \Delta V}{V_\chi} + (N_\phi + 1) \ln \frac{V_\phi - \Delta V}{V_\phi}\right)\right] \tag{2.102}$$

Note that in both move types, the pressure cancels, and so is not needed to be set for a move in concert.

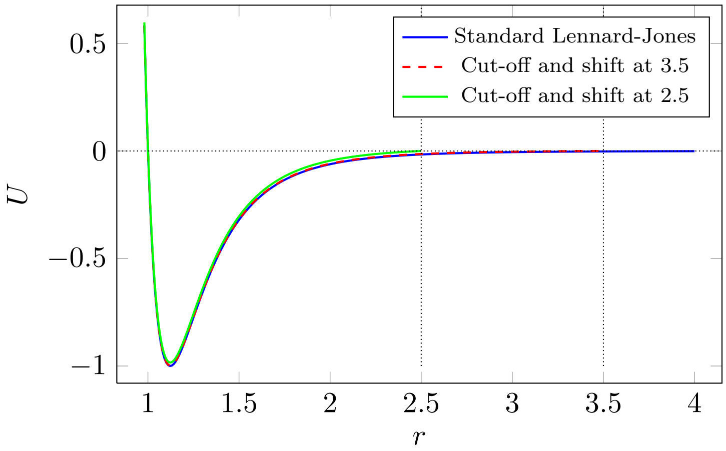

Figure 2.7: Gibbs cutoff compared to the normal cutoff used in the group.

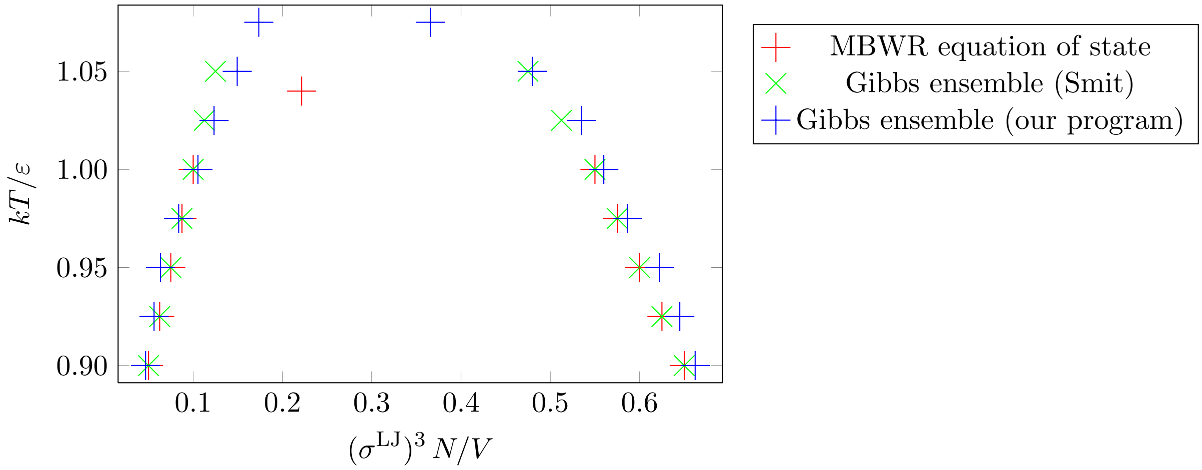

Figure 2.8: Proof of principle on the Lennard-Jones system (MBWR equation from [179], Gibbs ensemble results from [180] and our program).

Proof of principle

In order to ensure that the Gibbs ensemble implementation was working correctly, we tested it again against the Lennard-Jones system. As the smallest box size which can be simulated is one which has a side with length $2 \times r_{\text{cutoff}}$, it was necessary to use a smaller cutoff than usual, but balanced with keeping the shape as close to that of our larger cutoff as possible. We used a cutoff of $2.5\,\sigma$, which gives a very similar shape of potential to our normal cutoff of $3.5\,\sigma$ (as demonstrated in figure 2.7). As figure 2.8 shows, the implementation agrees closely with other Gibbs ensemble runs and the modified Benedict-Webb-Rubin equation.

Chapter 3

Encapsulation

3.1 Introduction

This chapter looks at the use of a central template particle to encourage the formation of capsid-like entities. Previous research [112] in the group has looked at the use of *torsional constraints* to encourage the formation of certain target structures. While this constraint is useful to prevent kinetic trapping into aggregates and is common to many proteins, we are more interested in the development of rules that could be used for non-torsional nano-technological applications. By using a central, encapsulated, template particle, we were able to form a high yield of clusters even in the absence of torsional constraints.

Other work from the group [46] has also touched on the use of structured shells of particles building up in layers to form a final structure. This templating was shown to be effective under certain conditions, but requires high precision control for the design of particles, so is not looked at in detail here.

First, we explore the biological motivation behind the idea of templating by encapsulation, then apply some of the basic concepts using our model to explore the formation of a wide variety of clusters. We look into the formation of a system which is possible without the use of an encapsulating particle or torsional constraints in the potential, the *icosahedron*, before moving onto at a system that has been shown to be difficult to form in the absence of torsional constraints, the *dodecahedron*. Finally, we take a look at the effect of changing from a simple, single, wide-patched particle to a non-isotropic narrower patch and then a central particle with multiple patches.

This chapter is based on a recently published paper [6].

3.1.1 Terminology

Throughout this chapter, we refer to the particles making up the external surface of the cluster as *shell* particles and the encapsulated/template internal particle as the *core*. Patches which bind the shell particles together are termed A-patches, those on the shell particles which bind to the core particle B, and the core particle's patch(es) are referred to as C. The Lennard-Jones well depth for interactions between shell particles is referred to as ε_{AA}, and that between shell and core particles as ε_{BC}. Interactions are also described in terms of the particle type on which the patches are and the particle type to which they bind; for example, shell-shell interactions refer to A-type patches (they are on the shell particle and bind to other shell particles), shell-core to B-patches, and core-shell to C-type patches. The Lennard-Jones diameter of the shell particles is known as σ_A^{LJ}, and that of the core particle as σ_B^{LJ}.

In all pictures from simulations, the colour scheme for the systems of interest is for the shell particles to be red with blue being used to represent the location of the patches which are able to bind together to create the shell structure (A-type patches), while green is used to show the location of patches of any interaction site on the shell which binds to the core (B-type patches). The core particles are shown in yellow. Note that the patches do not have a volume and are just schematically represented as such to show the binding sites (as demonstrated in figure 3.1).

When looking at the yields of different types of cluster, we used an adaptation of a diagram designed by Villar [178]. These show the kinetic yields of different clusters on a two-dimensional grid in different colours. The target cluster is coloured red, the equivalent cluster without a central particle green, and aggregates as blue. An aggregate is defined as any cluster having 25 or more particles in the icosahedral system, and any with 41 or more particles in the dodecahedral system (that is, greater than twice the number of particles in the target cluster's shell).

The terms 'centred' and 'encapsulated' are used interchangeably when describing clusters with a core particle, and the terms 'uncentred' and 'empty' are used to describe those without a core particle. The core particle is also referred to as the central particle in this chapter.

3.2 Biological inspiration

The motivation behind encapsulation came from biological systems, mainly from viruses.

We looked at viruses in detail in the introduction. Viruses come in many forms (see section 1.3.2), but the inspiration for this section is mainly from the assembly mechanisms of certain icosahedral viruses.

While viruses are a source of inspiration for us when looking at encapsulation, the majority of this chapter is concerned with the formation of clusters without torsional constraints, and this differs from the formation of viral capsids (which have a torsional dependence when binding due to the non-rotationally symmetrical positioning / shape of their sites of interaction).

3.3 Relevant research

3.3.1 Experimental systems

Experiments have shown that virus capsid proteins are not only able to encapsulate their genome [75], but conditions can be found where they assemble around non-biological objects, such as nano-particles [182–186]. In such assemblies, the nano-particles mimic the electrostatic behaviour of the nucleic acid.

As mentioned in chapter 1, some very interesting experimental work has been carried out by Burnley and Cox [109]. This looked at the use of both a flexible template (thought to be an analogue to the viral genome) and a globular template (a structureless central particle). Their findings were that a flexible template gave reasonable yields of correctly formed target cluster, while a globular template did not. This should be borne in mind when we examine our results, which clearly show that a structureless globular template is not only able to encourage assembly, but that, by being structureless and allowing particles to diffuse across its surface, it permits efficient assembly, especially in the case of the dodecahedron. We explore the rôle of this surface diffusion when comparing the assembly of an isotropic central particle with one with patches and structure (section 3.8.2).

The use of dissolvable templates has been shown by Caruso's group to be an effective way of controlling the size, composition, geometry and thickness of the formation of self-assembling shells [187]. Their method uses a colloidal template particle onto which species can be adsorbed using layer-by-layer deposition of charged species [188]. By choosing the material of the components carefully, the central particle can then be dissolved [189] or thermally decomposed [190], leaving a porous shell of the desired size. Though we are finding the limits of and useful methods to encourage encapsulated self-assembly, researchers looking into the application of these rules should consider some of the chemistry carried out by Caruso *et al.*

3.3.2 Simulation and theoretical research

Although less studied, simulations [127–130] and theory [140–142] have begun to address the assembly of filled viral capsids. Particularly relevant to the work in this chapter are the simulations of Hagan and co-workers, who have shown that the change in assembly mechanism from homogeneous nucleation of the capsid proteins in the empty capsid case to heterogeneous nucleation around a core enhances the ability of the capsids to assemble [127]. These simulations have modelled the genomic material as a spherical entity that has an attractive interaction with the capsid proteins. They neglect internal structure and degrees of freedom of all particles, in particular the core, which is polymeric *in vivo* in viruses. Some research has suggested that the polymeric nature of the genome may play an important rôle in the self-assembly of viruses. For example, diffusion along an RNA strand by capsid monomers has been shown to be relevant to the assembly mechanism by use of theory [141], as it gives a fast one-dimensional diffusion path over which capsid monomers can be brought together quickly. However, as mentioned previously, some viruses are able to assemble around a nano-particle core, so the exact nature of the template is not relevant [182–186]. This chapter concentrates on globular, patched central particles, rather than the use of polymers for encapsulation, as we are interested in nano-technological applications rather than modelling viruses.

3.4 Nano-technological applications

As previously mentioned, first-generation patchy colloids and nano-particles are unlikely to have the torsional component in the potential which is common to protein-protein interactions. Simulations have indicated that the one-component self-assembly of simple monodisperse targets is still feasible [111, 112, 191]; however, there can be significant differences in the mechanisms of assembly and the nature of the kinetic traps compared to torsionally constrained particles. In particular, the lack of a torsional component in the interparticle potential to enforce convexity in the growing clusters leads to disordered aggregation competing with correct assembly [111, 112]. Furthermore, as the size of the target structure increases, the difficulty of assembly increases much more rapidly when torsional constraints are not present. For example, the self-assembly of a simple 20-particle dodecahedron occurs readily with a protein-like potential with torsional constraints [192], but is impossible without [112].

We show that the use of encapsulating particles is a useful step in allowing the formation of clusters of certain shapes in the absence of torsional constraints. This work may be applied to certain pharmaceutical aims. The creation of structured clusters encapsulating

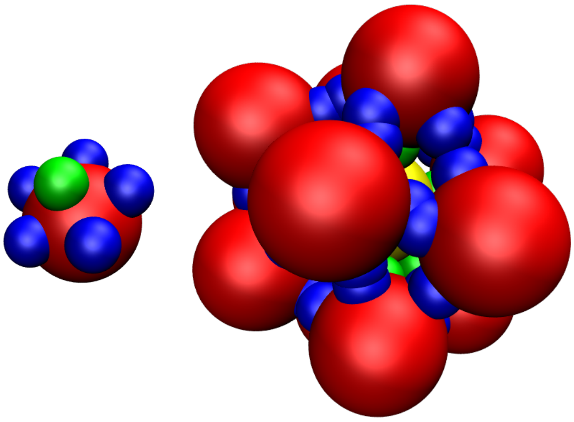

Figure 3.1: A shell particle and a completed cluster for the centred icosahedra forming system.

a drug for delivery to a host organism is a desirable aim of much research, as it allows specific targeting of cells (*e.g.* cancer tumour cells [193]) while preventing the drug from being inactivated by the body's enzymes [194]. The use of DNA clusters able to form nano-capsules has also recently been shown to be a possible way of encapsulating a drug compound [195]. The encapsulated cargo acting as the template for formation is a desirable aim in nano-technological assembly, as it would simplify the synthesis route of the final capsid.

This research also helps to understand some of the theoretical basis of the current experimental virus-based nano-technology, which includes use as delivery mechanisms [196], in immunoassays [88], and even as tools for finding out the connectivity of the central nervous system [89].

3.5 Icosahedra

Previous work in the group has already characterised the formation of empty icosahedral clusters [46, 112]. We begin by comparing the formation of empty icosahedral clusters to the formation of icosahedra around a central shell particle with a single isotropic attractive patch (which we refer to as IhC). When not varied, the encapsulated particle had a size which allowed efficient assembly and packing to occur (the internal particle has a diameter of 95 % of that of a shell particle). Icosahedral clusters are able to form in the absence of torsional bonding constraints and the initial investigation aimed to establish the differences in assembly mechanisms and stability of the two systems.

Variable	**Value**
$\frac{\varepsilon_{BC}}{\varepsilon_{AA}}$	1
σ_A^{PW}	0.35 radians
σ_B^{PW}	0.35 radians
$\frac{\sigma_B^{\mathrm{LJ}}}{\sigma_A^{\mathrm{LJ}}}$	0.95 (IhC) *or* 1.80 (DdC)
$\frac{N_{\mathrm{shell}}}{V}$	0.15 $(\sigma_A^{\mathrm{LJ}})^{-3}$
$\frac{N_{\mathrm{core}}}{N_{\mathrm{shell}}}$	0 (IhO and DdO) 1/12 (IhC) 1/20 (DdC)
$\lvert\boldsymbol{r}^{\mathrm{cutoff}}\rvert$	3.5 σ_A^{LJ} (dynamic) *or* 2 σ_A^{LJ} (Wang-Landau)
ζ (patch noise parameter)	0
$\frac{1}{\sigma_A^{\mathrm{Tors}}}$	0
$\frac{1}{\sigma_B^{\mathrm{Tors}}}$	0
Repeats for dynamic runs	5
Number of Monte Carlo cycles	10^6
Maximum number of complete clusters simulated	10 per run (dynamic) *or* 1 per run (Wang-Landau)

Table 3.1: Parameters used in the simulations of centred icosahedra and dodecahedra

3.5.1 Dynamical results

The investigation into the dynamical behaviour explored a large amount of parameter space. The regions where encapsulated icosahedra form efficiently have been delimited. We will look at the effect of the different variables in turn before examining the changes in assembly mechanism that were relevant to this system. All of the kinetic simulations were run for a million Monte Carlo cycles, and with the use of virtual move Monte Carlo steps [146]. The parameters used for the moves were those which previous research within the group had shown gave results the most similar to the diffusional behaviour of small colloid clusters in water [46] (see appendix D.1.1).

Each simulation contained 120 shell particles and 10 core particles, allowing a maximum yield of 10 centred icosahedra. When calculating the yield of clusters, we allowed a maximum of two defects in our final snapshot (*i.e.* two bonds that were not *formed*). A bond is defined as formed if the energy of attraction between two particles is at least 40 % of the well depth of perfectly aligned patches. In the cases where the central particle was not able to form bonds with all of the shell particles (due to size or patch width), we have also included the yield of clusters which have the correct size (13-mers in the case of icosahedra: final configurations from the program have shown that these were generally of the correct shape). All results are taken from the average of 5 runs, each with a different random seed (see chapter 2 for more details).

Unless otherwise stated, the values in table 3.1 were used for the simulations.

Strength of interactions between the shell and core particles (ε_{BC})

We looked at the effect of changing the strength of the interactions between the encapsulated particle and the external particles, ε_{BC}. This, unsurprisingly, has a strong effect on the balance of centred and uncentred icosahedral cluster yields. Figure 3.2 shows the effect of changing the strength of the interactions on the yield at a variety of temperatures.

At weak patch strengths, the formation of uncentred icosahedra is preferred due to the gain in free energy by releasing the central particle (entropically favoured while enthalpically disfavoured). This effect can be seen at patch strength ratios just above 0.2 where, as the temperature decreases, the predominant structure changes from empty to centred icosahedra. All other patch strengths have a reasonable window of formation of the centred target structure, with the window width increasing with increasing patch strength. The temperature at which the system becomes trapped into aggregates seems little affected by the change of the interaction strength. This suggests that encapsulation can be applied to other structures which are normally unable to be formed owing to aggregates acting as a

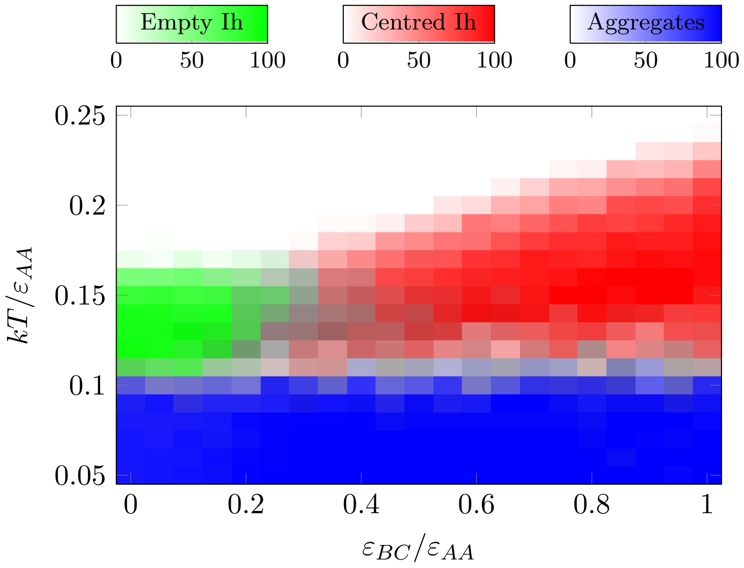

Figure 3.2: Plot of successful assembly of the centred icosahedron-forming system as a function of the interactions between shell and core particles (other settings are given in table 3.1).

kinetic trap.

Patch widths

As previously mentioned, there are three patch types that are relevant to this system: the patches which interact to form the shell, A, and those on the shell particle, B, which interact with the central particle's patch, C. As the central particle was initially configured to be anisotropically attractive, we looked at the effect of varying the shell-particle's A (shell-shell) and B (shell-core) patches. The patch widths of patches of types A and B were set at 0.35 unless they were being varied.

Shell-shell patch widths (σ_A^{PW}) Figure 3.3 shows the yield when varying the shell-shell patch width. Unlike the patch strength plots, there are many more limits to successful assembly visible here. Figure 3.4 labels each of the regions as explained below.

There are two kinetic traps and two regions of thermodynamic instability for our target cluster.

1. Our first kinetic trap is where large aggregates form. This occurs at low temperature in the middle ranges of patch width. It is also discouraged at very narrow patch widths. Here, the monomers become trapped in incorrectly formed clusters part way through assembly, and these clusters bond together: they are unable to become free and assemble correctly owing to the lack of thermal energy.

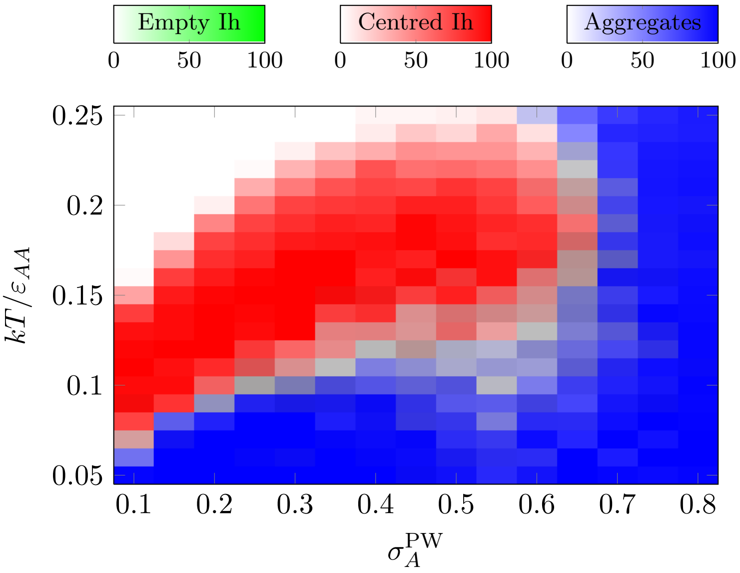

Figure 3.3: Plot of successful assembly of the centred icosahedron-forming system as a function of shell-shell patch width (other settings are given in table 3.1).

2. In the second kinetic trap, clusters part-assemble correctly, but all of the monomer ends up in the intermediates. The system becomes trapped in these part-assembled clusters which do not aggregate significantly (though some small aggregates are able to form). We call this arrested formation. This happens at low temperature to patches which are too narrow to form a liquid.

3. The first region of thermodynamic instability for our cluster occurs at very wide shell patch widths. Here, the system approximates Lennard-Jones spheres. Thus, we do not get clustering, but rather form liquid droplets. This behaviour is destabilised by decreasing the patch width. We look at this more in chapter 6.

4. The second region of thermodynamic instability of our cluster is when the formation of small aggregates (or monomers) is more favourable than the icosahedral clusters. This occurs when the free energy of the target clusters is higher than that of the small aggregates or monomers. This can be due to narrow patch widths, meaning that there is low entropy of the completed cluster, or weak bonds in the target cluster, at high temperatures.

The correct formation of clusters is thus a balancing act between the patch width needed to be sufficiently narrow to destabilise the liquid with respect to the clusters, but not so narrow as to destabilise the clusters with respect to monomers. The need to consider both the monomer and liquid phases' stability with respect to that of the cluster is an important theme when designing patchy particles.

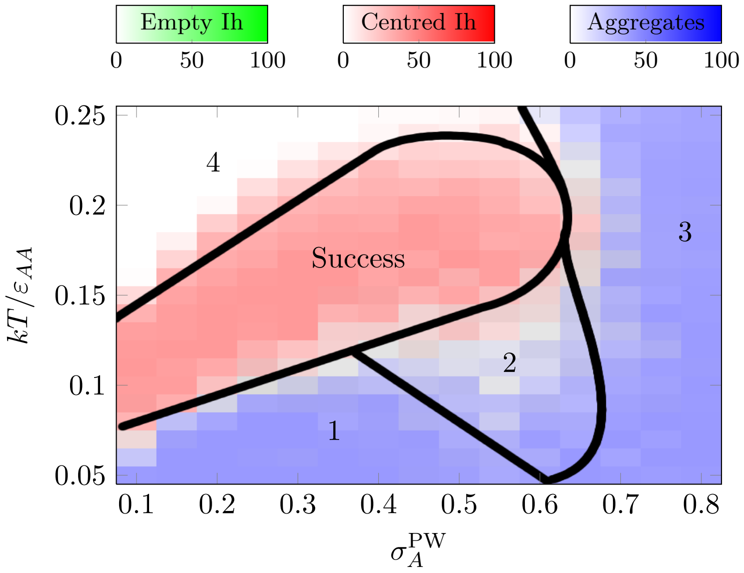

Figure 3.4: Schematic diagram to demonstrate the limits of assembly of the centred icosahedral system as a function of the shell-shell patch width. Labels 1-4 refer to regions of different types of aggregation, and are explained in the text.

The patch width of the shell-shell interactions is the first variable that we are able to compare to the non-templated icosahedral cluster forming system (known as *IhO*), the results for which are shown in figure 3.5. It is clear that the inclusion of a central particle greatly increases the range of conditions for which self-assembly is successful. Figure 3.6 shows the yield when we scale the temperature to account for the fact that the energy of the completed empty cluster is lower than that of the centred system. When compared with figure 3.3, it is clear that the centred system has an enhanced ability to self-assemble, especially at lower patch widths.

The effect of using a weaker interaction between shell and core particles In order to show the effect of competition between centred and uncentred icosahedral clusters, we decreased the well depth of the interaction between the shell and core particles (ε_{BC}) to half the strength of the interaction between shell particles (ε_{AA}). We then examined the effect on the yield plots as a function of shell-shell patch widths (σ_A^{PW}). The results can be seen in figure 3.7. It is clear that under these conditions, the kinetic stability of uncentred clusters becomes relevant at wide patch widths and lower temperature. The yield of uncentred clusters may be due to the assembly mechanism and is explained in section 3.5.1. Figure 3.8 shows the yields when we scale the temperature by the ratio of the energy of the original centred system and the energy of the system with weakened core-shell interactions. The results demonstrate a decrease in the parameter range at which we see successful assembly of the centred icosahedron with weakened core-shell interactions

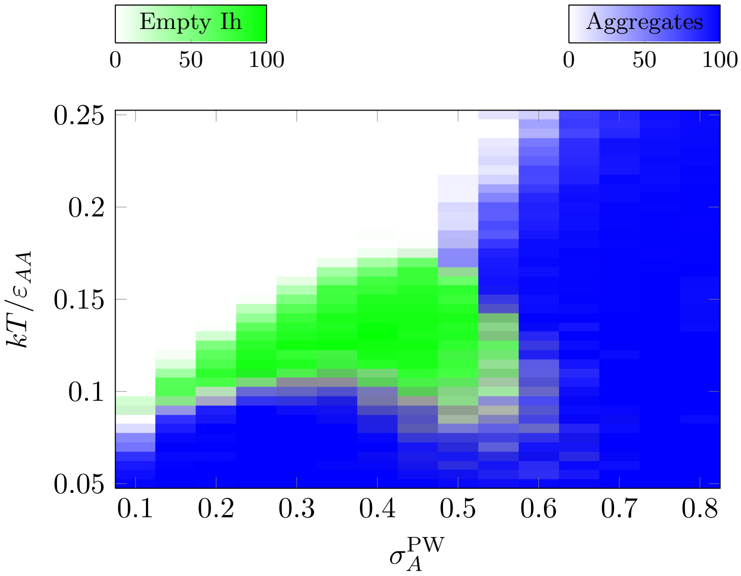

Figure 3.5: Plot of successful as a function of shell-shell patch width for the uncentred icosahedra forming system (other settings are given in table 3.1).

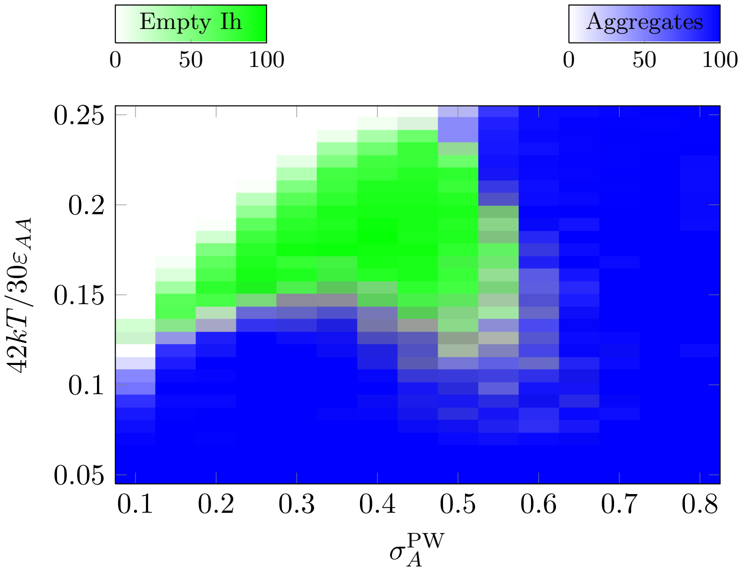

Figure 3.6: Plot of successful assembly for the uncentred icosahedra forming system as a function of shell-shell patch width with temperature scaled by the ratio of total energies of the IhO and IhC clusters, to allow better comparison with figure 3.3 (other settings are given in table 3.1).

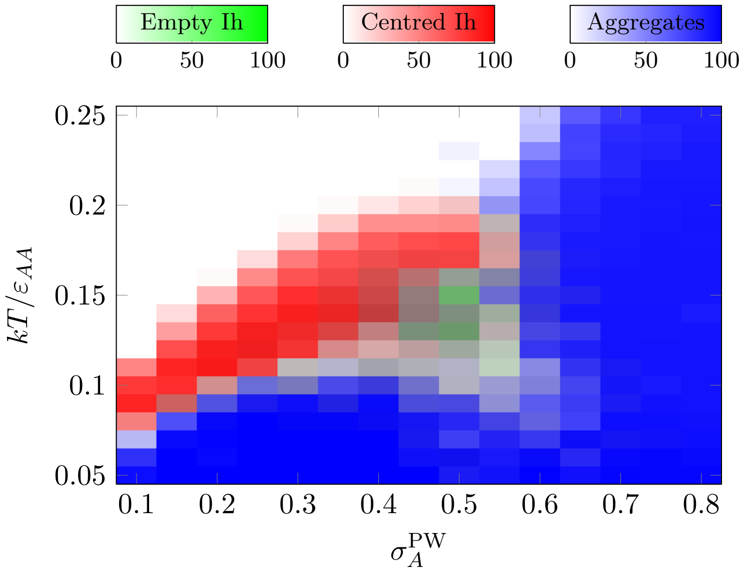

Figure 3.7: Plot of successful assembly as a function of shell-shell patch width with weakened shell-core interactions ($\frac{\varepsilon_{BC}}{\varepsilon_{AA}} = 0.5$; other settings are given in table 3.1).

compared to the original, IhC, system. The decrease in the region of successful assembly is partly due to trapping in uncentred icosahedra.

Shell-core patch widths (σ_B^{PW}) The patch width of the shell patches which bonded to the isotropic central particle was found not to limit the yield of icosahedra strongly (figure 3.9), except at certain extreme values. The yield does not vary strongly between patch widths of 0.2 and 0. At very wide patch widths, bonding with the central particle no longer requires shell particles to be facing the correct way, and so the interaction with the core no longer ensures that the shell grows with the correct curvature. In fact, it is possible for the core particles to bind to the shell-forming particles with the shell-forming-patches facing outwards rather than inwards, stabilising the formation of the liquid compared to the cluster and the monomer. This simultaneously decreases kinetic traps, as a part-formed cluster can rearrange more easily.

At narrow patch widths, the centred system has decreased entropy compared with wider patch widths, as less variation in conformation is permitted while being bonded. This effect destabilises the cluster with respect to the monomer. There is also less chance that when two particles come into contact they will bond, so narrow patch widths are both thermodynamically and kinetically limiting. Again, the theme of balance is important.

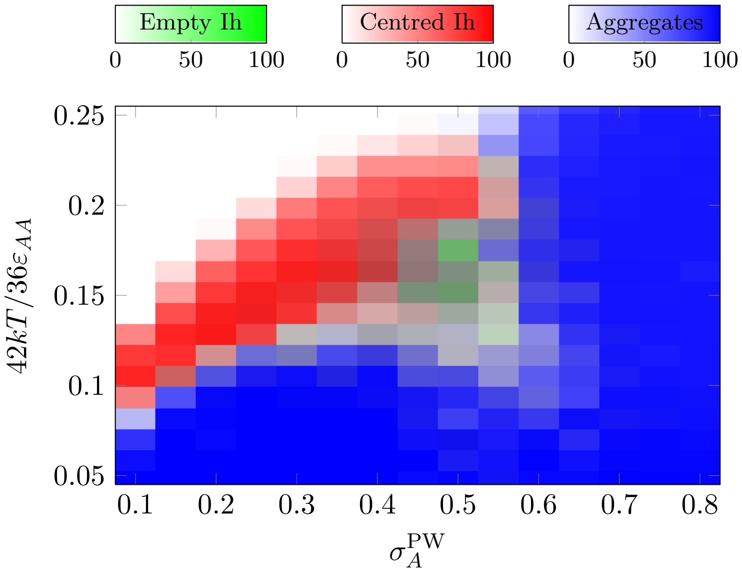

Figure 3.8: Plot of successful assembly as a function of shell-shell patch width with weakened shell-core interactions with temperature scaled to allow simple comparison with figures 3.6 and 3.3 ($\frac{\varepsilon_{BC}}{\varepsilon_{AA}} = 0.5$; other settings are given in table 3.1).

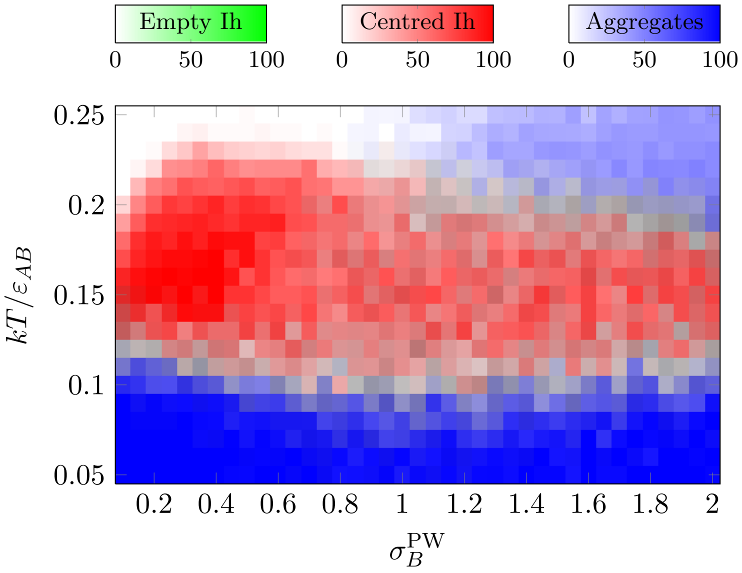

Figure 3.9: Plot of successful assembly as a function of shell-core patch widths (other settings are given in table 3.1).

Size of the core particle (σ_B^{LJ})

When defining whether a complete cluster has formed correctly, it is necessary to note that when the central particle is large compared to the shell particles, the formation of centred icosahedra is no longer feasible, as the shell particles are unable to form bonds (which we define as interactions with an energy of at least 40 % of the Lennard-Jones well depth) both with each other and the central particle (geometrically, this happens when the diameter of the core particle is greater than 140 % of the diameter of the shell particles). Similarly, when the central particle has a diameter too small (geometrically, the cutoff is at about 50 % of the diameter), then it may no longer bond to all of the shell particles, and instead may either bond to a set of particles close together on one edge or “rattle” about the cage, moving freely from the centre to the edge. We examine this rattling effect in section 3.5.2. For this reason, when examining the dependence of assembly yield on the size of the core particle we consider both our target structure (figure 3.10(a)), and clusters of size 13 (figure 3.10(b)). The latter 13-mers all appear to be completed clusters but with the incorrect number of shell-core bonds.

When the core particle is small, the system still forms an appreciable yield of icosahedral clusters. This yield is interesting, as it is not possible for the small core particles to bond to all of the shell particles. This result implies that the key process determining the product that forms is the initial nucleation of an icosahedral cluster on the core, and not the ability of all particles to bond to the core. In our kinetic yields, we also see the formation of some clusters with two core particles, though these are rare. The formation of empty icosahedra competes with centred icosahedra at very small core particle sizes. Although some small non-icosahedral particles form around the core, they are normally less stable than the empty icosahedra. This behaviour is different to the dodecahedral system, which we study later (see section 3.6.1).

At larger core particle sizes, we find that there is a sudden cutoff in the yield of the centred icosahedron (at a diameter of around 110 % of the shell diameter), but it is not replaced with uncentred icosahedral clusters. Here, centred clusters which have more than 12-shell particles can form. These clusters correspond to a variety of different structures stabilised by the strong core-shell interaction. As we use a 12 : 1 stoichiometric mixture of shell to core particles, sometimes these clusters are not closed owing to the system running out of shell-monomers, and often this results in there being unbound core particles. An example configuration found is shown in figure 3.11.

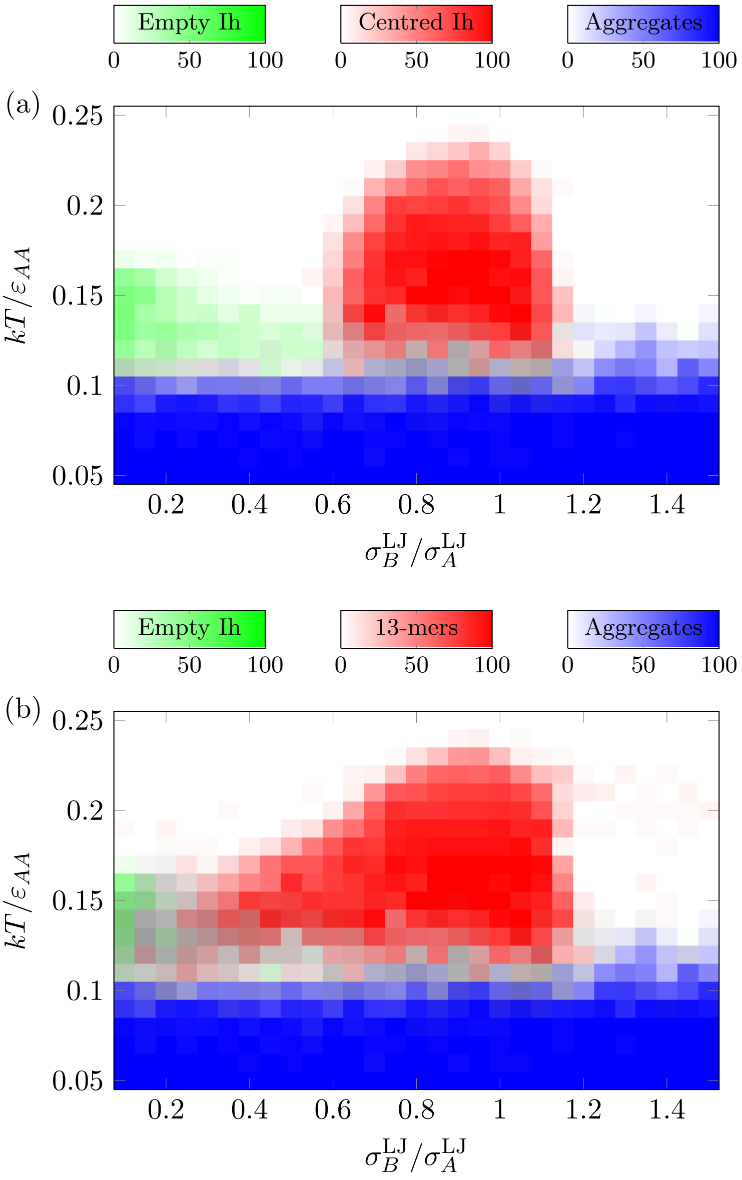

Figure 3.10: Plot of (a) yield clusters with the correct number of bonds (allowing up to two defects) and (b) yield of correctly sized clusters as a function of the size of the core particle for the centred icosahedra-forming system (other settings are given in table 3.1).

Density

The effect of changing the number density of particles was investigated at low densities for completeness (see figure 3.12). It is clear that changing the density has little effect on the total maximum yield, though it does broaden the temperature range that gives a reasonable yield, somewhat due to the reduction in the translational entropy of monomers. At even higher densities, there is the possibility that other thermodynamic phases, *e.g.* liquids, monomer crystals and cluster crystals, could become the stable phase (this possibility is investigated for some systems in chapter 6, but is not considered in detail here).

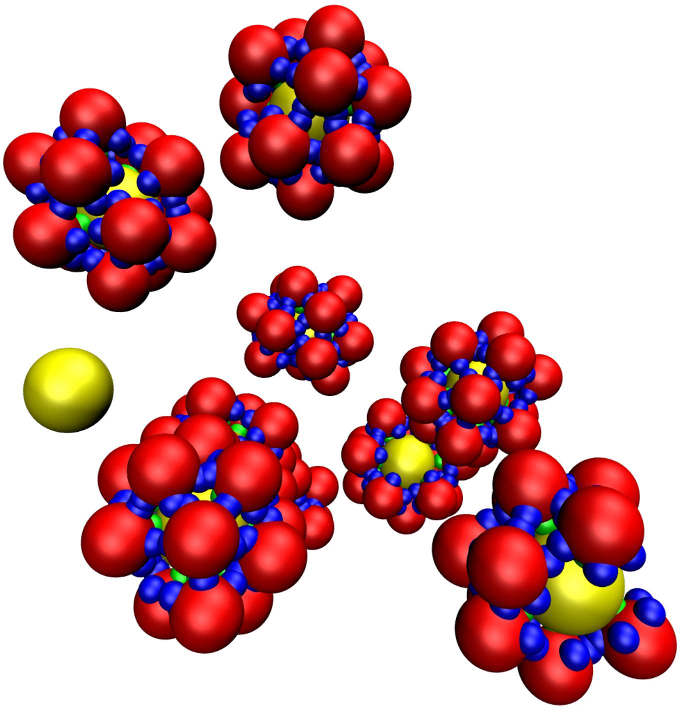

Figure 3.11: An example of the larger structures formed ($\sigma_B^{\mathrm{LJ}}/\sigma_A^{\mathrm{LJ}} = 1.3$, $kT/\varepsilon_{AA} = 0.15$; other settings are given in table 3.1). Note that the structure in the foreground is incomplete, and that there is an isolated free core particle.

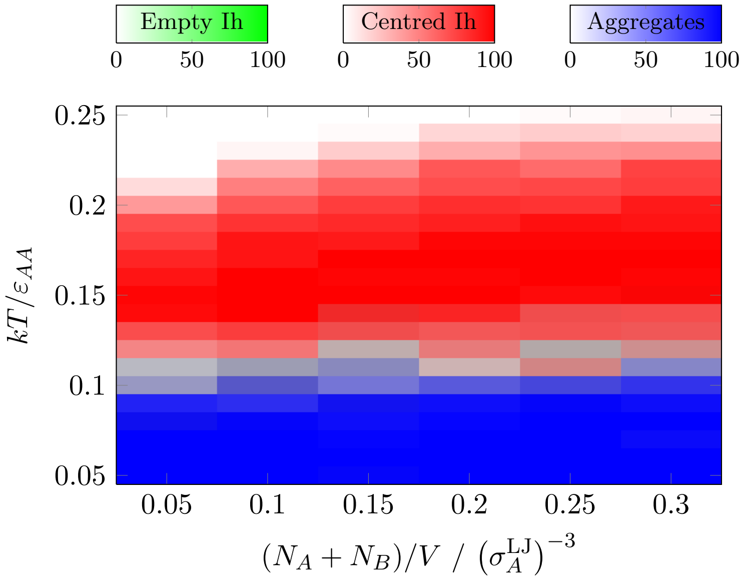

Figure 3.12: The effect of density on the assembly of the centred system (other settings are given in table 3.1).

Assembly mechanism

There is a noticeable change in the mechanisms of assembly for centred and uncentred icosahedra. Previous work on uncentred icosahedra has shown that there are two mechanisms of self-assembly: *budding* and *direct nucleation* [46]. In *budding*, a large cluster forms initially, and it then subsequently rearranges to release icosahedra.

The centred-templated system showed a preference for a direct nucleation-type mechanism (occurring around the core particle), while the uncentred system can also form under certain conditions by budding. The difference in assembly mechanisms can be seen in the time resolved yield plots (figure 3.13). Both were run under equivalent conditions (both have a shell-shell patch width of 0.35, the centred system is run at a temperature of $0.14\,\varepsilon_{AA}/k$, and the uncentred at the slightly lower temperature of $0.12\,\varepsilon_{AA}/k$, so that they both had similar yields). For the uncentred case, the system initially forms aggregates which then break down to form uncentred icosahedra. For the centred system, on the other hand, there is only a small transient population of aggregates. At wider patch widths and low temperatures, however, there was some evidence of an intermediate aggregate stage for the centred system. It is worthwhile noting that for most viruses, viral proteins are assembled by nucleation and the *budding* mechanism is unheard of[1] [197]. Thus the simple model is in agreement with experimental data within the region of high yield despite the lack of torsional constraints in our potential.

Robustness of assembly

As any experimental method to form patchy particles will involve a certain degree of "polydispersity" in the patch positions, we have also considered the effects of imperfectly positioning the patches. These were generated by adding a small random quaternion (Q_{random}, of maximum length ζ) to the quaternion representing the ideal position of each patch on each particle (Q_{orig}), followed by re-normalisation (see section 2.2.1 for details on quaternions):

$$Q_{\text{fin}} = \frac{Q_{\text{orig}} + Q_{\text{random}}}{|Q_{\text{orig}} + Q_{\text{random}}|} \tag{3.1}$$

The level of noise was varied by changing the maximum size of each of the random deviations (ζ). The random deviation was chosen for each patch in turn, so that they would give a patch positioning similar to that which is expected to result from mass-produced (and slightly noisy) patchy particles. This allowed us to see the difference in tolerance to noise in patch placement.

Figures 3.14 and 3.15 show the effect of noise on the assembly of centred icosahedra;

---

[1] Beware that the term 'budding' has a different meaning in virology

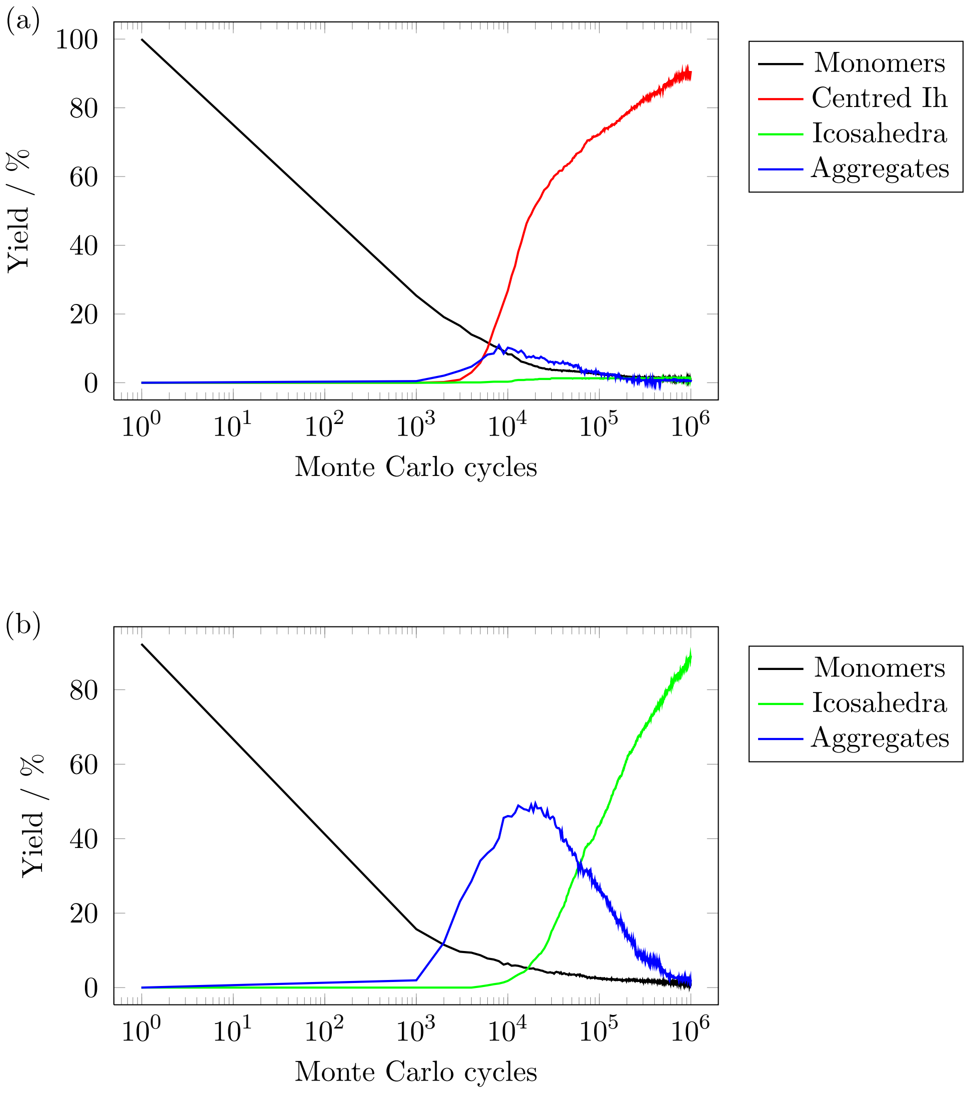

Figure 3.13: Yields as a function of time for the (a) centred ($kT = 0.14\,\varepsilon_{AA}$; other settings are given in table 3.1) and (b) uncentred system ($kT = 0.12\,\varepsilon_{AA}$; other settings are given in table 3.1).

again, we show both centred icosahedra and clusters of size 13. Unsurprisingly, the yield and range of correct assembly decreases as the noise increases. The persistence of 13-mers over centred icosahedra is due to the patch noise preventing the formation of all nearest-neighbour inter-particle bonds (again, due to our definition that a bond has an interaction of at least 40 % of the well-depth). Comparing these data with the uncentred system (figures 3.16 and 3.17) shows that templated assembly leads of a larger robustness against noise in particle production. As large scale synthesis is rarely perfect, a higher tolerance to defects is a useful property for a self-assembling synthetic system to have, again suggesting that templated self-assembly has significant practical advantages.

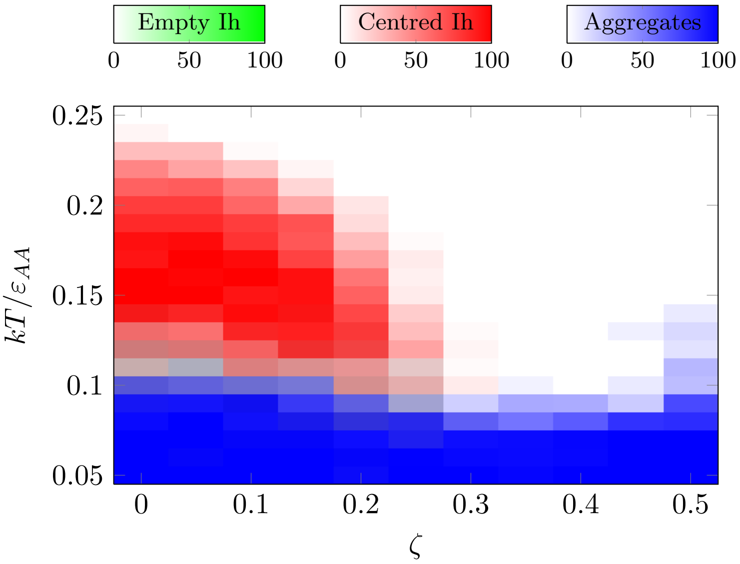

Figure 3.14: The effect of patch noise on the assembly of centred icosahedra (other settings are given in table 3.1).

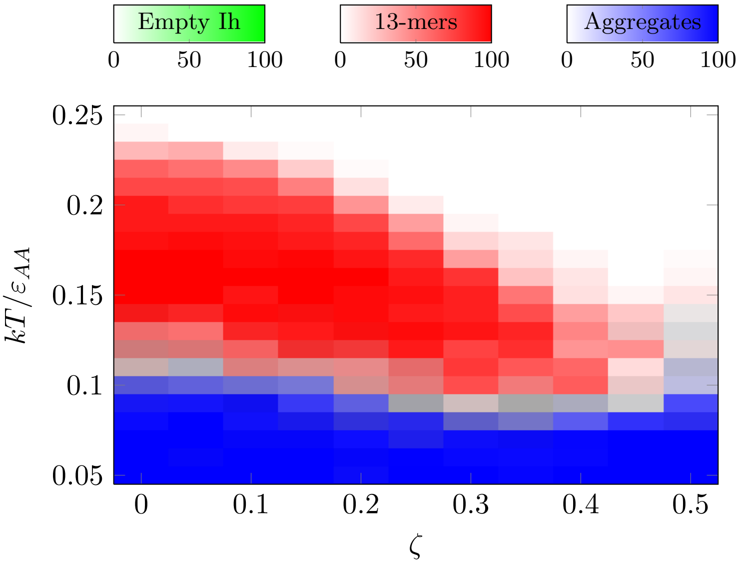

Figure 3.15: The yields of 13-mers formed in the centred icosahedra-forming system as a function of patch noise (other settings are given in table 3.1).

3.5.2 Wang-Landau results

As explained in detail in section 2.2.5, the use of the Wang-Landau algorithm to find thermodynamic information about a system is well-established [149, 198]. We use the fact that the transitions between different states of the system are usually accompanied by a change in energy, and so can be detected by the presence of a heat capacity peak, to delimit the regions of thermodynamic stability for different clusters.

We applied the Wang-Landau algorithm to the icosahedral system initially. We aimed

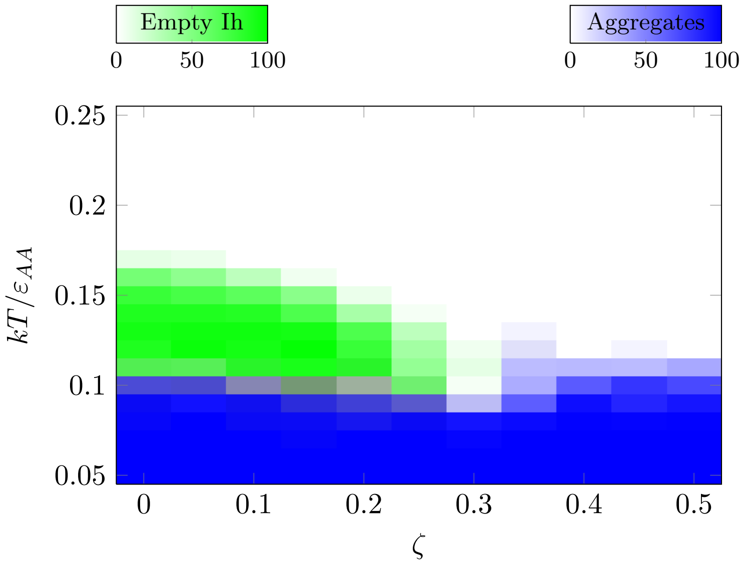

Figure 3.16: The effect of patch noise of assembly in the uncentred icosahedra-forming system (other settings are given in table 3.1).

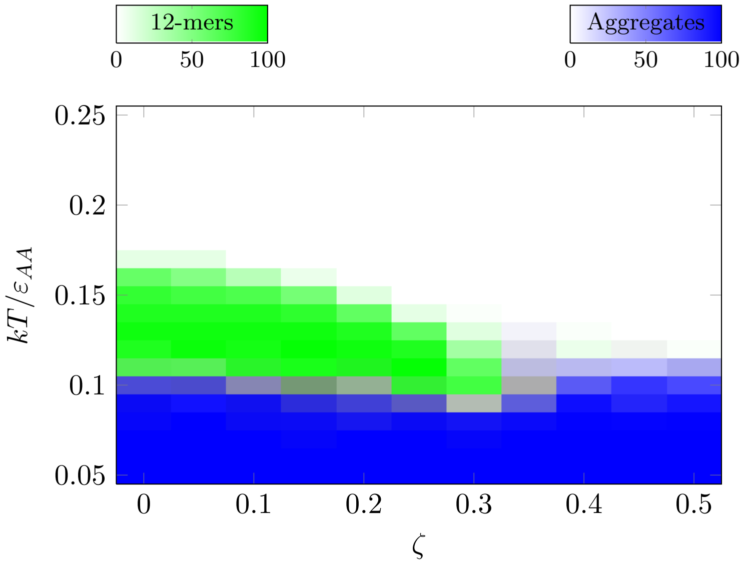

Figure 3.17: 12-mers formed in noisy assembly in the uncentred icosahedra-forming system (other settings are given in table 3.1).

to get thermodynamic data to supplement our understanding from our kinetic data. The results when different factors were changed are summarised below. All of our Wang-Landau simulations for the centred icosahedron-forming system are for a box containing 13 particles (12 shell and one core), at a shell number density of 0.15 / $(\sigma_A^{\mathrm{LJ}})^3$. We examine the effect of using the algorithm for a small system in appendix B.3.1. It was necessary to use a cutoff of 2 σ_A^{LJ}, so that the cutoff was less than half the periodic box width for the icosahedral system. We examine what effect this has on the system in appendix B.4.

Most runs use a minimum energy that is either approximately $2\,\varepsilon$ above the theoretical

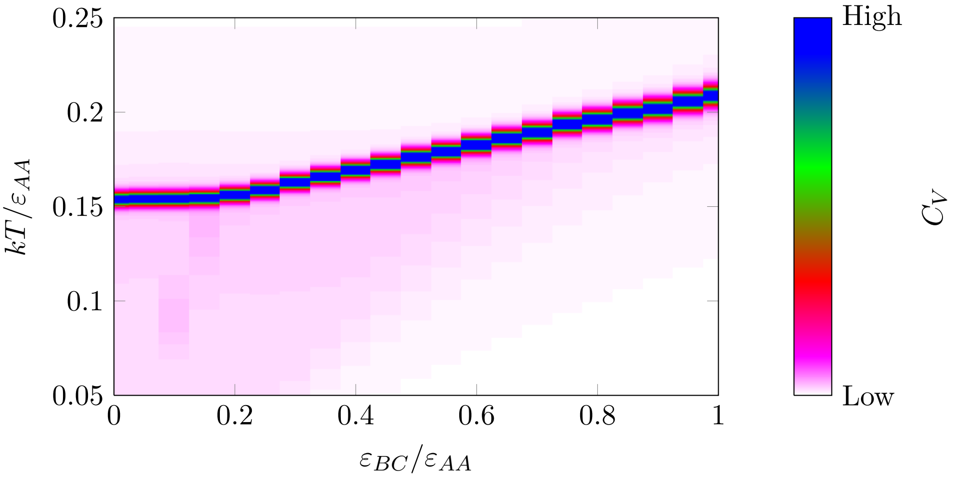

Figure 3.18: The heat capacity of a system of 12 shell and one core particle as a function of temperature and strength of the interaction between the core and shell particle (other settings are given in table 3.1).

minimum, or ε below the lowest energy achieved in a dynamic run. We explore the effect of changing the minimum energy in appendix B.5.1. All runs have at least 200 bins, which is shown in appendix B.5.2 to be sufficient to allow convergence to a sensible degree of accuracy in a reasonable time. A mixture of updating mechanisms were used, which did not affect the outcome; these are described in appendix B.5.3.

Interaction strength

The results (figure 3.18) show the prediction of the Wang-Landau algorithm when we change the strength of the interactions between the central and shell particles. As expected, they show an increase in transition temperature with increasing interaction strength.

The transition between monomer gas and icosahedral gas can clearly be seen in the plot (figure 3.18). At sufficiently low $\varepsilon_{BC}/\varepsilon_{AA}$, the main transition becomes associated with the formation of uncentred and not centred icosahedra (this is the region for which the higher energy transition temperature is no longer dependent on $\varepsilon_{BC}/\varepsilon_{AA}$ because the gain in energy when a core particle is in the centre no longer offsets its loss of entropy at this temperature). Only at much lower temperature do the centred icosahedra become more thermodynamically favourable. In this range of $\varepsilon_{BC}/\varepsilon_{AA}$, two transitions are resolved. On a log plot, at low interaction strengths, two transitions can be seen (figure 3.19). The one which occurs at the lowest temperature, labelled A, is the centre of the equilibrium between centred and uncentred icosahedra. The sharper, larger peak, B, is the transition between uncentred icosahedra and monomers.

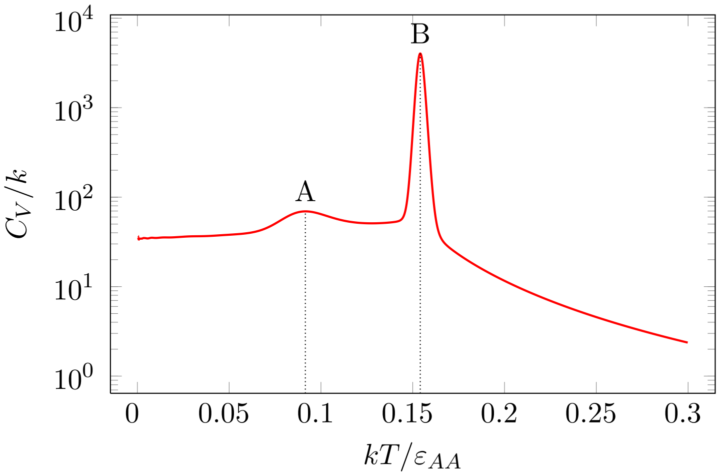

Figure 3.19: The heat capacity at low patch strength (other settings are given in table 3.1). The peak labelled A corresponds to the equilibrium between centred and uncentred icosahedra. The sharper, larger peak, labelled B, corresponds to the transition between uncentred icosahedra and monomers.

Relative particle size

The Wang-Landau results for the effect of changing the size of the central particle are shown in figure 3.20. Between central particles sizes of 0.1 and 0.2 σ_A^{LJ}, we see a transition between uncentred icosahedra and monomers. Then, at core particle sizes between 0.25 and 1.1 σ_A^{LJ}, we have the equilibrium between our centred target clusters and monomers. Up until a central particle size of 0.5 σ_A^{LJ}, the core particle is generally found towards the edge of the cluster, whilst at and above a core size of 0.7 σ_A^{LJ}, the particle is found approximately in the centre of the cluster. At core sizes of 0.55 and 0.6 σ_A^{LJ}, we see wider transitions with somewhat lower heat capacity peaks, as the particle is able to move between these two positions (visualisations of dynamic runs confirm this).

At higher core particle sizes, we see a sharp cutoff in our kinetic yield plot and a sudden change to a mixture of different-sized clusters. This change is reflected in the heat capacity, which exhibits wider, lower peaks (though this is also affected by the fact that we are only using 12 shell patches in our run; see appendix B.3.1). The products' structures are not monodisperse, as demonstrated in figure 3.11.

Weaker shell-core interactions For weaker core-shell interactions, the centred icosahedron is the most stable structure for a narrow range of core sizes, namely from 0.15 to 0.45 σ_A^{LJ}. The big difference from the centred icosahedron-forming system is that above this range, only uncentred icosahedra form.

It is worthwhile for experimentalists to take note of the decrease in other clusters present in this system. If they are only able to produce non-monodisperse core particle

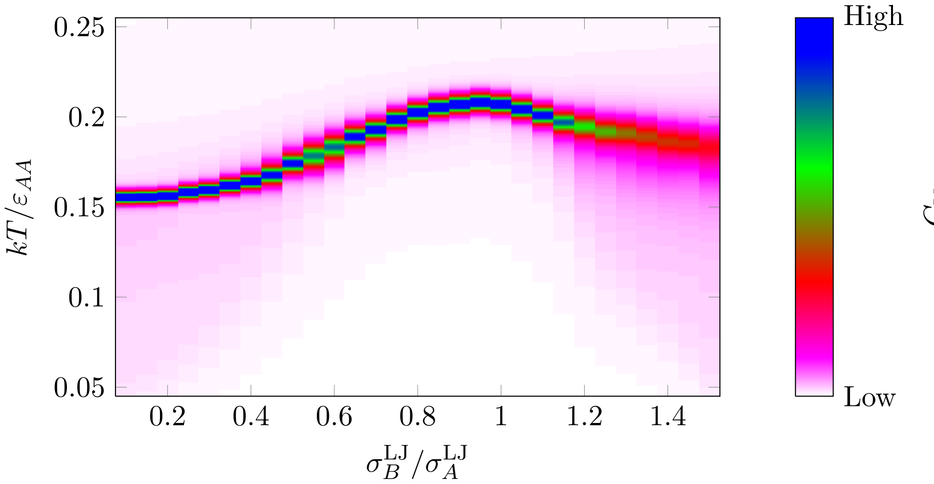

Figure 3.20: The heat capacity of a system of 12 shell and one core particle as a function of temperature and size of the core particle (other settings are given in table 3.1).

sizes, but wish to have only clusters with a given number of particles, then they should consider a system with weaker interactions between core and shell. This would allow them then simply to separate empty and filled clusters, or, by controlling the temperature, they could configure the system to give only yield of the filled cluster or the monomer (for $0.15\,\varepsilon_{AA} < kT < 0.17\,\varepsilon_{AA}$, only centred icosahedra and monomers are possible).

Patch width

The effect of patch width on the thermodynamics of the system can be seen in figure 3.21. These 13-particle simulations allow the thermodynamics associated with clustering to be explored, but cannot capture the thermodynamics associated with competition between aggregates and clusters. At large patch widths, there is some broadening of the heat capacity, which is evidence that the behaviour is changing, and kinetic simulations of larger numbers of particles suggest this may be due to liquid-like aggregates forming. However, we are unable to tell what is happening owing to our small simulation size. We look at the issues of larger simulations in appendix B.3.

3.5.3 Simple approximation-based thermodynamics

Our simulation results on the effect of patch strength can be used to test the validity of a thermodynamic approximation. This approach is further used in the chapter on hierarchical assembly (chapter 4) for the prediction of assembly mechanisms. A simple diagram of the thermodynamic equilibrium lines can be created using the following approximations.

1. The relative entropy of a set of clusters is dominated by the positional entropy of the clusters, such that we can approximate any change in entropy as being directly

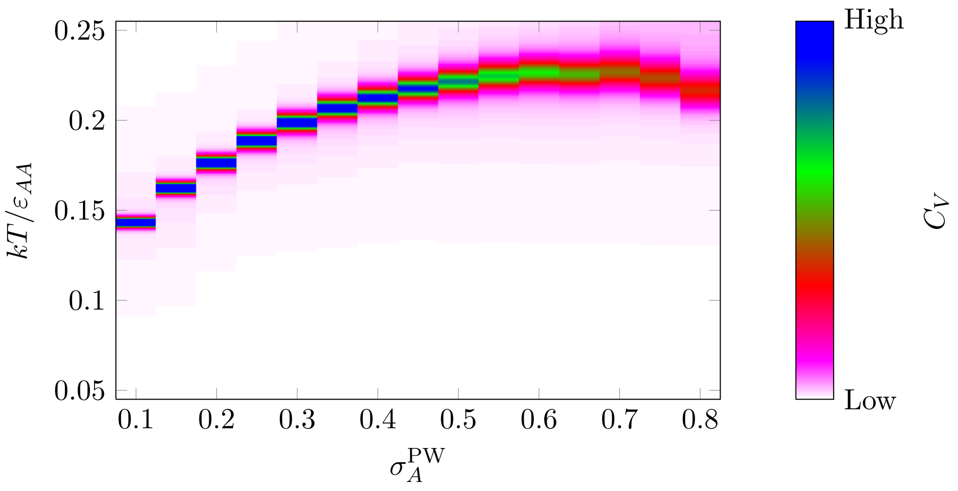

Figure 3.21: The heat capacity of a system of 12 shell and one core particle as a function of temperature and patch width of the shell-shell interactions (other settings are given in table 3.1).

proportional to the change in the number of clusters present (n),

$$\Delta_{\mathrm{trs}} S \propto \Delta_{\mathrm{trs}} n. \tag{3.2}$$

2. The enthalpy of clusters comes entirely from the intra-cluster bonds formed and all bonds are perfectly formed, such that the change in internal energy comes entirely from the change in the number of formed bonds (m_{ij}) of each patch pair ij with bonding energy E_{ij},

$$\Delta_{\mathrm{trs}} U \propto \sum_{ij} \Delta_{\mathrm{trs}} m_{ij} E_{ij}. \tag{3.3}$$

At the transisition temperature, the change in Helmholtz energy is zero, $\Delta_{\mathrm{trs}} A = \Delta_{\mathrm{trs}} U - T_{\mathrm{eq}}\, \Delta_{\mathrm{trs}} S = 0$. Therefore, the transition temperature, T_{eq}, is predicted to obey

$$T_{\mathrm{eq}} \propto \frac{1}{\Delta_{\mathrm{trs}} n} \left(\sum_{ij} \Delta_{\mathrm{trs}} m_{ij} E_{ij} \right). \tag{3.4}$$

To find the proportionality coefficient, we need to know T_{eq} for one set of parameters. It is then easy to calculate the dependence of T_{eq} on the relative interaction strengths of the different patches. For the centred icosahedral system, we have overlaid the expected shape of the equilibrium lines diagram over the Wang-Landau heat capacity data (figure 3.22) as a function of interaction strength, scaled to the middle of the transition from the Wang-Landau results (figure 3.22), at $\varepsilon_{BC} = \varepsilon_{AA}$.

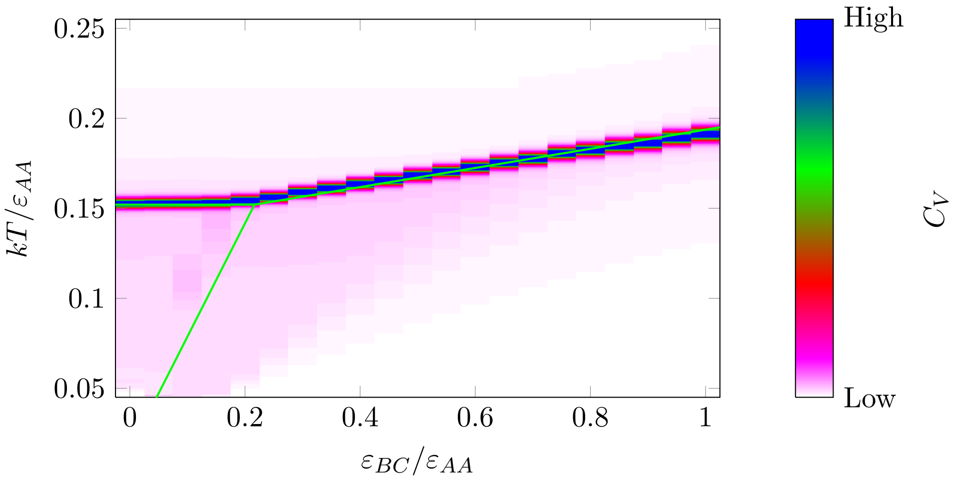

Figure 3.22: Simple approximation of phase boundaries compared to the Wang-Landau results (other settings are given in table 3.1).

This simple approximation technique exhibits clear agreement with our data. At lower interaction strengths, there is some divergence from our results, especially the position of the centred/uncentred icosahedron transition line, but it is clearly a fair approximate system for finding the transition lines.

3.6 Formation of an otherwise inaccessible shape: the dodecahedron

Having shown that formation of patchy particle clusters with a target cluster that could be accessed using other non-torsional assembly methods is possible, we then moved onto looking at the formation of shapes which previous research has shown are rarely formed for a variety of reasons [112]. The dodecahedron is the prime example of this type of system, and one which we consider in detail here.

The formation of a dodecahedral cluster has been shown previously not to be possible without the use of torsional constraints [112]. The temperatures at which kinetic aggregates form are above the temperature at which the complete cluster forms. There is no window in temperature over which the dodecahedron is the only stable species with respect to monomers. There are only three bonds per particle, so the correct formation of the structure does not have as much enthalpy per particle as the icosahedral system, and the loss in entropy is significantly higher (as there are 20 particles in the shell rather than 12 for the same gain in energy of the completed cluster).

With templating, however, we were able to form centred dodecahedral clusters (figure 3.23) without any torsional constraints being used. An analysis of our findings is given

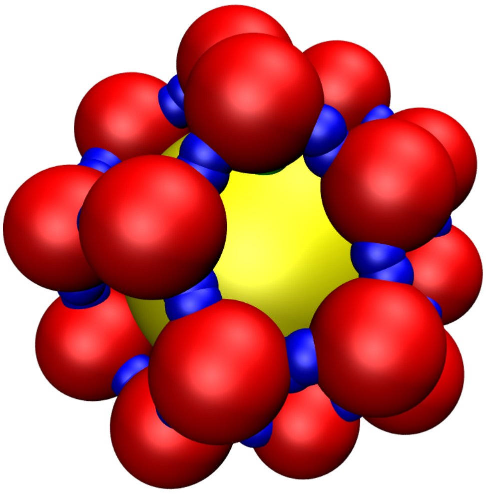

Figure 3.23: Centred dodecahedron.

below. We begin by taking a look into the effect of varying the patch width on the yields of dodecahedra, and then look more at the effects of interaction strength and core particle size.

3.6.1 Dynamic yields

As with the icosahedral system, we begin by examining how the results of dynamic Monte Carlo simulations depend on the parameters of our model.

Patch width

The effect of changing the patch width on the yield of dodecahedra was similar to that of the icosahedral cluster forming system, with the patch widths affecting the formation of aggregates and clusters in a similar manner. We mainly highlight the differences between the two systems here. We again had two patch widths which we could change, the shell-shell interaction, A, and the shell-central particle interaction, B.

Shell-shell patch width (σ_A^{PW}) When varying the patch width of the shell particles, we found a similar effect to that seen with the icosahedron, but with no interference from the uncentred system (which is kinetically inaccessible). At wide patch widths, aggregates are still preferentially formed (figure 3.24). Unlike in the centred icosahedra forming system, there is a clear white region in figure 3.24 between successful formation of the target cluster and aggregation at wide patch width where neither aggregates nor dodecahedra are formed. This region is caused by the formation of non-dodecahedral clusters around a core particle. An example is shown in figure 3.25. This type of structure is less common in the centred icosahedra-forming system, as the difference in angle between a triangular and a square

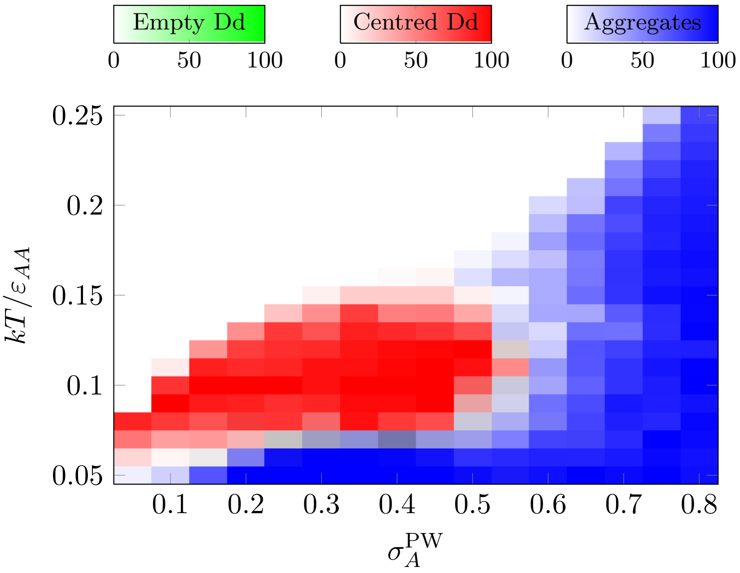

Figure 3.24: Plot of successful assembly of the centred dodecahedron-forming system as a function of the shell-shell patch width (other settings are given in table 3.1).

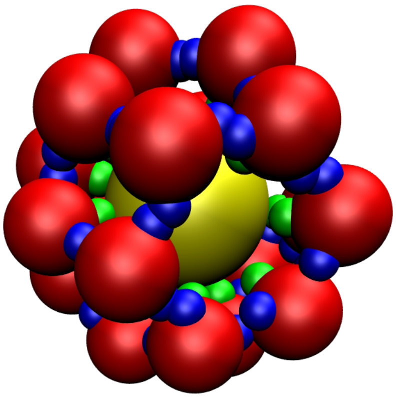

Figure 3.25: An example cluster formed at wide patch widths showing hexagonal and square faces ($kT = 0.1\,\varepsilon_{AB}$, $\sigma_A^{\mathrm{PW}} = 0.6$; other settings are given in table 3.1).

face is $0.17\,\pi$ radians, whereas to change from a pentagonal face to a square or hexagonal face (as occurs in a dodecahedron) requires a distortion of only approximately $0.1\,\pi$ or $0.06\,\pi$ radians, respectively. We analyse the alternative structures in more detail later in this section.

Shell-core patch width (σ_B^{PW}) As in the icosahedral system, we found that the width of the patch on the shell particles which pointed to the central particle had little effect except at extremes (figure 3.26). It is worth noting that the maximum yield occurs across a wider range of patch widths than for the icosahedral system (figure 3.9), but that, outside of this range, the yield decreases more quickly than it did for the icosahedral system.

At wide shell-core patch widths, in addition to the structures seen at wide shell-shell

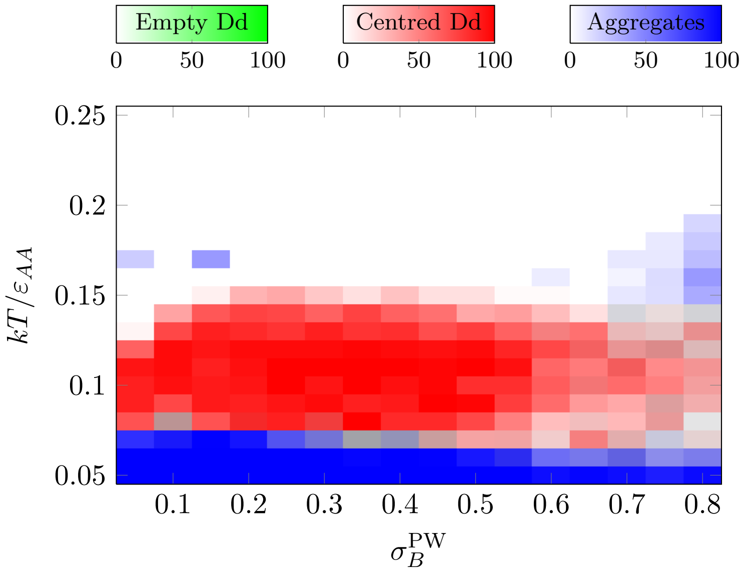

Figure 3.26: Plot of successful assembly of the centred dodecahedron-forming system as a function of shell-core patch width (other settings are given in table 3.1).

patch widths (which are less common, but still present), the main competition to assembly comes from the formation of large clusters with multiple encapsulated cores (figure 3.27). If the formation of multiply-cored clusters were a problem in practice, then ensuring some form of repulsion between the cores should help destabilise such structures, and would, in fact, increase the range over which centred dodecahedra could form.

Comparison with “torsional” results

As we are unable to compare the centred dodecahedron directly with our uncentred non-torsional system, we look at the effect of varying the torsional patch widths on the uncentred dodecahedral system here (torsional constraints are explained in section 2.1.1). We vary the reciprocal of the torsional patch width (σ_A^{Tors}), rather than the patch width, so we can see the effect of having an effectively infinite torsional patch width (equivalent to the non-torsionally constrained case).

Figure 3.28 demonstrates that torsional constraints in our potential clearly destabilise the formation of aggregates. It is worthwhile noting that in all cases, the formation of dodecahedral clusters actually occurred over a narrower range of conditions compared to our encapsulated system. This shows that, in certain cases, templating is a more effective approach to ensuring the formation of target clusters than the use of torsional constraints. However, total repression of aggregation is only possible with narrow torsional patch widths, so this approach may have advantages in some conditions.

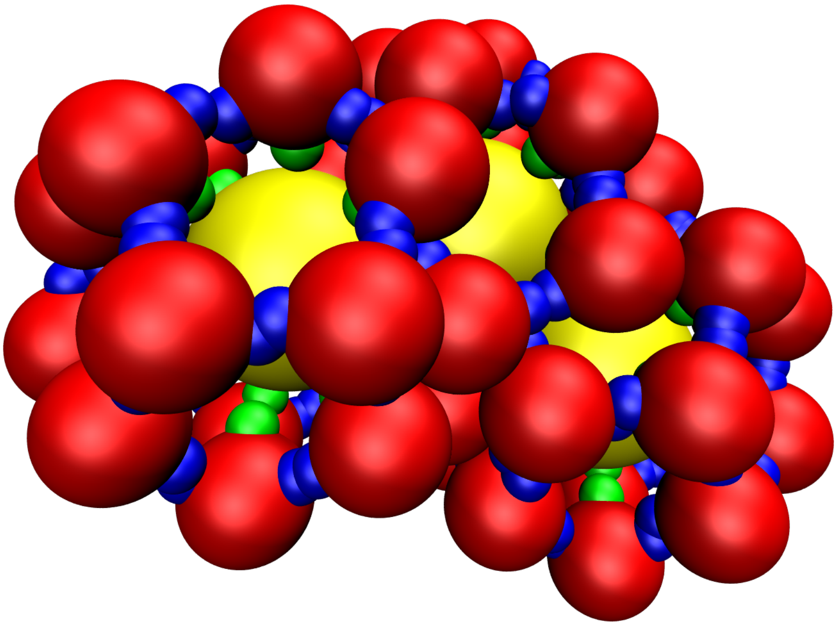

Figure 3.27: An example cluster formed at wide shell-core patches showing multiple encapsulated cores ($kT = 0.1\,\varepsilon_{AB}$, $\sigma_B^{\mathrm{PW}} = 0.8$; other settings are given in table 3.1).

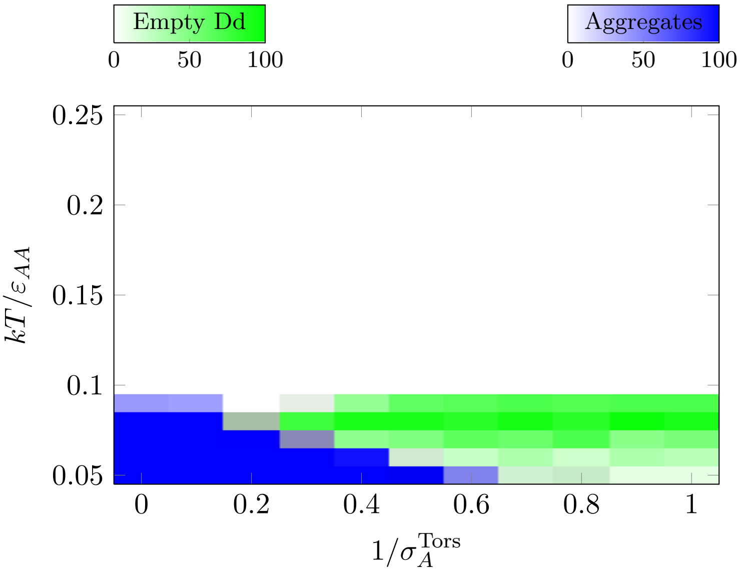

Figure 3.28: Plot of successful assembly of the uncentred dodecahedron-forming system with torsional constraints as a function of the reciprocal of the torsional patch width (angular patch width fixed at 0.35 radians, other settings are given in table 3.1).

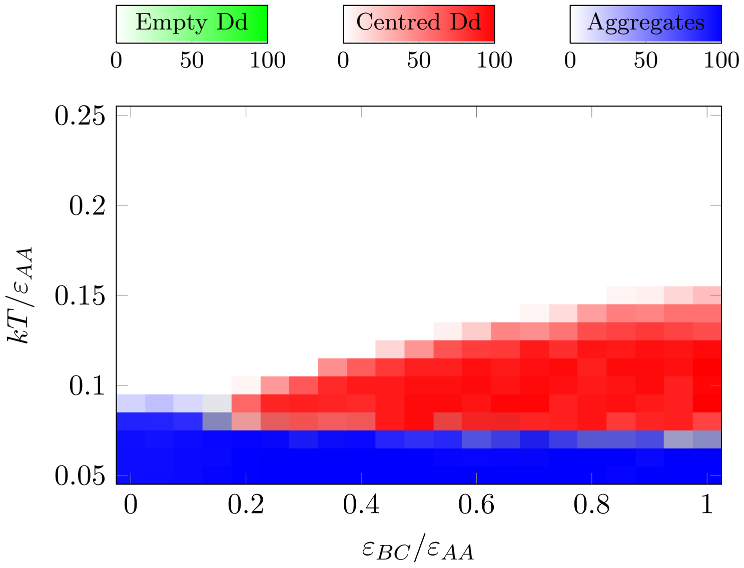

Figure 3.29: Plot of successful assembly as a function of the interaction strength between core and shell particles (other settings are given in table 3.1).

Interaction strength

The effect of the strength of the interaction between the core and shell particles (ε_{BC}) in dodecahedra is more interesting that its icosahedral counterpart. Figure 3.29 is the key to understanding why templating is a useful strategy. We are able to open up a window where the temperature of clustering is above the kinetic temperature of aggregation. Increasing ε_{BC} increases T_{clus} without changing the temperature where aggregates are stable with respect to monomers, T_{agg}, significantly. In order to see a reasonable yield of centred dodecahedra, the shell-core interaction must be at least 20 % of the shell-shell interaction.

Core particle size

As with icosahedra, the definition of completed centred dodecahedra excludes certain closed clusters where the central particle is unable to bond to all of the shell particles. However, unlike in the icosahedral system, all clusters of size 21 are centred dodecahedra. The results in figure 3.30 show that, unlike in the icosahedral case, the yield of centred dodecahedra depends strongly on the size of the internal particle, with only a narrow range of sizes allowing correct assembly. Particles with three patches are likely to have more plasticity in the structures that they can form (this is exemplified by the wide range of fullerene structures), and thus the final structure formed in this system depends strongly on the size of the central particle, which lowers the free energy of a shell with a certain number of particles.

The dodecahedral system shows further interesting behaviour at small central particle

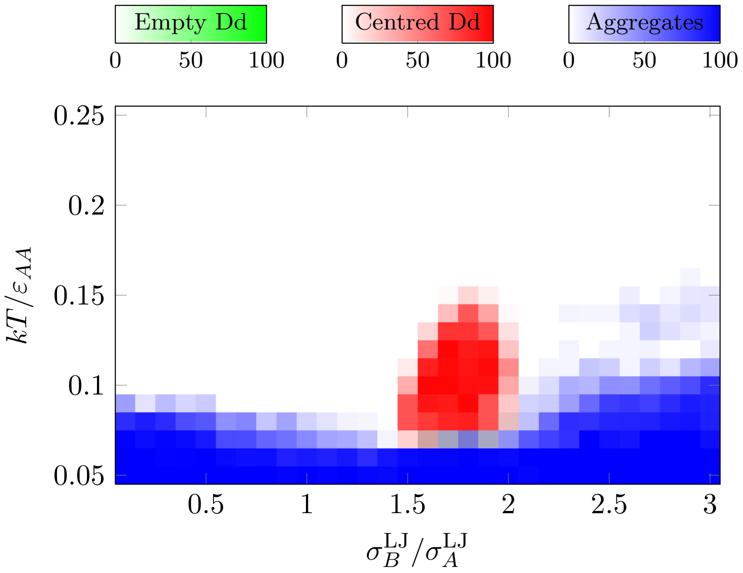

Figure 3.30: Plot of successful assembly as a function of the size of the core particle (other settings are given in table 3.1).

sizes. The system is able to form shells that fit around the central particle, but that are no longer dodecahedral. Figure 3.31 shows the yield of clusters of relevant sizes as a function of the size of the internal particle. It is clear that there exists a region where clusters with between 15 and 17 members are formed in preference to completed dodecahedra, though no specific sized cluster is ever the only member present. Other sized clusters are also present in our simulations, and generally those with an even number of shell particles are more stable than those with an odd number. We consider the thermodynamic stability of these clusters in section 3.6.2. This behaviour is somewhat similar to results for the assembly of virus capsids around gold nano-particles, where changes in their packing are seen as a function of the nano-particle size [183], although these seem to keep their icosahedral symmetry and just vary the T-number, while our system changes to become less symmetric at smaller core particle sizes.

As seen before, these smaller clusters can form by straining the bond angle slightly to form squares as well as pentagons on the surface of the core particle. An example cluster is shown in figure 3.32(a). The clusters have many different forms, with a different connectivity of square and pentagonal faces. Generally, two square faces do not share an edge, as this induces more strain, but some clusters were formed in which this did occur. At larger core particle sizes, there is similarly a number of larger clusters formed around the central particle. These have hexagonal and pentagonal faces. An example cluster is shown in figure 3.32(b).

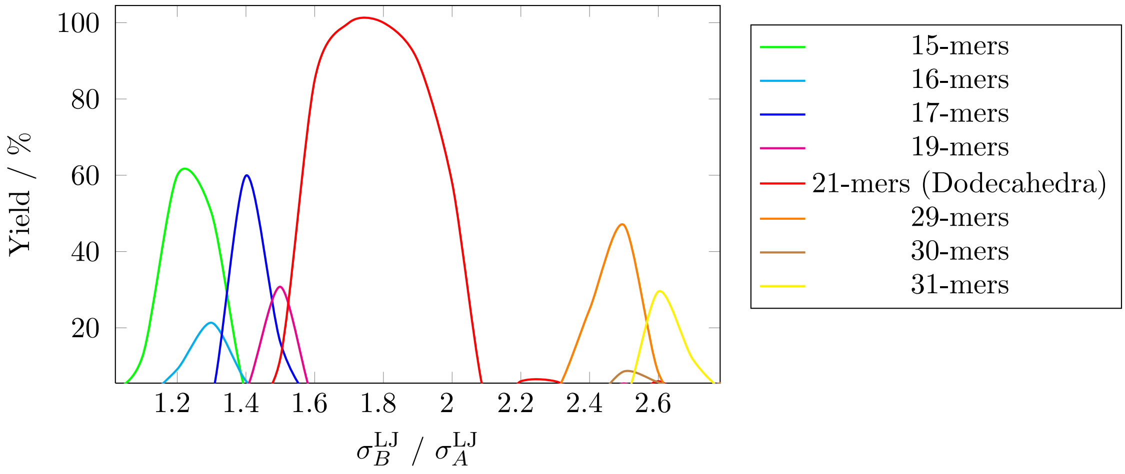

Figure 3.31: A smoothed plot of the yields of selected clusters as a function of the relative size of the internal particle (at $kT/\varepsilon = 0.11$; other settings are given in table 3.1). Note that yields below 5 % are not shown.

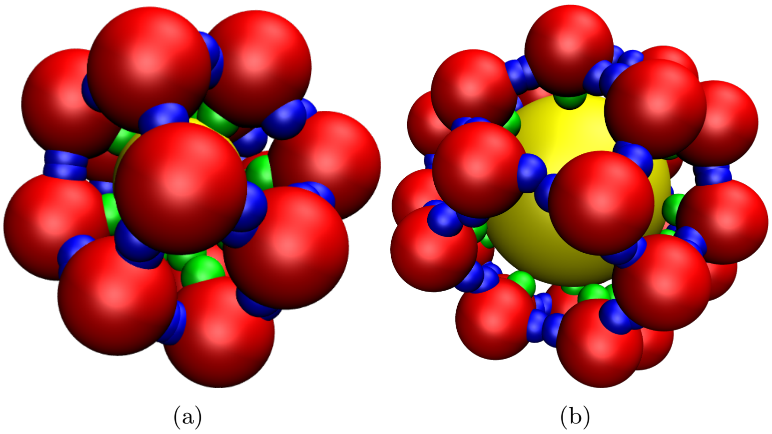
(a) (b)

Figure 3.32: An example cluster formed of (a) 15 particles with a core particle of 110 % the size of the shell particle ($kT/\varepsilon = 0.11$), (b) 30 particles with a core particle of 260 % the size of the shell particle ($kT/\varepsilon = 0.11$). Other settings are given in table 3.1).

3.6.2 Thermodynamic data

Thermodynamic data were obtained using the Wang-Landau algorithm. The effect of changing the different variables can be seen below. These runs are all completed using 200 bins to sample energy space, as this allows thermodynamic data to be obtained in a reasonable time. The cutoff for the Wang-Landau runs was 2.5 σ_A^{LJ} (the effect of this is considered in appendix B.4), and all runs are at a shell number density of 0.15 $(\sigma_A^{\mathrm{LJ}})^{-3}$.

Patch width

The thermodynamic data (figure 3.33) show some different aspects from the icosahedral system (figure 3.21). Similar to the icosahedral system, as the patch width is decreased, the

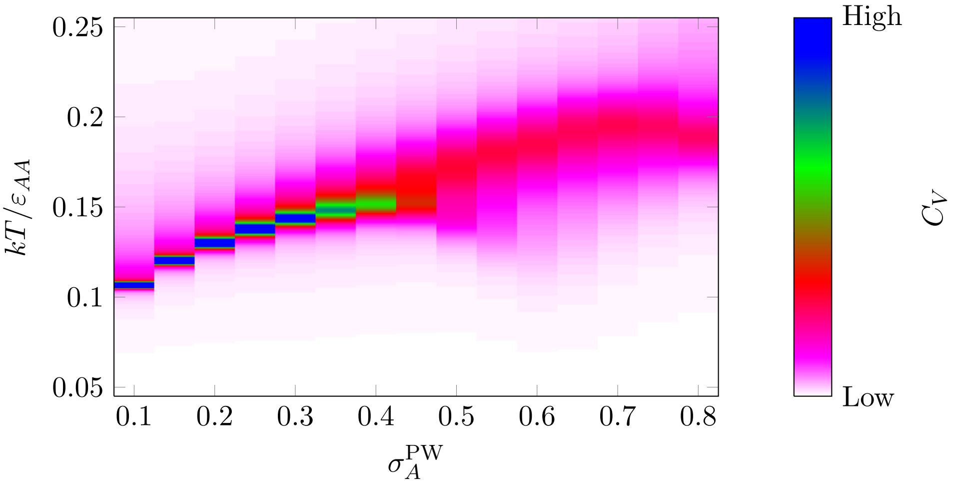

Figure 3.33: The heat capacity of a system of 20 shell and one core particle as a function of temperature and patch width of the shell particle (other settings are given in table 3.1).

transition temperature decreases because of the lower vibrational entropy of the clusters and the transition becomes sharper due to the increased entropy difference between monomers and clusters. The change to the transition at low patch width can be seen in figure 3.34.

As shown previously, at larger patch widths, the formation of non-dodecahedral small clusters becomes possible (figure 3.27), and at very large patch widths, liquid-vapour phase separation can be seen. Some of these changes are also reflected in the heat capacity data (note that as the Wang-Landau results are for a 21-particle system, the size of clusters is inherently restricted). The heat capacity peak splits into two at the point where non-dodecahedral and dodecahedral clusters are stable (figure 3.35). As the dodecahedral cluster has less strain (more negative internal energy), but also has slightly lower entropy, it is the lower peak. At very wide patch widths, we only see one peak, corresponding to the transition to disordered clusters: the size of these clusters is limited by the number of particles in the simulation, so this transition may not correspond directly to a transition in the bulk.

Interaction strength

The effect of changing the strength of the interactions between the core and shell particles is especially interesting in this system (figure 3.36) compared to the icosahedral system.

At very low interaction strengths, we are able to resolve a marked change in behaviour (figure 3.37). With a shell-core interaction of $0.1\,\varepsilon_{AA}$, there is a clear, sharp peak, which we have shown in previous plots to correspond to a transition between monomers and dodecahedral clusters. At lower interaction strengths, there is a broad peak, which occurs as a result of the formation of disordered clusters with little structure. These disordered

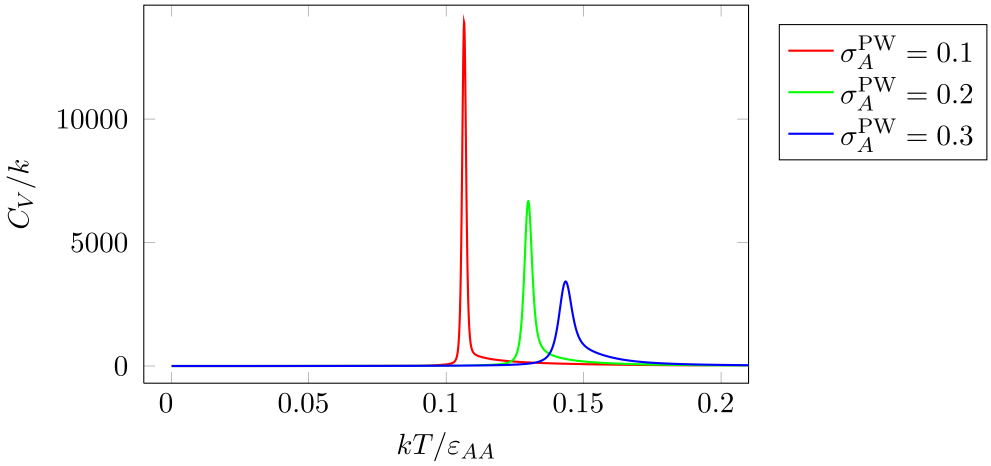

Figure 3.34: A plot of the heat capacity at low patch widths (other settings are given in table 3.1).

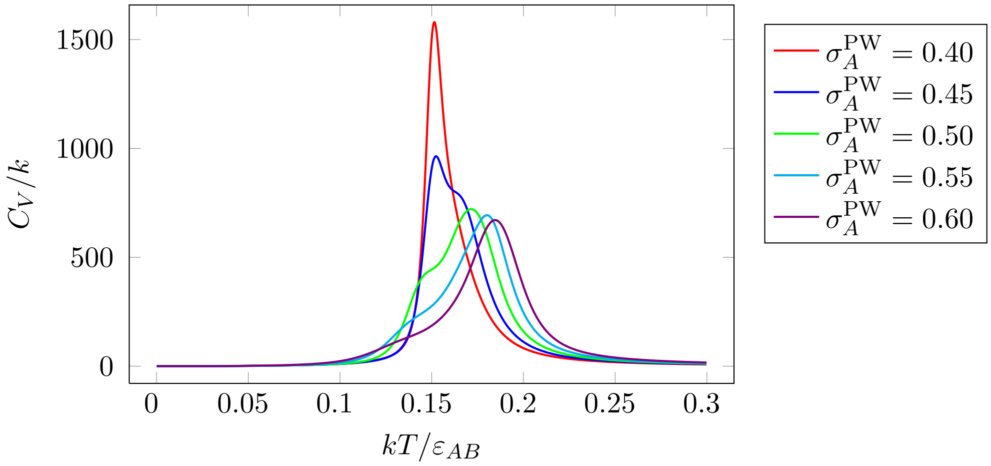

Figure 3.35: A plot of the heat capacity at higher patch widths (other settings are given in table 3.1).

clusters may not have more than 21 particles in our simulations, so this transition may not correspond exactly to the transition in a bulk system.

Core particle size

The Wang-Landau algorithm gives much clearer results using different settings from the previous runs (figure 3.38). These are at a lower resolution and with a higher minimum energy. The results with the same settings are explained in appendix B.5.1.

As mentioned in the dynamic data section on particle size (section 3.6.1), the lack of uncentred dodecahedra and the formation of non-dodecahedral clusters formed around a central particle leads to different behaviour than in our icosahedral system. At very small core particle sizes, we see one transition corresponding to the formation of unstructured clusters; again, these are limited by the number of particles in our simulation box. As

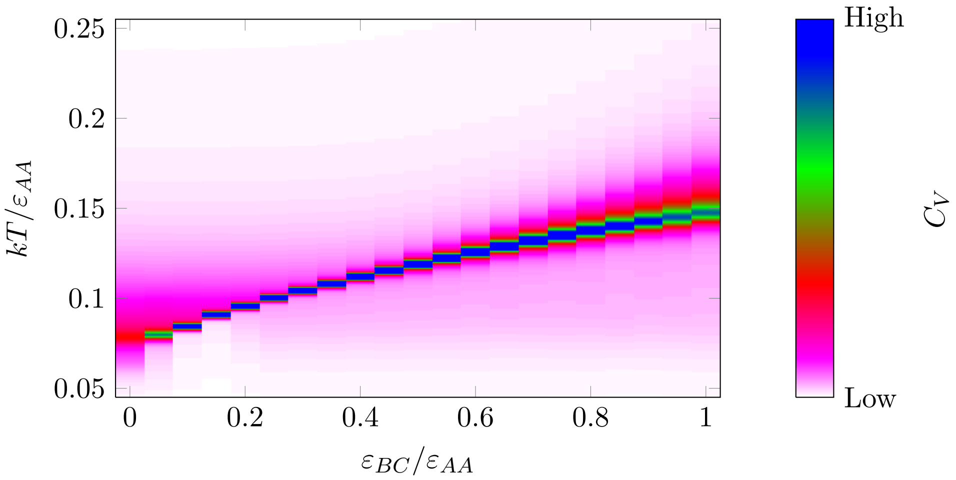

Figure 3.36: The heat capacity of a system of 20 shell and one core particle as a function of temperature and interaction strength between the shell and core particles (other settings are given in table 3.1).

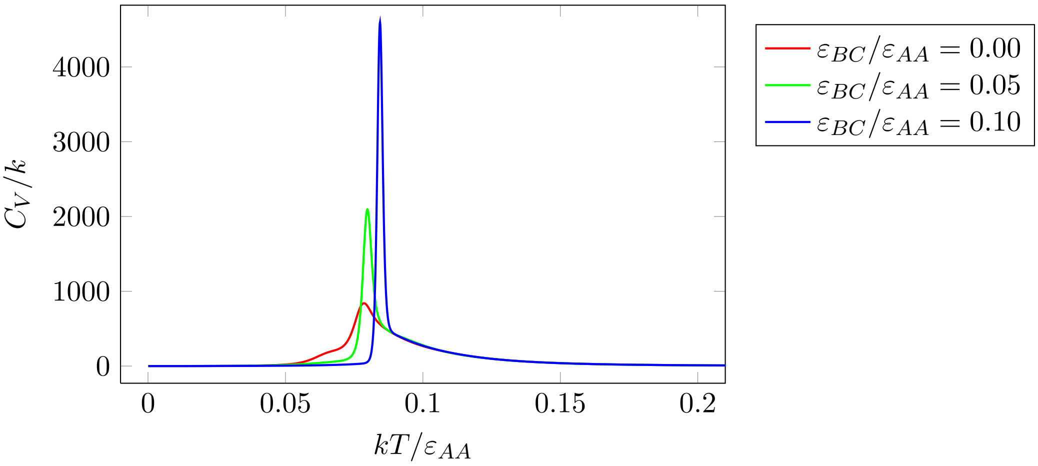

Figure 3.37: A plot of the heat capacity of the dodecahedral system at low core/shell interaction strength (other settings are given in table 3.1).

we increase the core particle size, we see a region of thermodynamic stability of smaller non-dodecahedral clusters with a central particle, as predicted by our kinetic results. There appear to be two transitions, one corresponding to the transition between monomers and centred, non-dodecahedral clusters (the higher temperature), and one occurs where small unstructured clusters are formed from the remaining shell particles (see appendix B.5.1 for further details). Centred dodecahedral clusters are stable where the core particle's radius is between $1.5\,\sigma_A^{\mathrm{LJ}}$ and $2\,\sigma_A^{\mathrm{LJ}}$. This transition is clearer in the diagram in appendix B.5.1, though it contains many more artefacts as a result of our settings. At very large core particle sizes, the formation of polydisperse core containing clusters occurs. Using 20 shell particles may have an effect on the exact location of this transition, but we have clearly shown it to occur in our dynamic yield plots.

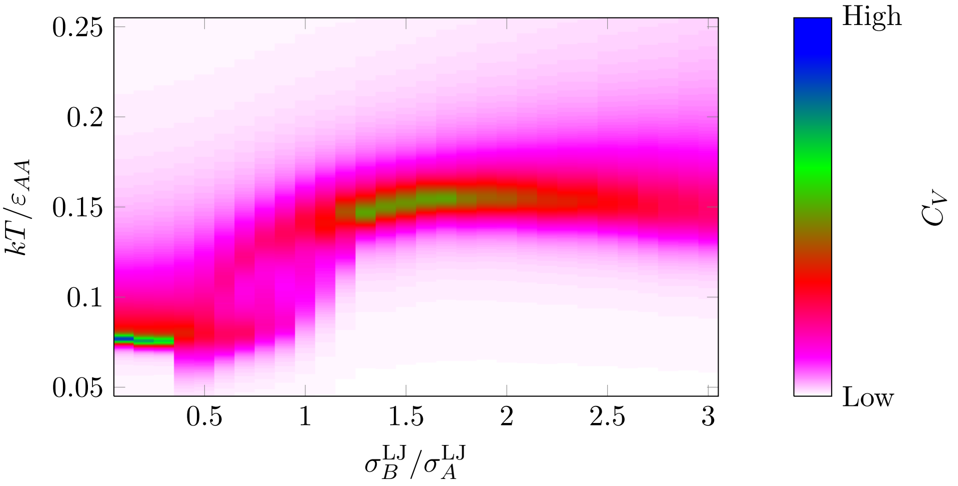

Figure 3.38: The heat capacity of a system of 20 shell and one core particle as a function of temperature and the relative size of the core and shell particles (other settings are given in table 3.1).

3.7 Analysing the structure of the different aggregates

As shown in figure 3.4, there are two main forms of aggregates that we encounter in our simulation. The first are thermodynamic and a result of phase separation into liquid and vapour (the density used for simulations lies in the two-phase region of the phase diagram) that occur at large patch widths, where the potential is becoming close to isotropic. The second are kinetic aggregates (glass-like) formed as a result of trapping at low temperature.

In this section, we look at the structural differences between the different aggregates, especially concentrating on the behaviour of the core particle in an encapsulated system.

3.7.1 Icosahedral system

We used the simulation results from our kinetic simulations in which we varied the patch width of the shell-shell interactions in order to explore the differences in our aggregates' structures. We first considered the location of the core particle, then looked at the underlying morphology of some of the aggregates.

Core particle

The core particles can be in fundamentally different environments from the shell particles. Figure 3.39 shows the location of the core particles as a function of patch width. They show that generally, the core particle is in the same state as the bulk (figure 3.3), except in the area in which a liquid-like structure should form. The upper right corners of figures 3.39(a) and (b) show that the core particles are not as common in the liquid-like aggregates as the shell particles. They are generally excluded and this can be seen in figure 3.40(b).

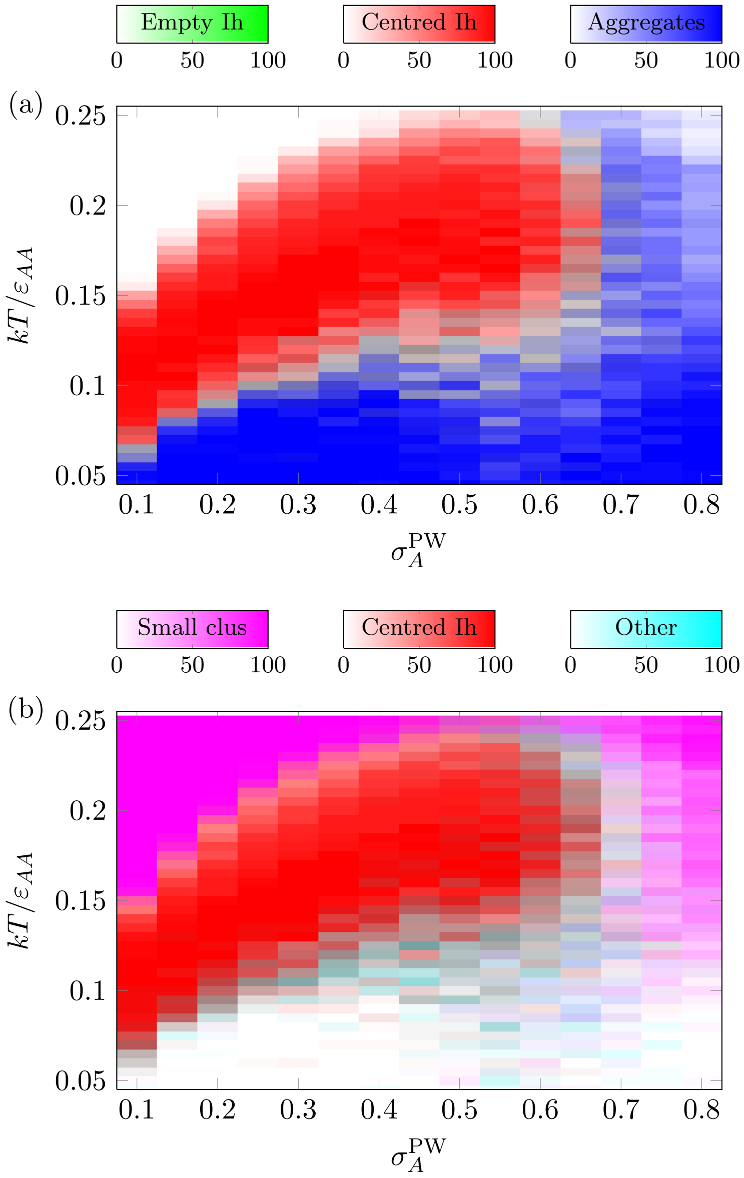

Figure 3.39: Amount of central particles located in (a) the target cluster and aggregates (b) small clusters (12 or fewer particles) and clusters larger than centred icosahedra but not large enough to be included as aggregates (14-24 particles). Other settings are given in table 3.1.

Morphology

The radial distribution functions (RDFs) of the aggregates and the completed target clusters can be seen in figures 3.41, 3.42 and 3.43. The three clear peaks in the icosahedral cluster region's RDF correspond to the correct distances for the shell particles within an icosahedron. It is interesting that in both the glass and the liquid region, these first three peaks can be seen, as there is still some icosahedron-like structure present in the aggregates, though it is less well defined in the liquid than in the glass. At very high patch widths, the system would recover the Lennard-Jones behaviour, where the liquid does not have

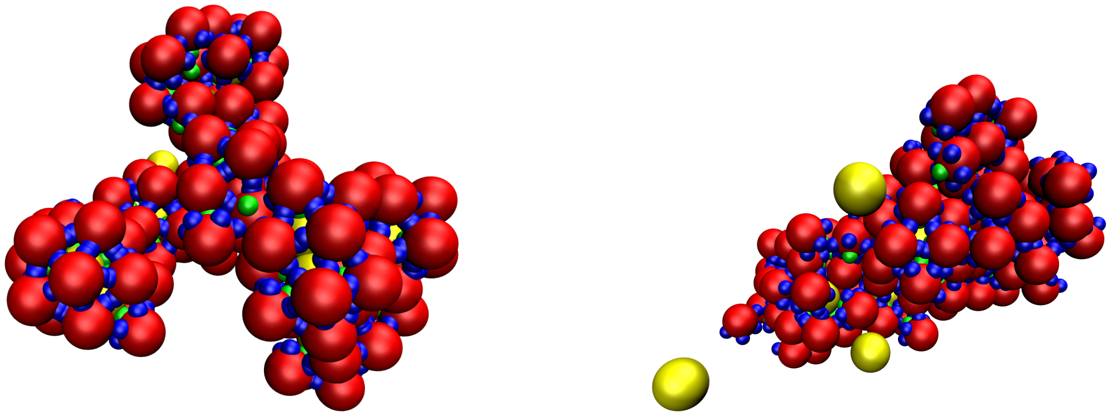
(a) (b)

Figure 3.40: The structure of the (a) glass-forming region ($\sigma_A^{\mathrm{PW}} = 0.35$, $kT/\varepsilon_{AA} = 0.05$; other settings are given in table 3.1) and (b) liquid-forming region ($\sigma_A^{\mathrm{PW}} = 0.8$, $kT/\varepsilon_{AA} = 0.2$; other settings are given in table 3.1). Note the free core particles in the latter case.

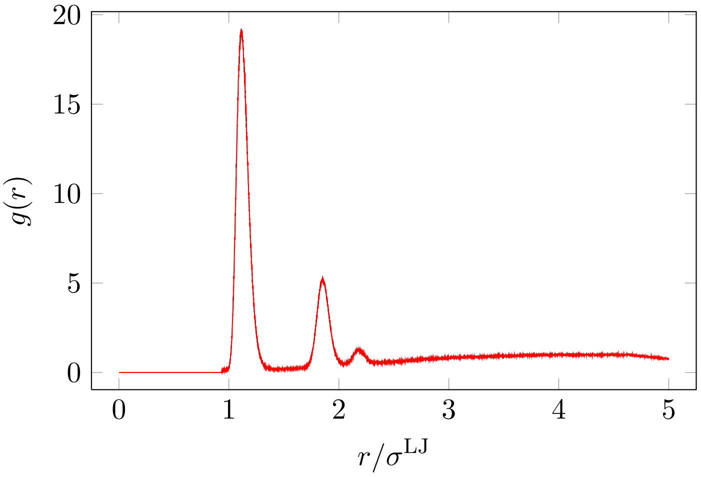

Figure 3.41: Icosahedral forming region radial distribution function ($\sigma_A^{\mathrm{PW}} = 0.35$, $kT/\varepsilon_{AA} = 0.15$; other settings are given in table 3.1).

such a clear structure.

The differences in structure can also be seen in figure 3.40. The liquid region has a number of core particles (shown in yellow) excluded from the aggregate, which is not the case for the glassy aggregate, and the glassy aggregate has clearer icosahedral cluster-shaped structure when compared to the liquid.

3.8 Introducing structure to the encapsulated particle

As we are trying to find the most efficient assembly method possible using encapsulated particles, it is important to explore the effect of changing the interactions and patches of the central particle. We did this is two ways.

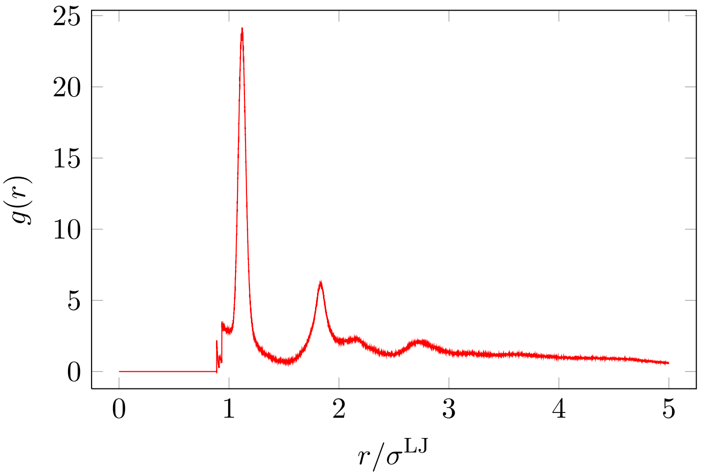

Figure 3.42: Glass forming region radial distribution function ($\sigma_A^{\mathrm{PW}} = 0.35$, $kT/\varepsilon_{AA} = 0.05$; other settings are given in table 3.1).

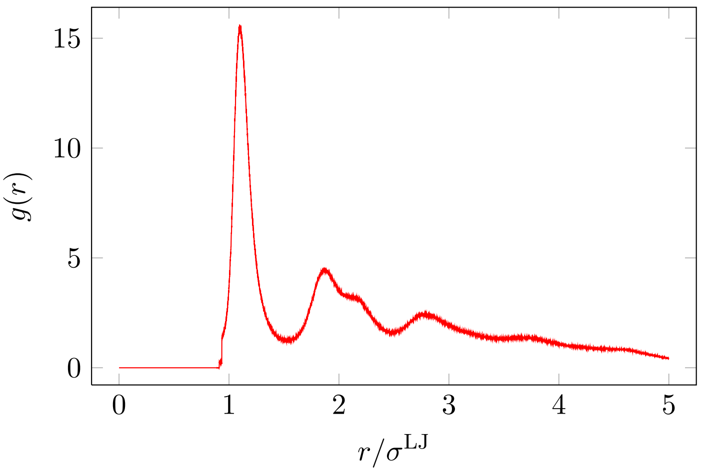

Figure 3.43: Liquid forming region radial distribution function ($\sigma_A^{\mathrm{PW}} = 0.8$, $kT/\varepsilon_{AA} = 0.2$; other settings are given in table 3.1).

3.8.1 Changing the patch width of the encapsulated particle

Initially, we looked at modifying the encapsulated particle from an isotropic single patch to a single patch of varying width. The aim was to locate the point at which the curvature of a partially-formed shell encouraged assembly significantly such that the interaction with the encapsulating particle was no longer the driving force for further growth. The width of the patch was originally set to 100 radians. As this is the standard deviation of the Gaussian in radians, this is equivalent to an isotropic attraction.

By looking at the effect of changing the patch width, we were able to see the surface coverage needed to encourage clusters to form, and thus to learn something about the position of the barrier to the formation of the target structure.

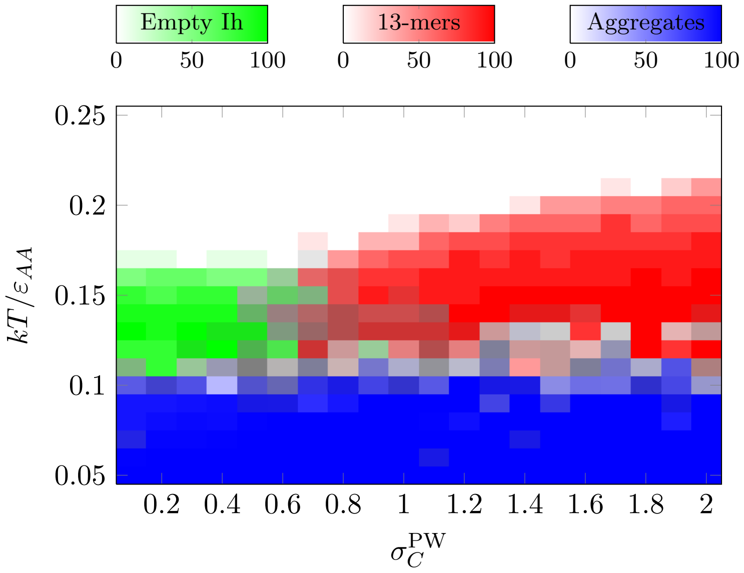

Figure 3.44: Plot of successful assembly of clusters of size 13 as a function of the core-shell patch (other settings are given in table 3.1).

Icosahedral system

With the icosahedral system, there was competition between the formation of centred and uncentred icosahedra. In figure 3.44, we plot the yield of clusters of size 13 (which are all centred icosahedra in shape, though not all shell particles may bind to the core). The centred cluster was able to form at a wide range of patch widths, suggesting again that the initial nucleation around the core is the key step causing centred rather than uncentred icosahedra to form.

Dodecahedral system

In the dodecahedral system, we looked at the effect of wide (figure 3.45) and narrow patches (figures 3.46). We plot only 21-mers, rather than our true target, as these are all dodecahedral in shape, but, as before, the central particle is not always able to bond to all of the shell particles. The absence of competition between centred and uncentred systems meant that we were able to see the effect of the internal patch more clearly.

It is clear that as long as there is a central particle with a reasonable patch width, it is possible to form a cluster. Any patch with a width below 2.6 radians on the central particle is unable to interact with *all* shell particles simultaneously with an attraction of at least 50 % of the well depth, and one with a width below 1.4 radians is unable to interact with an attraction of at least 10 % of the well depth with all shell particles simultaneously. Yield is seen all the way down to a patch width of 0.8 radians, which cannot interact significantly with all particles in the cluster, suggesting the presence of the central particle

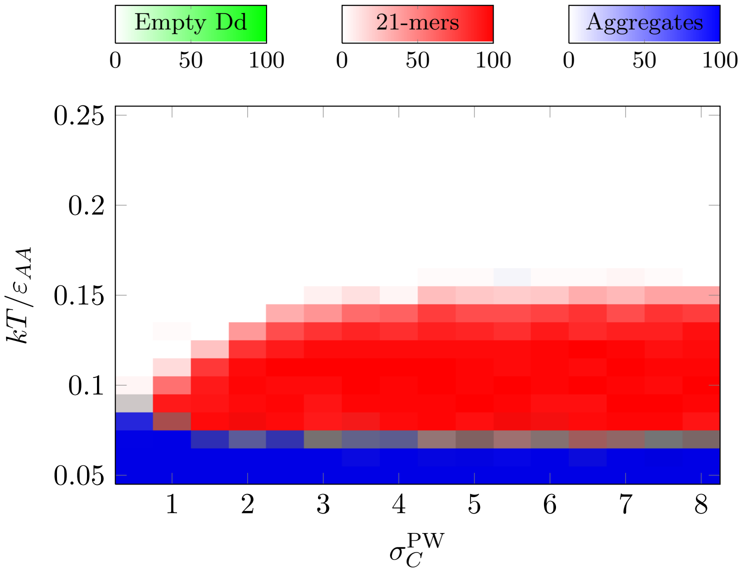

Figure 3.45: Plot of successful assembly of 21-mers in the dodecahedral system as a function of core-shell patch width (other settings are given in table 3.1).

is important for initially templating the growth and that interactions with the growing cluster not energetic adsorption, are sufficient to lead to completion of the shell. Note that the excluded volume of the central particle will have some effect in preventing certain incorrect configurations.

We have shown that it is not important for the attraction of the central particle to be fully isotropic, provided the patch width is sufficiently wide to encourage the initial correct curvature of the growing clusters and to provide an efficient energetic stabilisation of the cluster so that $T_{\text{clus}} > T_{\text{agg}}$. At narrow patch widths, we also limit the surface diffusion of particles, which we look at in more detail in the next section.

3.8.2 Using multiple patches

So far, we have looked at isotropic and singly-patched particles as the templating centre, as this is simpler and likely to be easier to synthesise. As the methods to create patchy particles become more well-developed, more complicated arrangements and accurate placement of patches will be possible. Therefore, we examined how a “perfect” template (one with specific patches that point directly at the shell particles in the target structure) compares to an isotropic centred particle. By widening the patch widths of the multiply-patched template particle, we are able to move from a regime of distinct multiple patches back to the original isotropic system, allowing us to detect if there is a significant change in the assembly behaviour between the two systems. This system also allows us to examine the rôle of surface diffusion in our assembly. As the patches become narrower, we expect that

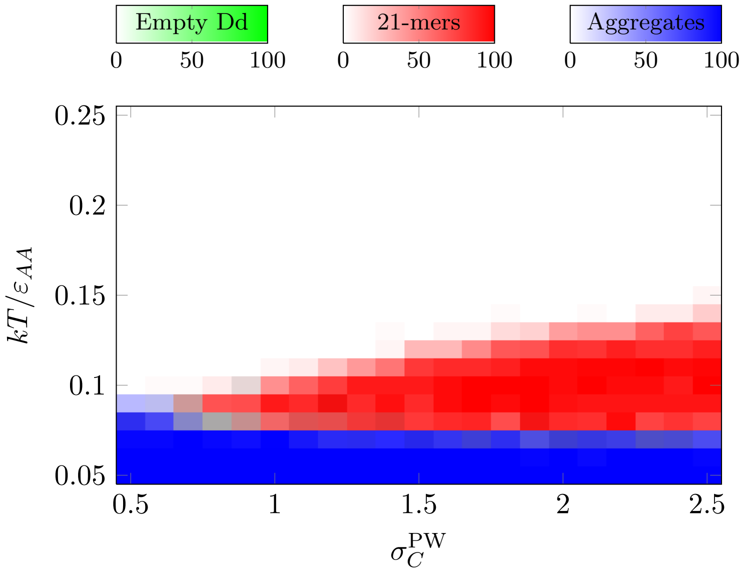

Figure 3.46: Plot of successful assembly of 21-mers in the dodecahedral system as a function of core-shell patch width at lower patch widths (other settings are given in table 3.1).

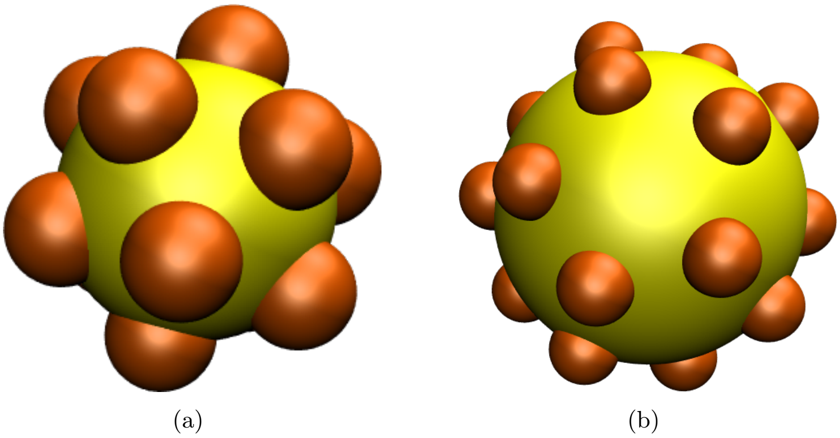
(a)
(b)

Figure 3.47: Core for (a) multiply-patched centred icosahedra-forming system (IhCM), (b) multiply-patched centred dodecahedra-forming system (DdCM).

surface diffusion of the shell particles on the encapsulating particle will become increasingly slow, thus perhaps limiting assembly. This effect should become important at different patch widths for the two different systems we examine here (as the patches will overlap less well for the icosahedral system than for the dodecahedral system at narrow patch widths). The effect on surface diffusion, however, may be offset by the fact that particles can only bind in the correct position for the target structure. The multiply-patched core particles can be seen in figure 3.47.

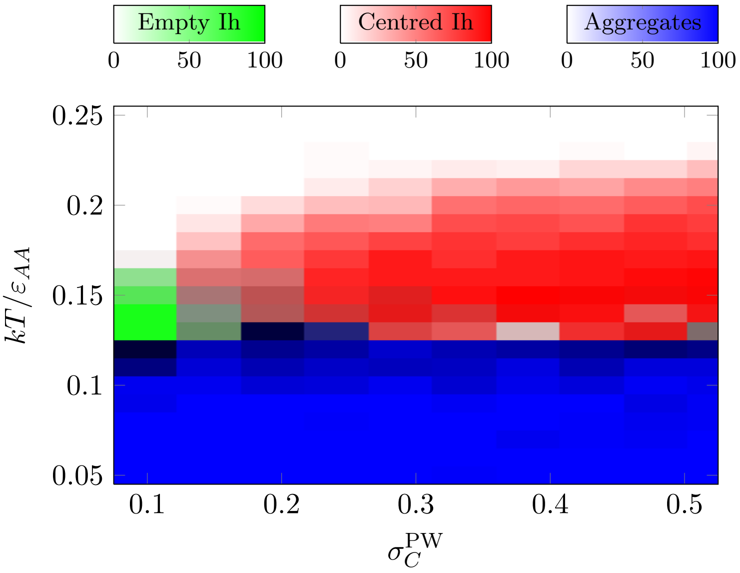

Figure 3.48: Plot of successful assembly as a function of core-shell patch width for the centred icosahedra-forming system with a multiply-patched core (other settings are given in table 3.1).

Icosahedral system

The use of a multiply-patched central particle was initially tested on the well-characterised icosahedral system. The final yields (figure 3.48) show little difference between the unstructured, isotropically attractive patched system at wide patch widths, and the specifically designed system. This means that less stringent designs are just as effective in the formation of icosahedral clusters. It is clear that below a patch width of approximately 0.3 radians, the yield is affected for the worse.

Dodecahedral system

We then tested the effect on the dodecahedral system. As this has proved to be effectively impossible to access in a one-component system without some form of torsional constraint, we felt it would give more interesting results. It is interesting to note that there was again little difference between the multiply- and singly-patched central particles (figure 3.49). This behaviour is somewhat different from what has been shown in the macroscopic case of forming a dodecahedron from magnetic components [109]. The macroscopic research showed that perfect templates formed dodecahedra (this is somewhat similar to a multiply-patched central particle), while the use of a globular template (unstructured, somewhat like our isotropic central particle) did not give a yield of the target cluster.

The effect of the patch width starts to affect the yields below 0.2 radians for the dodecahedral system, while in the icosahedral system, it becomes important at patch

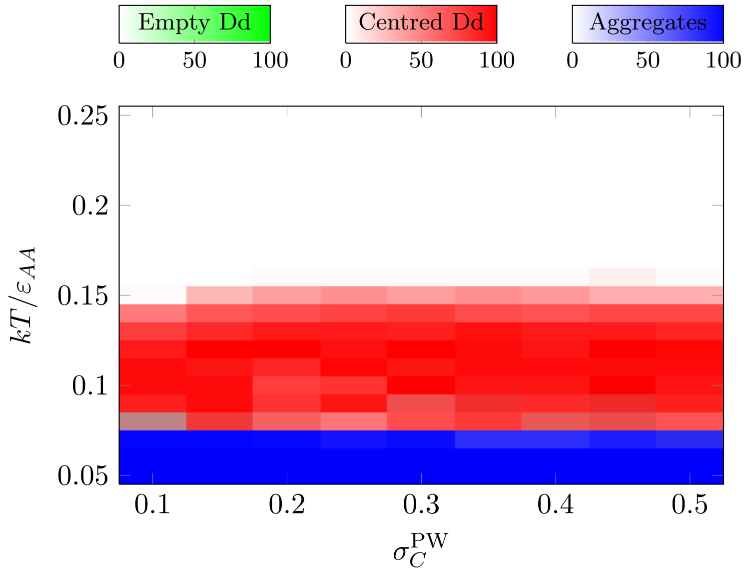

Figure 3.49: Plot of successful assembly as a function of core-shell patch widths for the centred dodecahedra-forming system with a multiply-patched core (other settings are given in table 3.1).

widths below 0.3 radians. This effect could be thermodynamic – the target cluster has lower vibrational entropy as the patches become narrower– or kinetic due to the slower surface diffusion. Interestingly, for the icosahedral system, it leads to uncentred icosahedra being preferred at the narrowest patches. However, this is not a feasible alternative for the dodecahedral system.

3.9 Kinetic stability of centred dodecahedra

As we are interested in the effective stability of completed dodecahedra clusters, we created a number of systems which had full yield of dodecahedra, and then ran the simulations for a million Monte Carlo cycles to see if they remained in that state.

We concentrated on a few variables to show the difference between kinetic accessibility and kinetic stability of the systems. In order to ensure that adequate statistics were obtained, we used eight different starting configurations of completely formed dodecahedra, and had three repeats for each configuration (24 runs per point in all). Configurations were chosen such that they would be kinetically stable for all initial patch widths used. When considering the effect of the size of the central particle, simulations were only performed for a central particle that was approximately equal to or smaller than the ideal, as it was not possible to create a system where both the shell and core particles were bonded correctly for larger cores.

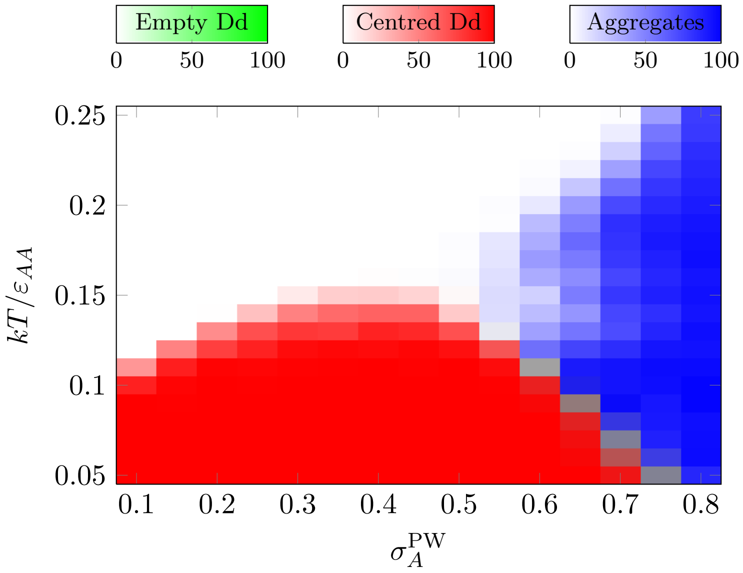

Figure 3.50: Stability of formed centred dodecahedra as a function of shell-shell patch width (other settings are given in table 3.1).

Patch width

Figure 3.50 shows the stability of formed dodecahedra when the patch width is varied. Compared to figure 3.24, there are no aggregates at low temperature and intermediate patch width, supporting our previous classification of them as kinetic traps. The liquid still forms at wide patch width, suggesting that it is indeed a thermodynamic limit, and clear phase re-entrance is demonstrated (see chapter 6).

Central particle size

We have previously shown that the thermodynamics of the dodecahedral system contains more complexity as a function of the size of the central particle than the icosahedral system, owing to the possibility of creating smaller centred clusters with some strain (section 3.6.1). Our yield plot in figure 3.51 (again only shown for clusters of size 21, all of which are dodecahedral in shape) supports the existence of a region of stability for the small clusters we found in the earlier section. The decrease in the upper temperature of reasonable kinetic stability of the formed centred dodecahedron as the size of the internal particle decreases can be explained in terms of the existence of the more stable smaller structures we have seen previously. At very low temperature, we see kinetic trapping into 21-mers, in the same region as we previously saw aggregates.

It is interesting to note that uncentred dodecahedra are observed at very low core particle sizes. This formation only occurs in a small region with a very small central particle, which is able to escape from the shell without interfering with the stability of the

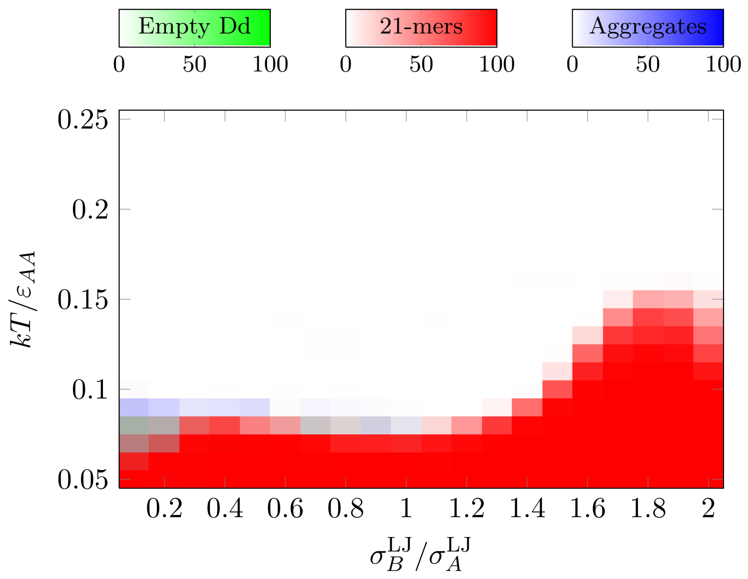

Figure 3.51: Stability of formed clusters with 21 members (other settings are given in table 3.1).

completed cluster. Previous research within the group [46] has already looked into the stability of the formed empty cluster alone.

Interaction strength

When varying the interaction strength, we have shown that the kinetic yield is strongly dependent on an attraction between the core and shell particles. As figure 3.52 shows, the kinetic stability of the centred dodecahedra depends again on this attraction. At the point of there being no attraction, the core particle still remains within the dodecahedron; even though it is not bonded, it is kinetically stable. This stability is encouraging for the use of a dissolving core: once the template particle is removed, the structure remains stable over a certain range, despite it being thermodynamically unstable with respect to aggregation.

3.10 Encapsulation conclusions

In our simulations, we are straightforwardly able to increase the number of particle types and types of patches, and to introduce specificity into the patch-patch interactions. By contrast, even though the synthetic strategies for producing patchy particles are rapidly improving, some of the particles whose behaviour we have analysed here would be very challenging to synthesise, particularly in terms of the control of patch position and identity. DNA-mediated interactions provide a potential route to achieve the required specificity in the patch-patch interactions [199]. However, our results also indicate that the use of templated self-assembly provides a means to increase greatly the repertoire of structures

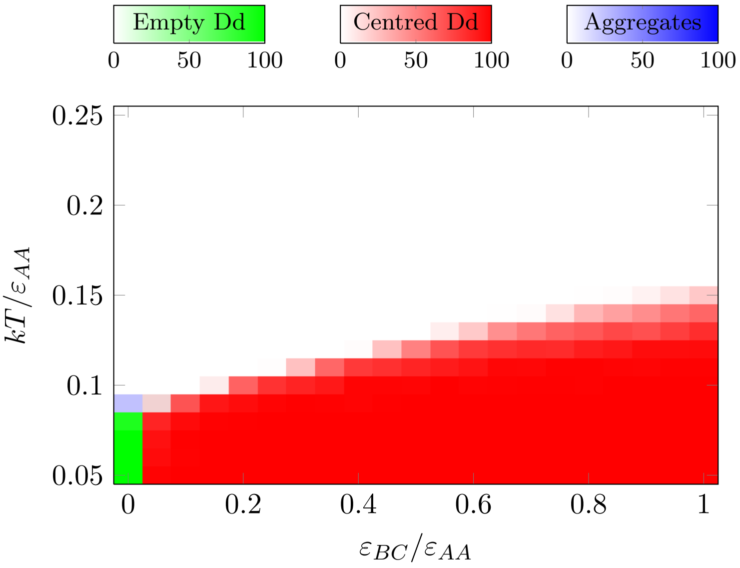

Figure 3.52: Stability of formed centred dodecahedra as a function of patch strength (other settings are given in table 3.1). Note that the *empty dodecahedra* found here do contain the central particle still, but it is not bonded to any of the shell particles.

into which such particles could assemble. We have highlighted several advantages that are attainable using encapsulation:

Kinetic traps: Aggregation and other traps affect encapsulated systems less than their uncentred counterparts, and this allows the formation of otherwise inaccessible clusters. In particular, tuning the strength of the interaction with the template creates conditions where aggregation does not compete with assembly.

Robustness: Encapsulated systems are more robust to noise in positioning of shell patches when compared to their uncentred counterparts, and are also able to form when the central particle is not isotropically attractive.

Simple templating: Simple isotropic templates are sufficient within our model. More complicated templates with specific interactions did not lead to a better yield than globular isotropic templates. Thus, particle design is easier than might intuitively be expected.

This chapter is similar in content to the following article: A. J. Williamson, A. W. Wilber, J. P. K. Doye, and A. A. Louis. Templated self-assembly of patchy particles. 2011 Soft Matter 7 pp 3423-3431. Reproduced by Permission of The Royal Society of Chemistry (RSC).

Chapter 4

Hierarchical self-assembly

4.1 Introduction

Hierarchical assembly is an assembly pathway in which the final cluster is built from a non-monomeric intermediate. Hierarchical assembly is ubiquitous in biological systems [200]. For example, the formation of the turnip crinkle virus follows a pathway of monomers forming dimers forming hexamers, which are bonded to the genetic material of the virus [93]. Whether this biological prevalence is due to the evolutionary pathways common to protein complexes [201] or because of assembly efficiency is unclear. In this chapter, we look at the utility of hierarchical pathways in ensuring monodisperse assembly.

Hierarchical assembly is also an analogue to multiply-staged chemical synthesis, and has many similar issues and limits (see section 4.1.4).

4.1.1 Requirements for efficient assembly

When designing a self-assembling system, it should be noted that the difference in free energies between the target and other competing structures should be maximised. This can be achieved by increasing the free energy of the competing structures or by decreasing the free energy of the target structure. These thermodynamic strategies can also be combined with kinetic ones, encouraging the preferred path of assembly by the use of fast initial steps. In designing hierarchical assembly, finding the means of speeding up these initial steps compared to other pathways is one of the elements crucial for success.

There are many kinetic and thermodynamic traps that must be avoided in order to achieve efficient self-assembly. In the simulation of non-torsional patchy particles, we have found that the most limiting of these is the kinetic formation of large aggregates that behave like a glass. This occurs below the (kinetic) aggregation temperature (T_{agg}), while clustering occurs below the (thermodynamic) clustering temperature (T_{clus}). The

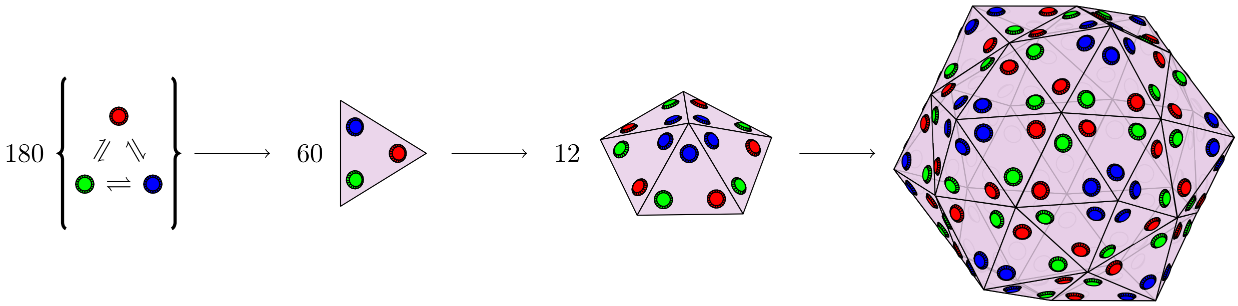

Figure 4.1: A schematic of the assembly of a polio virus. Mechanism based on research by Hogle [97].

region in parameter space with the clustering temperature at the upper bound and the aggregation temperature at the lower bound delimits the window of possible successful cluster formation. In the event that the aggregation temperature is above the clustering temperature, it is very unlikely that clusters can ever be formed, as may be the case with certain target structures built in an all-at-once fashion [112].

4.1.2 Experimental research

The classic example of self-assembly is that of a virus capsid (see figure 4.1 and chapter 1), which is able to form in a wide range of biologically relevant conditions [202], often doing so in a hierarchical manner with very high yields [203]. This is by no means the only biologically relevant hierarchically self-assembling system: in fact, the human body is built of many complicated structures that are formed by similarly hierarchical pathways [200].

DNA-based nano-structures are another area of active research in the scientific community [204, 205]. Such hierarchically self-assembling systems have great potential for use in programmable nano-technology [206].

In our investigation into the utility of hierarchical assembly, we looked at the effect of the choice of intermediates and the design of particles. Our aim was to find a system that was capable of robustly building a certain cluster across a wide range of conditions.

4.1.3 Simulation-based research

Given that the area of self-assembling synthetic particles is still in its infancy, much of the current research into the dynamics and methods of assembly is unsurprisingly simulation- and theory-based [4, 207, 208].

Our previous work

Previous work within the group [4] has shown that it is easier for a system to form in the presence of a torsional constraint and that, in the absence of these constraints, a simple all-at-once assembly method does not give reasonable yield of certain target structures [112].

In non-torsionally constrained systems, the free energy of the aggregate is lower than for those with torsional constraints. A non-torsionally constrained system also has a smaller barrier to aggregate formation. Thus, the formation of aggregates is thermodynamically and kinetically more favourable in the absence of torsions. This factor strongly influences which shapes and sizes of cluster can be formed by patchy particles with non-torsionally constrained interactions. We examine the effect of this thoroughly in this chapter.

Without a torsional constraint, recent simulations have shown that the use of a central particle encapsulating the surrounding allows a high yield of certain otherwise inaccessible cluster shapes (see chapter 3), as it decreases the free energy of the target structure. We have also shown the utility of building up more complicated structures by encapsulation of smaller structures inside larger ones [46].

Other work

Other researchers have looked at the utility of self-assembly in forming monodisperse target structures [13, 208], often using biology as a guide and an inspiration. DNA self-assembly and simulation is also well documented [209, 210].

4.1.4 Is hierarchical assembly the answer?

Although it is intuitively reasonable that a hierarchical assembling system should give the ability to assemble things that otherwise would not assemble, whether they can do so in a quick manner with a high yield of product is not so obvious. We have identified a series of potential advantages and disadvantages of using a hierarchical assembly route, as detailed below.

Potential advantages

Control Fine-tuning the yields and rates of simple small steps should permit full control of the overall process. As the control of simple steps is much easier to deal with than attempting to balance many factors concurrently, as would be required in an all-at-once assembly mechanism, this should lead to hierarchically assembling systems being more readily controlled. Notwithstanding that designing nanoparticles which need to undergo

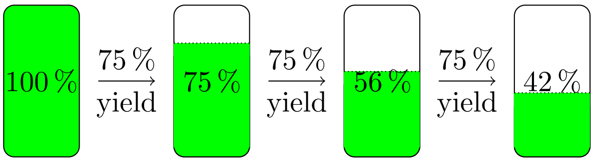

Figure 4.2: A schematic of the limiting nature of attrition. Following only three steps with 75 % yield, there is only an overall yield of 42 %.

a series of small steps may add some complexity to their initial manufacture, it is not expected that this will be an insurmountable barrier to manufacturing.

Ease of pathway selection By designing particles to form the desired intermediates quickly, it should be possible to disfavour certain structures kinetically, by ensuring that undesirable pathways are slower.

Potential disadvantages

Attrition The main disadvantage of any multi-stage process is the fact that the losses at each stage are magnified. This is known as *attrition* and is exemplified in figure 4.2 for a simple three-stage process with 75 % yield at each step.

Blocking In addition to the effect of the well-characterised effect of attrition in chemical synthesis, there are additional issues in nano-particle formation. This is not only because losses decrease the maximum possible yields, but also because incorrectly formed structures often actually inhibit the ability of the correctly formed intermediate structures to join together to form the final structure. If an incorrectly formed structure irreversibly bonds with a correctly formed structure, neither of them will be able to produce the desired product. We term this *blocking*.

Assembly competence In order to minimise attrition, it will be necessary to ensure that systems are able to form mainly the correct intermediate. This constraint leads to the idea of *assembly competence*, best illustrated by example.

A non-assembly competent intermediate is able to form stable structures which are unable to bond together to form the final structure. For example, the two pentagons in figure 4.3 demonstrate how two pentagonal intermediates can be produced by different interaction matrices. The first has a cyclical bond structure (patch A (blue) can only bond to patch B (green)). The remaining patch in this example (patch C) is coloured yellow. The interactions are such that a completed pentagon of AB bonds can only form in one isomer, with all the type-C patches pointing upwards. This is an assembly competent

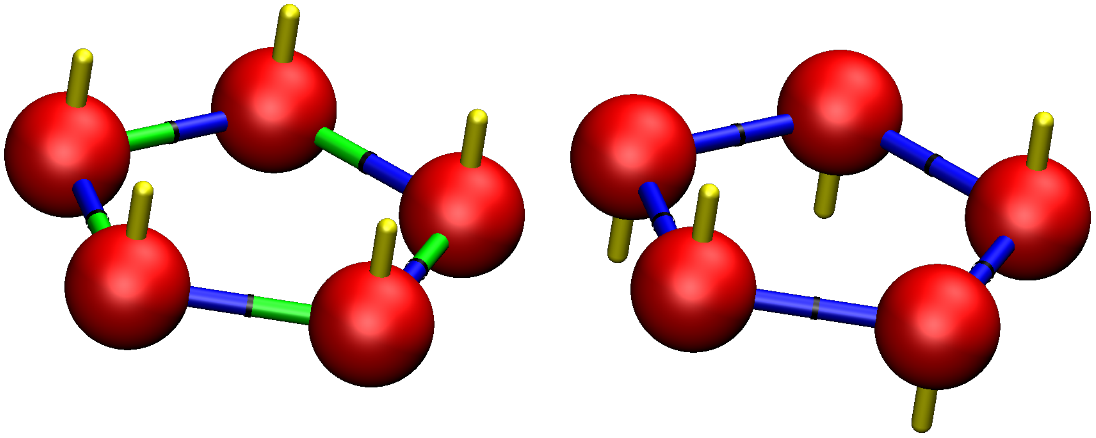

Figure 4.3: An example of assembly competent (left) and non-assembly competent (right) intermediates. The assembly competent intermediate is only able to form a loop with all of the yellow patches pointing in the same direction, unlike the non-assembly competent intermediate, which may form in many different ways.

intermediate. On the other hand, the structure on the right of figure 4.3 shows how a naïve attempt to create a hierarchical assembly method can lead to difficulties. The monomers have three patches: two of type-A (blue) and one of type-C (yellow). The interactions are tuned such that patch A can only bond to other patches of type-A. This means that many different isomeric completed pentagonal intermediates may be formed, with the type-C patches pointing either up or down, and only one is in fact the useful isomer (with a 1 in 16 chance of this version being assembled). Any misformed intermediates in the non-assembly competent system could cause blocking and prevent the final structure from being formed in a high yield.

This effect of assembly competence leads to a requirement that interactions between monomer building blocks may have to be specified more precisely than in an all-in-one assembling system. Incorrectly bonded particles in nucleated assembly pathways are able to dissociate and then reassociate correctly. In hierarchical assembly, this is harder, as the intermediate is more stable with respect to monomers.

Possible purification steps While we have looked mainly at one-pot type assembly methods, it is possible for the impact of the attrition to be decreased by using purification steps in a multi-step synthesis. We look at the effect of such a purification step by starting our simulations from a box of fully formed intermediates. Naturally, a real-world purification step, such as centrifugation, would require energy and is thus a disadvantage in practice.

4.1.5 Simplified thermodynamics

As shown in section 3.5.3, we are able to estimate the thermodynamic equilibrium points using simple thermodynamics. These approximations allow us to plot a schematic thermodynamic plot of transition temperature against any of the different bond energies of the relevant system, giving a feel for the stability regions of different structures and allowing us to estimate the formation pathways of these structures. We can reasonably assume that in areas of the phase diagram where the stability of a single type of small intermediate is greater than that of monomers, a hierarchical-type assembly pathway will be followed, while in regions where more than one intermediate is more stable, we are likely to see trapping and a downhill reaction pathway. We test this theory in the following section.

This simple thermodynamic approach is able to estimate the relative temperatures of different transitions, but not their absolute values. Therefore, we use a transition temperature obtained from simulation to fix the temperature scale. The point we take is usually from a simulation where the different interactions are of equal strength where there is an equal probability of a randomly chosen particle being in our target cluster or being an unassociated monomer. We can then compare our simplified approximation and actual simulation yields.

4.1.6 Method

We looked at the yield obtained from reasonable length simulations (10^6 Monte Carlo cycles), both with and without the use of torsional constraints in the potential. Initially, we looked at the effect of changing the patch width when all interactions were the same strength to find out where the best yields were found, and we then used our schematic thermodynamic plots in deciding which patch strengths to vary in order to test whether hierarchically formed structures are easier to form than their non-hierarchical counterparts.

Unless otherwise stated, the values in table 4.1 were used for the simulations.

4.1.7 Yield plots

We again use an adaptation of the yield plots designed by Villar [178]. Red is used for the target cluster, blue for aggregates (defined as clusters larger than twice the size of the target cluster), green shows yield of triangles, cyan shows squares, and magenta is used for dimers. We then plot our expected transition lines onto the diagram. Equilibrium transitions are shown as solid lines, while those which show the limits of stabilities of small non-monomeric clusters with respect to monomers are plotted as dotted black lines. These should correspond approximately to changes in mechanism.

Variable	**Value**
ε	All equal
σ^{PW}	0.35 radians (icosahedral systems) *or* 0.45 radians (other systems)
σ^{LJ}	All equal
$\frac{N}{V}$	0.15 $(\sigma^{\mathrm{LJ}})^{-3}$
$\frac{1}{\sigma^{\mathrm{Tors}}}$	0 (without torsional constraints) *or* 1 (with torsional constraints)
Repeats for dynamic runs	5
$\lvert \boldsymbol{r}^{\mathrm{cutoff}} \rvert$	3.5 σ_A^{LJ} (dynamic) *or* 2 σ_A^{LJ} (Wang-Landau)
Number of Monte Carlo cycles	10^6
Maximum number of target clusters possible	10 per run (icosahedra and 24-sized clusters) *or* 20 per run (other 12-sized clusters)

Table 4.1: Parameters used in the simulations of the hierarchical systems.

Settings for dynamic runs are in table 4.1. Clusters are classified as formed when they have the correct number of particles and up to a maximum of 2 unformed bonds in the case of the icosahedral systems, and one in the case of the other clusters. A bond is defined as formed if the energy of interaction between two particles is at least 40 % of the well depth.

4.1.8 Images

In all pictures from simulations, the colour scheme for the systems of interest is for the particles to be shown as red and the locations of the patches are shown in blue (type A), green (type B), yellow (type C), cyan (type D) and magenta (type E). As in the previous chapter, although the patches do not have volume, they are schematically represented as having such for clarity.

4.2 Results

We looked at the formation of a series of different target geometries and the effect that hierarchical assembly had on the efficiency of self-assembly. We then used the results to assess the efficiency of hierarchical assembly and to find any generic rules for hierarchical assembly.

The cluster shapes that we investigated all had identical vertices and could be thought of as being constructed from regular polygons. We considered two different sizes of clusters: those formed from 12 patchy particles and those formed from 24 patchy particles.

4.2.1 Twelve-particle systems

Initially, we looked at clusters that were formed from 12 particles. In order to test hierarchical assembly, we looked at the yields and formation mechanisms of three clusters with differing numbers of patches per particle. They were the truncated tetrahedron (3 patches per particle), the cuboctahedron (4 patches per particle) and the icosahedron (5 patches per particle). The results are detailed below, commencing with the cluster with the most bonds per monomer.

Icosahedron

Much of our previous work concentrated on the formation and behaviour of self-assembling icosahedral clusters [6, 112] Here, we consider a particle design that could potentially allow icosahedra to form hierarchically from four triangles. The particles are designed so that the triangles will be assembly competent. We term this system *IWC* and compare its behaviour to other self-assembling systems which form icosahedral clusters.

IWC Our assembly competent icosahedron-forming single component hierarchical system consisted of a patchy particle with 5 patches of 5 different types. The interactions were configured such that patch A (blue) could only interact with patch B (green), patch C (yellow) with patch D (cyan) and patch E (magenta) could only interact with other E-type patches. The completed cluster can be seen in figure 4.4.

We looked at the effect of changing the values of the interactions ε_{CD} and ε_{EE} with respect to the remaining patch pair ε_{AB} (see figure 4.5). AB interactions are responsible for the formation of triangles, while CD and EE patch interactions form the bonds between them. If $\varepsilon_{AB} > \varepsilon_{CD}$ and $\varepsilon_{AB} > \varepsilon_{EE}$, then there is the possibility of a region in which hierarchical assembly may occur via an AB triangle intermediate. Equivalently, if $\varepsilon_{CD} > \varepsilon_{AB}$ and $\varepsilon_{CD} > \varepsilon_{EE}$, then a region in which hierarchical assembly may occur via

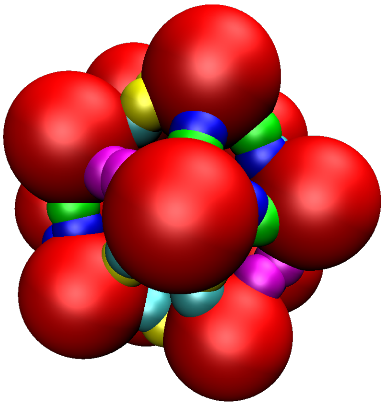

Figure 4.4: The IWC-type icosahedron.

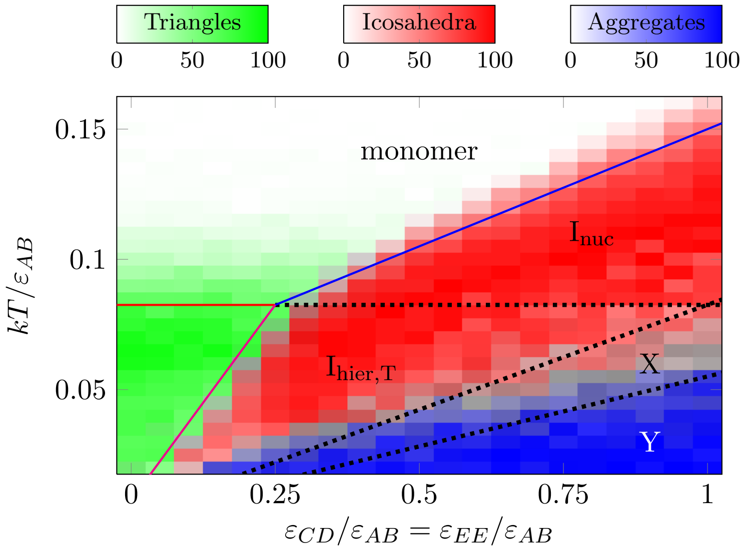

Figure 4.5: IWC yield plot as a function of the interactions which join AB-triangles together. $\mathrm{I}_{\mathrm{hier,T}}$ refers to the region in which icosahedra are expected to form hierarchically via triangles, and $\mathrm{I}_{\mathrm{nuc}}$ is the region where a nucleation type mechanism is expected. X is the region where CD or AB type triangles are both stable intermediates, and Y is the region where the formation of any bond is downhill in free energy. Settings not specified in the graph are given in table 4.1.

a CD triangle intermediate would be expected. Our proposed mechanism for successful hierarchical assembly of the IWC system is

$$12\,\mathrm{A} \rightleftharpoons 4\,\mathrm{A}_3 \rightleftharpoons \mathrm{A}_{12}. \tag{4.1}$$

For the point at which $\varepsilon_{AB} = \varepsilon_{CD} = \varepsilon_{EE}$, we have behaviour similar to our archetypal icosahedron, where decreasing temperature changes the kinetic phase from monomers to icosahedra to aggregates. As ε_{CD} and ε_{EE} decrease, the clustering temperature T_{clus} and aggregation temperature T_{agg} decrease due to lower energetic stabilisation. At sufficiently low values of ε_{CD} and ε_{EE}, AB triangles become stable with respect to monomers at a higher temperature than icosahedra. The regions where icosahedra are formed are designated as $\mathrm{I}_{\mathrm{hier,T}}$, and $\mathrm{I}_{\mathrm{nuc}}$. In the region labelled $\mathrm{I}_{\mathrm{nuc}}$, the intermediates are not stable with respect to monomers, so we expect a nucleation type mechanism of assembly. Conversely, in the region marked $\mathrm{I}_{\mathrm{hier,T}}$, AB triangles are stable with respect to monomers, but not with respect to the formation of icosahedra. Therefore, we expect AB triangles to form as an intermediate to assembly. In the region marked X, both AB and CD triangles are stable with respect to monomers, but they are unable to co-assemble successfully into an icosahedron. Assembly via 4 AB triangles, or via 4 CD triangles is possible, but a mixture of triangles may block correct assembly. In the region marked Y, dimers and triangles are stable with respect to monomers and any bond formation is downhill in free energy. This leads to blocking by the formation of mixed systems.

The yield plot (figure 4.5) clearly shows good assembly in the regions of nucleation and hierarchical assembly from one component. However, in regions where more than one intermediate is more stable than monomers, we see a decrease in yield, and eventual formation of aggregates. This is because more than one pathway is initially possible and so we get blocking of assembly.

Hierarchical and non-hierarchical assembling regions The designation of areas in figure 4.5 as hierarchical or non-hierarchical has been based on simple thermodynamic arguments about the stability of intermediates. To confirm the validity of our approximations, we show the time evolution of the numbers of monomeric patchy particles as well as the amount of triangular intermediate and completed clusters that are formed for parameter values where we expect the dynamics to be hierarchical (figure 4.6(a)) and non-hierarchical (figure 4.6(b)) on the basis of figure 4.5.

The graphs in figure 4.6 clearly support the classification of figure 4.5 into hierarchical and non-hierarchical assembly. Both were taken from conditions of similar overall yield. In figure 4.6(a), the two-step nature of the assembly process is clear, with monomers forming trimers and then trimers assembling into icosahedra. The peak in the yield of triangles at 10^4 Monte Carlo cycles in figure 4.6(a) shows its prevalence as an intermediate, while this peak is missing from the non-hierarchical regime's pathway (figure 4.6(b)), and trimers are only one of a series of intermediates that are formed in the pathway to the icosahedra and have no special stability compared to other intermediates. It is interesting to note that the number of cycles taken to reach a stable maximum yield was an order of magnitude higher for the hierarchical assembly region compared to the non-hierarchical region, suggesting that hierarchical assembly regimes may not be the fastest assembly mechanisms overall. This difference is perhaps not surprising, given that the hierarchical assembly involves two

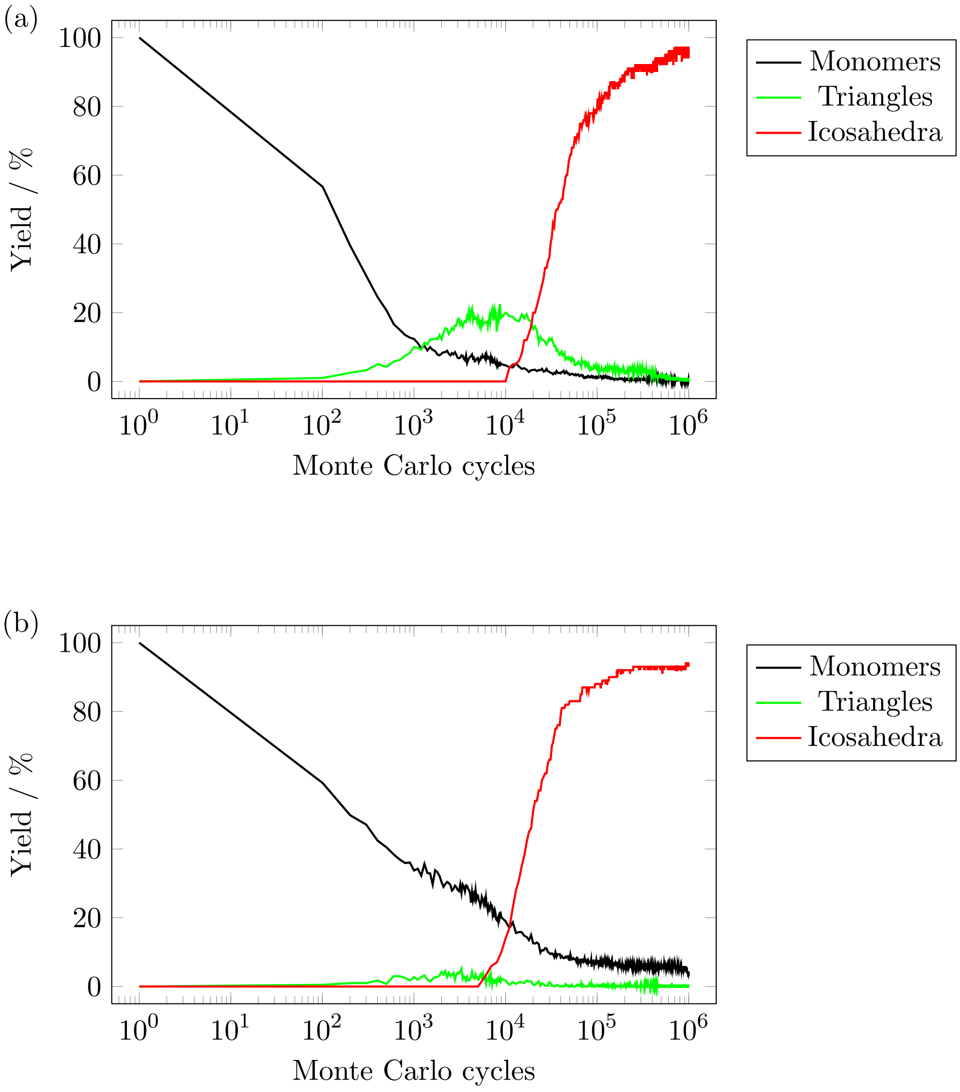

Figure 4.6: Time resolved yields of a run following (a) the hierarchical assembly route ($kT = 0.075\,\varepsilon$, $\varepsilon_{CD}/\varepsilon_{AB} = \varepsilon_{EE}/\varepsilon_{AB} = 0.4$) and one following (b) the non-hierarchical assembly route ($kT = 0.105\,\varepsilon$, $\varepsilon_{CD}/\varepsilon_{AB} = \varepsilon_{EE}/\varepsilon_{AB} = 0.75$) for the IWC system. Settings not specified are given in table 4.1.

steps and that the diffusion of trimers (as must occur in the second step) is slower than that of monomers. The graphs are typical for other points in the relevant regions.

IhO versus IWC We compared this IWC system to an icosahedron-forming system made of particles with 5 equivalent patches (which we term *IhO*) in order to see the overall effect of monomers capable of assembling hierarchically compared to clusters which were not. Each of the patches on one monomer in the IhO system was able to bind to any other single patch on another monomer. In figure 4.7, we compare the yield as a function of patch widths and temperatures for these two systems. It is noteworthy that the region of successful assembly for the IWC system (figure 4.5) occupies a wider range

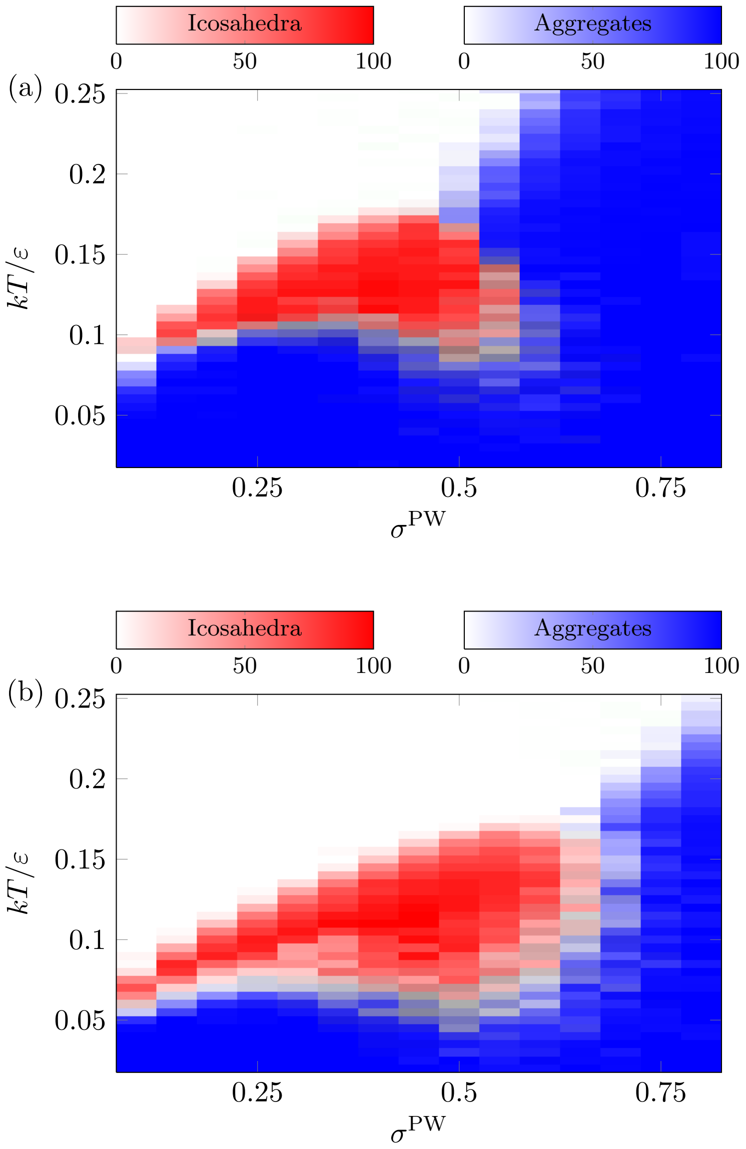

Figure 4.7: Comparison of yields of (a) IhO and (b) IWC as a function of patch width. Settings not specified in the graph are given in table 4.1.

of parameters with assembly enhanced at lower temperature and wider patch width, and is mainly due to a suppression of aggregation. Aggregation results from the non-torsional nature of the interactions that allows particles to come together to form structures that are not on the pathway to icosahedra. In the IWC system, the greater specificity of the interactions, in particular that they can only form assembly competent triangles, reduces the number of pathways available to aggregation. This clearly shows that a system able to form hierarchically has significant advantages over its non-hierarchical analogue. However, it should be noted that the regions of hierarchical and non-hierarchical assembly for the IWC system do not show a marked difference in overall yields.

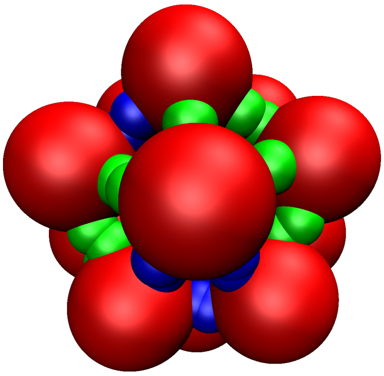

Figure 4.8: An icosahedron made from IW-type particles.

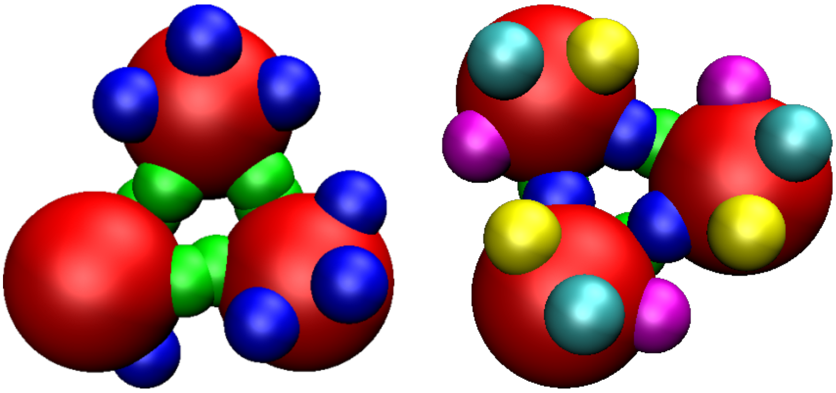

Figure 4.9: The IW intermediate (left) is not assembly competent, as it can form non-useful isomers which block assembly that have the same free energy as the correctly formed intermediate. The IWC intermediate (right), on the other hand, may only form one isomer, and thus it is assembly competent.

IWC versus IW: The importance of assembly competence In order to explore the importance of assembly competent intermediates (section 4.1.4), we compared the IWC system with a non-assembly competent analogue, the *IW* system. This had monomers with two patch types ($2A$[blue], $3B$[green], in the cyclical order $AABBB$), with each patch only able to bond to the same type (*i.e.* A to A and B to B). The target cluster is shown in figure 4.8. The system can form triangular intermediates held together by AA interactions; however, there are several isomers possible for such triangles. Those with the remaining B patches not all on the same side of the triangle (see the example in figure 4.9) have a structure that is incompatible with further assembly into an icosahedron. These isomers would be expected to cause some blocking and prevent correct cluster formation in the hierarchical regime. The intermediate is not assembly competent.

Figure 4.10 shows the yields of triangles, icosahedra and aggregates as a function of temperature and relative patch strengths and should be compared with figure 4.5. Figure 4.10 shows that the IW system does not have much successful assembly in the region where hierarchical assembly might be expected. This behaviour occurs because this

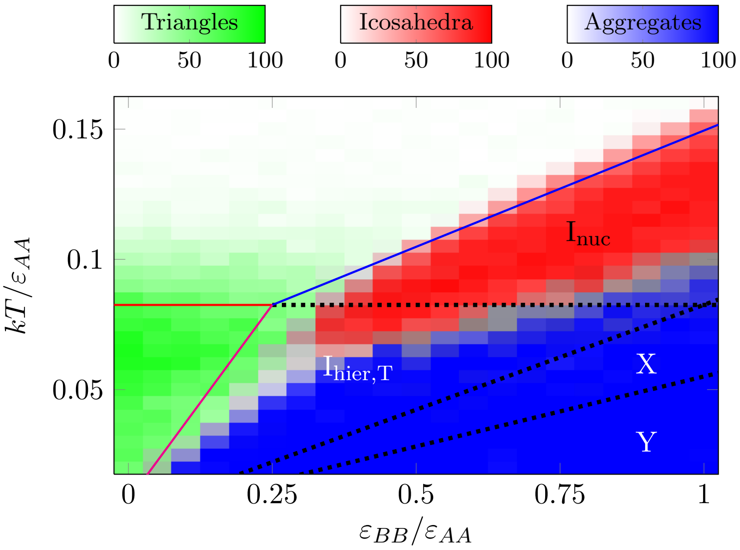

Figure 4.10: IW yield plot at patch width $\sigma^{PW} = 0.35$. Labels as in figure 4.5. Settings not specified in the graph are given in table 4.1.

is the region where non-assembly competent triangles are stable with respect to monomers (one would estimate that 3/4 of triangles formed are not assembly competent), and so any that form are unlikely to break-up in this region, as would be necessary if those particles were to become part of icosahedra. By contrast, in the region where we expect nucleation to be the main pathway, and no intermediate is stable with respect to monomers, any non-assembly competent clusters will only exist transiently and so cannot block successful assembly. Thus, we can conclude that designing the interactions so that intermediates are assembly competent is vital if hierarchical assembly is to occur.

Cuboctahedron

A cuboctahedron is an Archimedean solid and can be formed from patchy particles with four bonds per particle (figure 4.11). There are two allowed interactions in our implementation: those between A and B patches and those between C and D patches. We look at two possible constituent clusters: AB triangles and CD triangles (squares formed from alternating AB and CD bonds are also possible, but generally are less stable than the triangles, which would be expected to form faster). AB triangles are disfavoured by decreasing the strength of AB interactions with respect to CD interactions. Thus, a hierarchical assembly pathway via CD triangles is expected to become feasible on such a decrease of AB interaction strength.

Figure 4.12 shows the simulation results and a simple estimate of the assembly pathways (see section 4.1.5). The window between aggregation and cluster formation, in which the self-assembly of cuboctahedral clusters was successful, was not as wide as the IWC system.

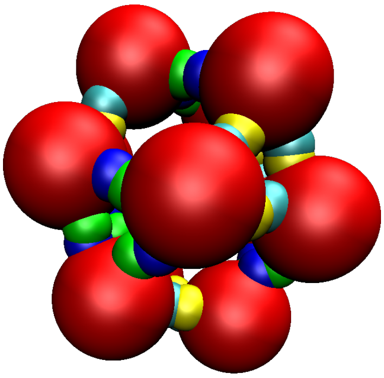

Figure 4.11: A cuboctahedron.

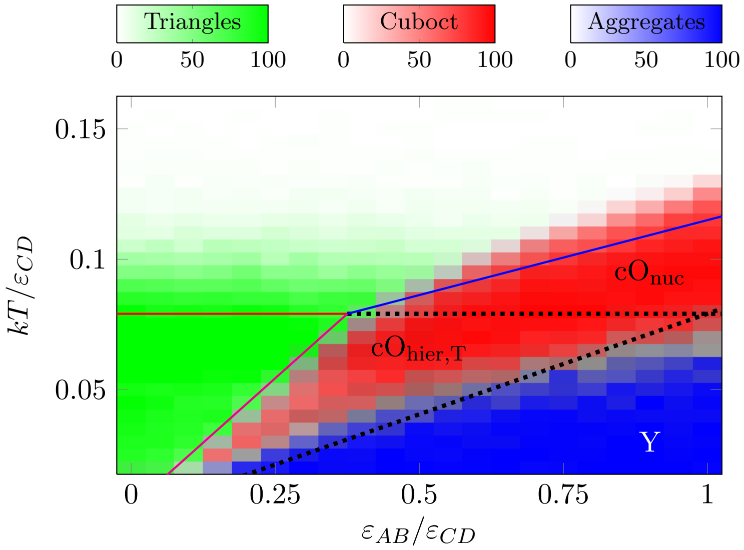

Figure 4.12: Cuboctahedron yield plot. $cO_{hier,T}$ refers to the region in which icosahedra are expected to form hierarchically via CD-triangles, and cO_{nuc} is the region where a nucleation type mechanism is expected. Y is the region where AB triangles and CD triangles are more stable than the monomer. Settings not specified in the graph are given in table 4.1.

This decrease in the window of successful formation is due to the decrease in the number of patches, and thus the specificity of the particles. Time-resolved simulations show that, within the hierarchical region, we still have successful assembly via an intermediate.

Truncated tetrahedron

The truncated tetrahedron is an Archimedean solid which can be formed from a patchy particle with three patches of three different types (figure 4.13). The interactions were configured such that A-type patches could only interact with B-type patches and C-type patches could only interact with other C-patches. By varying the ratio of the strength of the CC interactions compared to the AB interactions ($\varepsilon_{CC}/\varepsilon_{AB}$) between zero and one, it is possible to encourage hierarchical assembly via triangular intermediates (varying the

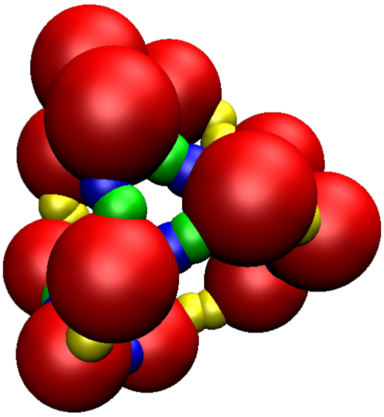

Figure 4.13: A truncated tetrahedron.

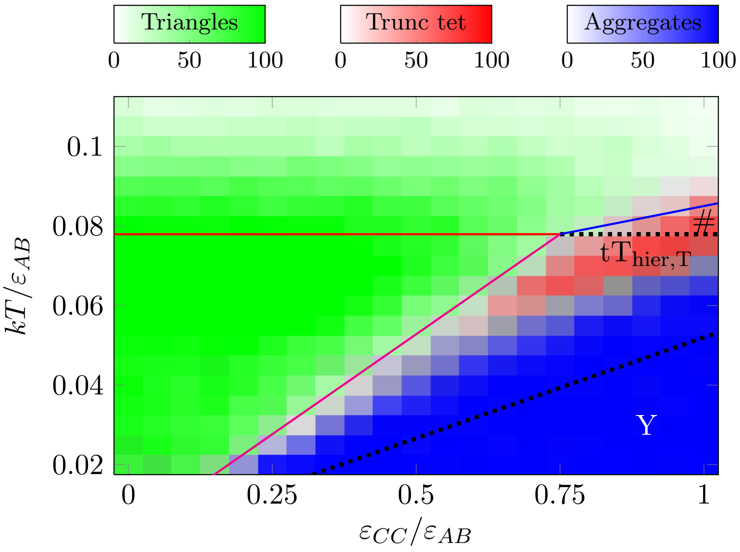

Figure 4.14: Truncated tetrahedron yield plot. The region labelled $\mathrm{tT}_{\mathrm{hier,T}}$ is the expected hierarchical region, # shows the nucleation region. Y is the region where both dimeric and triangular intermediates are more stable than monomers. Settings not specified in the graph are given in table 4.1.

inverse ratio between zero and one encourages dimers as an intermediate).

Figure 4.14 shows the results of simulations where the ratio of patch strengths was varied. Little successful assembly is seen outside of the hierarchical region. The formation of aggregates occurs across a wide range of the explored region. This prevalence of kinetic aggregates is due to the low number of bonds per particle. The window of assembly is lower than the cuboctahedron-forming system, and much lower than the IWC-forming system. The temperature at which the clusters form is strongly affected by the number of bonds per particle (clusters are mainly driven by the low internal energy of the target structure and are not entropically favoured), while the kinetic aggregation temperature is less strongly affected by the number of patches per particle. The number of possible configurations of non-overlapping particles that satisfy the energetic requirement for the patches to be bonded decreases rapidly as we increase the number of patches that must be bonded. If we assume that the intermediate forms faster than an aggregate, we can treat it like a patchy particle, and the strong dependence of the window of formation on the number of monomer patches becomes amplified. The intermediate of the truncated tetrahedra has three unbonded patches, while the cuboctahedron's intermediate has six, and the IWC-forming system's intermediate has nine. Thus the number of possible configurations decreases even faster than it would do for a system unable to form a stable hierarchical intermediate.

Comparison with a non-hierarchical system When the interaction table of the truncated tetrahedron is set such that all patches are of type A and are attracted equally to one another (the non-hierarchical assembling version of the system), then the yield is very low (less than 2 % correctly formed structures under optimal conditions), and aggregation is more common. While this does not definitively show the utility of hierarchical assembly, it does show that particles designed to assemble hierarchically into clusters of designed size and shape can in some instances be more successful than those which are not.

4.2.2 Twenty-four-particle systems

We used the 12-particle cluster forming systems to test the basic principles of hierarchical self-assembly. We were then interested in exploring the potential for hierarchical self-assembly to extend the range of structures that can form. We looked at a series of systems which could stably form 24-particle Archimedean clusters. In order of increasing bonds per particle, these are the truncated octahedron (3 patches per monomer), the rhombicuboctahedron (4 patches per monomer) and the snub cube (5 patches per monomer).

Truncated octahedron

Our particles which have the potential to assemble hierarchically to form a truncated octahedron possess three different patch types (A, B and C). The interactions are chosen such that patch A is attracted to patch B and patch C to patch C (see figure 4.15). Two hierarchical pathways can be envisaged. Firstly, the system could form squares (stabilised by AB interactions), six of which could then join together to form a completed truncated octahedron. Secondly, a completed cluster may form hierarchically with 12 dimers as intermediates (stabilised by CC interactions).

The truncated octahedron cluster represents a challenging target to form. In the absence of torsional constraints in the potential, truncated octahedra do not form. This

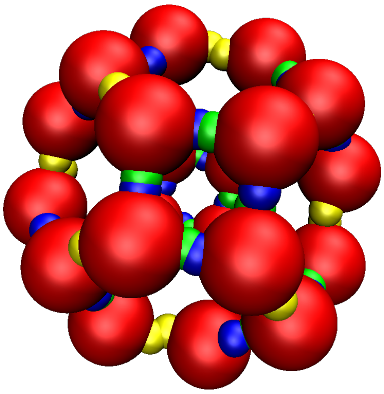

Figure 4.15: A truncated octahedron.

result is a similar to our findings for the dodecahedron [112], and is the result of the low number of bonds per particle and wide angles between the patches. With only three patches per particle, the clustering temperature in the absence of torsions lies below the kinetic aggregation temperature, and thus there is no temperature window for which the truncated octahedron will form from monomers. As aggregates form more rapidly, truncated octahedra never form with an appreciable yield, even in the regions where they are thermodynamically more stable.

With torsional constraints In the presence of torsional bonding constraints, it is possible to form truncated octahedra by both hierarchical and non-hierarchical assembly methods. The torsional constraints repress the formation of large aggregates by increasing their free energy. Figure 4.16 shows the results for the torsionally constrained bonding case. Note the suppression of aggregates, even in the low temperature regions where they are present for our other non-torsional simulations, but that in regions where more than one pathway is possible, we see trapping of the system into small clusters which are unable to interact.

Without torsional constraints The yield of truncated octahedra attempting to form hierarchically, in the absence of torsional constraints, was very low (under 2 %). Therefore, it was not possible to compile a yield plot. The reason for the low yield could be down to two possible limiting factors:

1. The rate of formation of productive intermediates is slow compared to other pathways that form alternative non-productive structures.
2. Assembly from squares into truncated octahedra is slow, and trapping into non-productive structures is common at this stage.

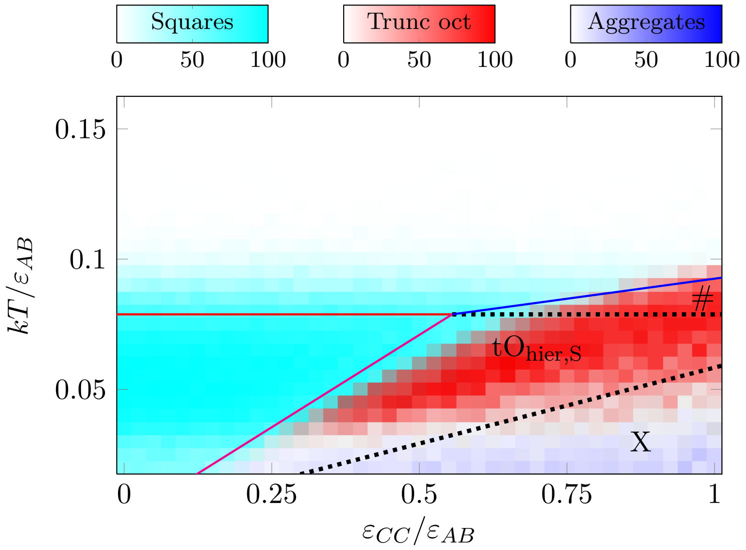

Figure 4.16: Yield of truncated octahedra with torsional constraints. # labels the expected region of nucleation. $\mathrm{tO}_{\mathrm{hier,S}}$ labels the estimated region of hierarchical assembly via squares. X denotes the region where both dimers and squares are more stable than monomers. Settings not specified in the graph are given in table 4.1.

In order to test which of these factors was the limiting one, we ran a series of simulations starting with a system of fully formed squares, along with simulations testing the formation of squares only (see section 4.2.3). The yield of truncated octahedra for the systems starting with fully formed intermediates was still low, but improved when compared to the systems starting from monomers (maximum yields of up to approximately 5 % compared to those less than 2 %, respectively). This is perhaps surprising as, in regions where the square intermediate is stable with respect to the monomers, we can treat the square as a patchy particle with four patches, and thus we are looking at a regime where six patchy particles with four patches need to come together to form a final cluster. This pathway is analogous to the assembly of an octahedral cluster, which forms in good yield. The difference between the two systems is the larger distance between patches in the square intermediate compared with the octahedra-forming monomer. The larger distance between patches changes the number of accessible configurations which satisfy the energetic requirement for the patches to be bonded without particles overlapping. Also, in simulations, we see that the square intermediates are able to open up and form long chains which are common to the aggregate phase, even when we start the particle in this configuration.

The alternative pathway of hierarchical assembly via dimers was also attempted, but there was again no appreciable yield. This implies that it is not just the choice of intermediates that is limiting the rate of formation of truncated octahedra, but also the final desired structure. The wide angle between patches and low number of interactions per monomer particle favour the formation of aggregates.

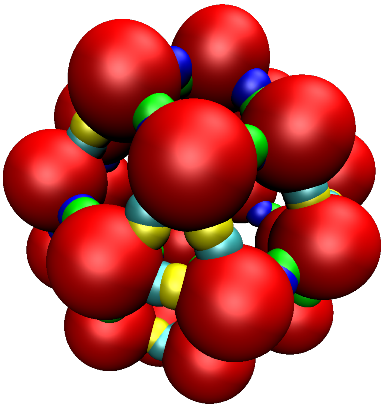

Figure 4.17: A rhombicuboctahedron.

Rhombicuboctahedron

The patchy particles we designed to be potentially capable of forming a stable rhombicuboctahedron had four patches of four different types. Patches A and B are able to bond, as are patches C and D. By A and B patches alone, the particles form squares, while by C and D alone, triangles are formed. By controlling the relative strengths of the interactions, we can encourage two possible hierarchical routes, and thus compare the strengths of squares and triangles as intermediates.

Unlike the truncated octahedron case, rhombicuboctahedral clusters could be formed both with and without torsional bonding constraints. Each particle had four patches, which increases the clustering temperature with respect to the kinetic aggregation temperature, allowing the cluster to be formed both in the presence and in the absence of torsional constraints. Since the aim of this chapter is to show the utility of hierarchical self-assembly in the absence of torsional bonding constraints, we have only analysed those simulations in detail.

Without torsional constraints, there are two possible assembly pathways which allow hierarchical assembly via assembly competent intermediates. If the interactions between A and B patches are stronger than those between C and D patches, then squares are the expected intermediates, while the inverse favours triangles. Figure 4.18 shows the results of the two different assembly pathways. It clearly shows that hierarchical assembly with triangles as an intermediate is more efficient than the pathway with square intermediates. This difference may be due to less attrition during the formation of triangles compared with squares (see section 4.2.3 for further investigation of this effect). Note the presence of aggregates in the square forming regions at low temperature that are absent from the triangular region.

Hierarchical assembly pathways again give a similar yield to non-hierarchical assembly

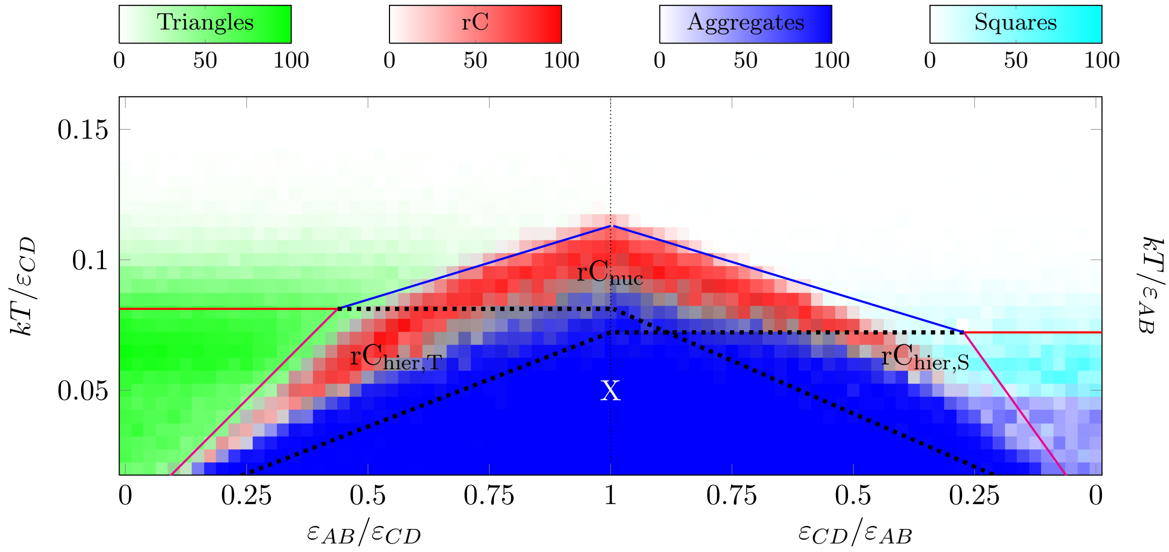

Figure 4.18: Rhombicuboctahedron yield plot at patch width $\sigma^{\text{PW}} = 0.45$. rC$_{\text{nuc}}$ is the expected nucleation region, rC$_{\text{hier,T}}$ and rC$_{\text{hier,S}}$ are the expected triangular and square-intermediate hierarchical assembly regions, and X shows the area where both squares and triangles are more stable than monomers. Settings not specified in the graph are given in table 4.1.

regimes. In areas where both triangles and squares are more stable than monomers (the downhill region), there is little useful assembly, and mainly aggregates form.

Snub cube

Our final 24 particle system is that of a snub cube. To assemble snub cubes, particles have five patches (labelled A to E). The interactions are configured such that patches A and B may bond, as may C patches with their own equivalents on other particles, and D- and E-type patches may also bond. A and B patches alone give square intermediates, C-patches give dimers, and D- and E-patches alone give triangles.

Here, we see an even larger window of formation between T_{clus} and T_{agg}, because the cluster is stabilised by five patches per particle. Figure 4.20 shows the results of simulations when the interaction strengths of CC and DE bonds were varied with respect to AB bonds. At low CC and DE interaction strengths, hierarchical assembly via squares is encouraged. The yield in the hierarchical region was poor when compared to rhombicuboctahedra in a similar region, and is partly due to the formation of strained AB triangles which can then assembly into distorted icosahedra. At higher temperatures, this competition is not an issue, as assembly into distorted icosahedra is reversible. However, at lower temperatures (in the hierarchical assembly regime), the formation of distorted icosahedra results in kinetic trapping of the intermediate. The existence of kinetic traps is a major issue for the

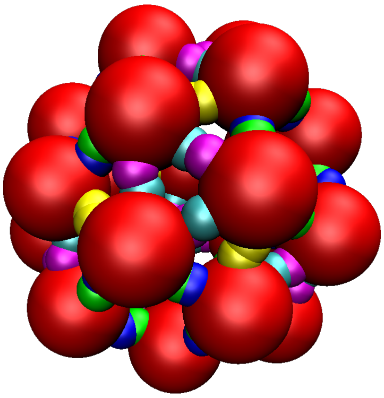

Figure 4.19: A snub cube.

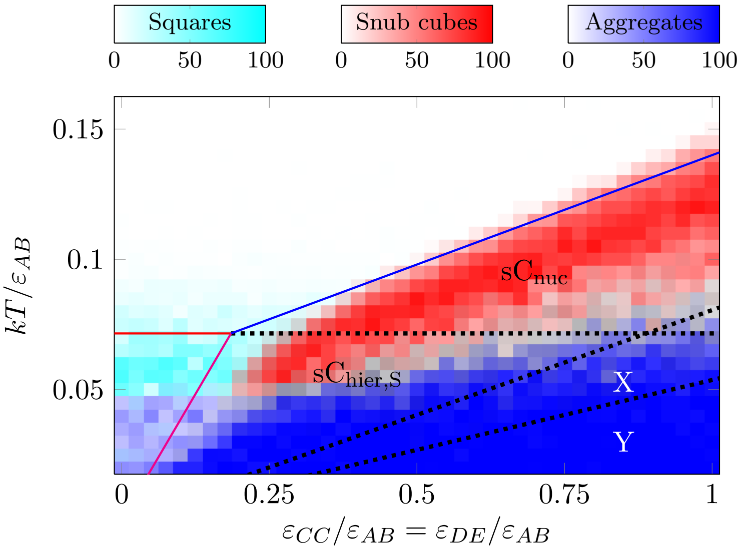

Figure 4.20: Snub cube yield plot (square intermediates) at patch width $\sigma^{PW} = 0.45$. sC_{nuc} is the expected nucleation region, $sC_{hier,S}$ is the expected square-intermediate hierarchical assembly regions, X shows the area where both squares and triangles are more stable than monomers, and Y shows the regions where squares, triangles or dimers are more stable than monomers. Settings not specified in the graph are given in table 4.1.

design of self-assembling systems, and the design of monomers to prevent the formation of unwanted structures is important.

Is it more important to have a high yield of intermediate or a wide range of conditions where intermediates can form reliably? Often, when looking at the yield plots of certain structures as a function of patch width, we see patch widths which offer very high yield at a narrow range of temperatures (which we term *patch width region of high yield*), and others which give a lower yield, but over a wider range (which we term *patch width region of wide yield*). Our results so far have been obtained using a patch width of 0.45 radians (except for the icosahedral results, which were run at 0.35 to enable

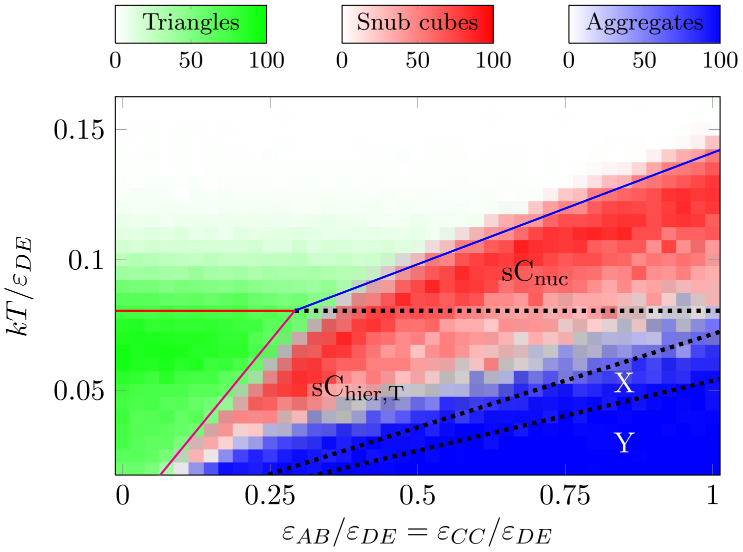

Figure 4.21: Snub cube yield plot (triangular intermediates) at patch width $\sigma^{\mathrm{PW}} = 0.45$. $\mathrm{sC_{nuc}}$ is the expected nucleation region, $\mathrm{sC_{hier,T}}$ is the expected triangular-intermediate hierarchical assembly region, X shows the area where both squares and triangles are more stable than monomers, and Y shows the regions where squares, triangles or dimers are more stable than monomers. Settings not specified in the graph are given in table 4.1.

comparison between different systems in regions where there was a reasonable yield), which corresponds to a regime of *wide* yield for squares, and to a regime of both *wide* and *high* yield for triangles. In order to see the effect of moving from a regime of *wide* yield to one of *high* yield of the intermediate, we tested the snub cube system's behaviour at a patch width of 0.35 radians in a region where we expect hierarchical assembly via squares (this corresponds to a region with a higher yield of squares, but aggregation of squares occurs at a higher temperature). Figure 4.22 clearly shows that the hierarchical assembly region for the narrower patch width is much smaller (see figure 4.20), and the maximum overall yield remains constant at approximately 95 %. This demonstrates that as long as the intermediate has an appreciable yield, hierarchical formation is possible, and that the wider the range of appreciable yields, the wider the range of hierarchical assembly. The maximum yield does not seem to depend on the maximum yield of the hierarchical intermediate: rather, it is a function of the final structure, and just requires a yield of correctly formed intermediates above a certain proportion. We look at the differences in the reliability of different intermediates in the next section.

4.2.3 The importance of intermediate choice

We have shown that it is important to have assembly competent intermediates when we are trying to follow a hierarchical assembly path. Results on the snub cube system have shown that the wider the range in which an intermediate forms successfully, the wider the

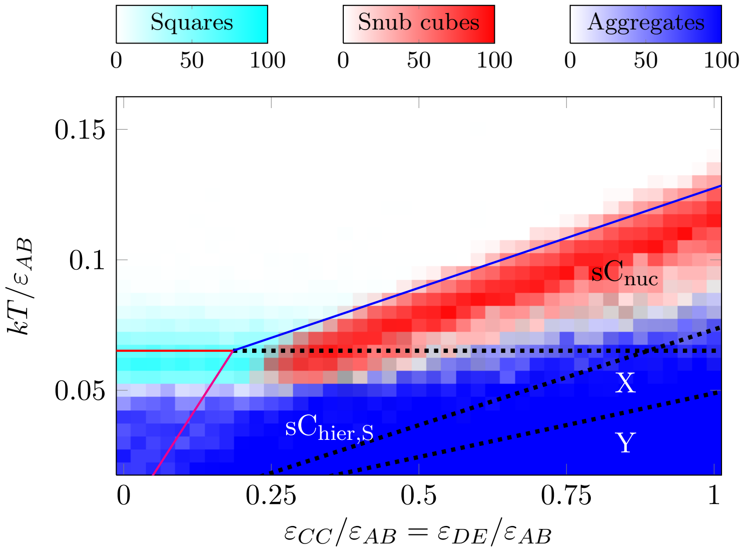

Figure 4.22: Snub cube yield plot with a narrower patch width of $\sigma^{\mathrm{PW}} = 0.35$. Regions are labelled as figure 4.20. Settings not specified in the graph are given in table 4.1.

range in which we can form the target structure hierarchically, and that the maximum yield of the intermediate is not relevant. On a patch width yield plot, there are three limits on the formation of a small cluster:

1. The maximum temperature where the cluster is unstable with respect to monomers (above successful assembly).
2. Where a liquid-like aggregate becomes more stable than the gas of clusters (to the right of the diagram).
3. Where the formation of kinetic aggregates, including chains, is faster than the correct formation of monomers (the region in the bottom left of the diagram).

We tested the formation of homomeric-bonding dimers, and cyclically bonding triangles and squares (which would be assembly competent were they to have another patch). The results are detailed below.

Dimers

Dimers are able to form successfully across a wide range of conditions (see figure 4.23). Formation of larger aggregates is simply not possible since, at low patch width, particles are limited to bonding to only one other particle. However, without the use of some form of torsional constraint, a dimer is neither an assembly competent nor an assembly-incompetent intermediate, as it cannot be trapped into either compatible or incompatible geometry owing to the free rotation about the dimeric bond. This free rotation also means that the

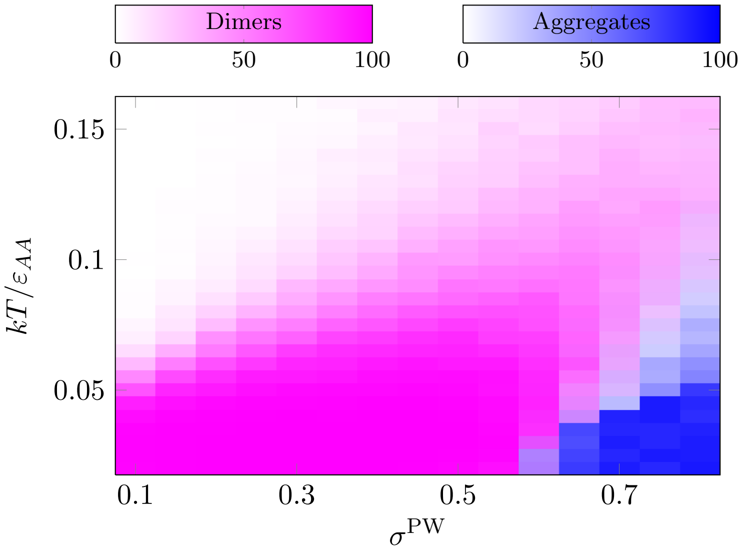

Figure 4.23: Dimer yield plot. Settings not specified in the graph are given in table 4.1.

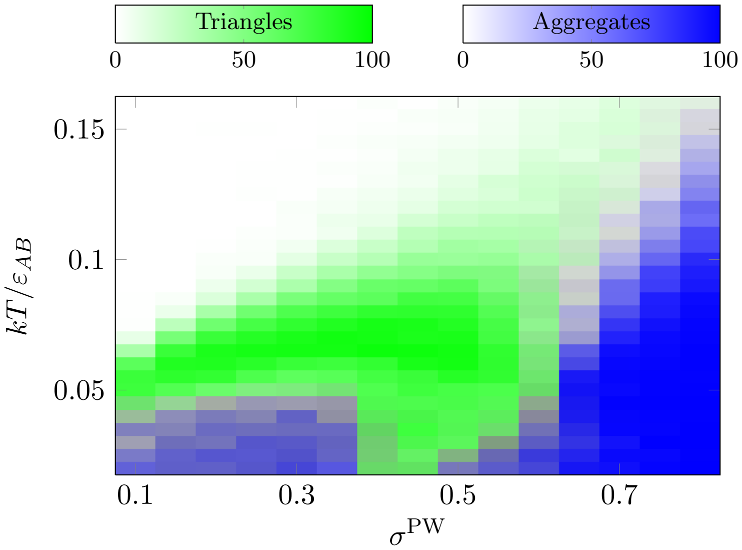

Figure 4.24: Triangle yield plot. Settings not specified in the graph are given in table 4.1.

complexity of the second step is not much reduced compared to the assembly of targets directly from monomers. Little gain was seen in regions where hierarchical assembly via dimers was attempted.

Triangles

Triangles are the smallest and thus probably the fastest forming assembly competent intermediates. Our results on small clusters show that assembly via triangular intermediates gives the highest yield. Figure 4.24 shows the yield of triangular intermediates at different patch widths. It is worth noting that the highest and widest yield regions overlap strongly, and this is one of the underlying reasons for the effectiveness of triangles as a hierarchical intermediate.

Owing to the small internal angle between bonds, the formation of liquid-like aggregates is disfavoured (and moved to high patch width), though some kinetic aggregates of chains may still form (seen at narrow patch width at the bottom left). These chains normally occur when two clusters of size two come into contact with one another and bond, and the low temperature means that it is difficult for these to break apart, despite the fact that a triangular cluster and a monomer would be a lower free energy state. The chains, when formed, are often longer than the width of our simulation box, but none were observed that bound to themselves.

Squares

In the absence of torsional constraints, the temperature window of formation of squares is narrow, since the formation of long chain-like aggregates is possible. The formation of aggregates is more pronounced compared to triangles. One of the reasons for this aggregation is that once a trimer has assembled, it only requires the correct orientation of the two bonds to form a third and complete the triangle, whereas once a tetramer has formed three bonds, it must have the correct orientation of all three bonds for the fourth to be able to complete the square. The need for three bonds to orientate takes longer than the two required for the triangle to form. This means that it is more likely for the tetramer to bind to another cluster and form a chain. Figure 4.25 shows the yield of cluster formation when the patch width is varied. It should be noted that the patch widths which offer the highest maximum yield of squares are in the range 0.3–0.4 radians, while patch widths of 0.4–0.5 radians give a lower maximum yield, but a wider range of square formation.

The clustering temperature for the formation of squares is also considerably lower than that for the formation of triangles. This decrease in clustering temperature can be explained using the simple thermodynamic arguments with which we have estimated the transitions throughout this chapter. Both triangles and squares have two bonds per particle, but the triangles only require three particles per cluster, giving them a higher entropy when compared to the square clusters.

The weakness of squares as a hierarchical intermediate is clear in all systems that were placed in conditions where we expected hierarchical assembly via square intermediates. This poor performance of squares as a hierarchical intermediate is due to the slow formation of squares compared to aggregate formation and trapping.

An example of a structure formed in the square system can be seen in figure 4.26. This aggregate is a long chain. Other observed structures included five-membered and eight-membered rings. The chains were often longer than the periodic simulation box, and

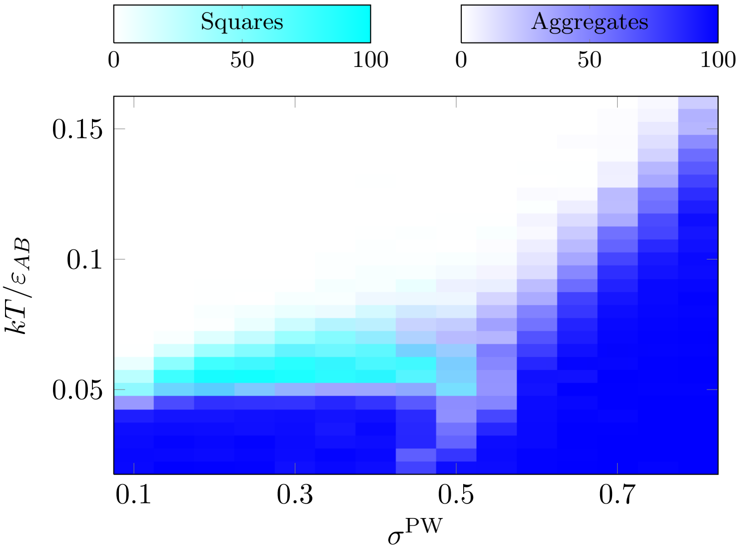

Figure 4.25: Square yield plot. Settings not specified in the graph are given in table 4.1.

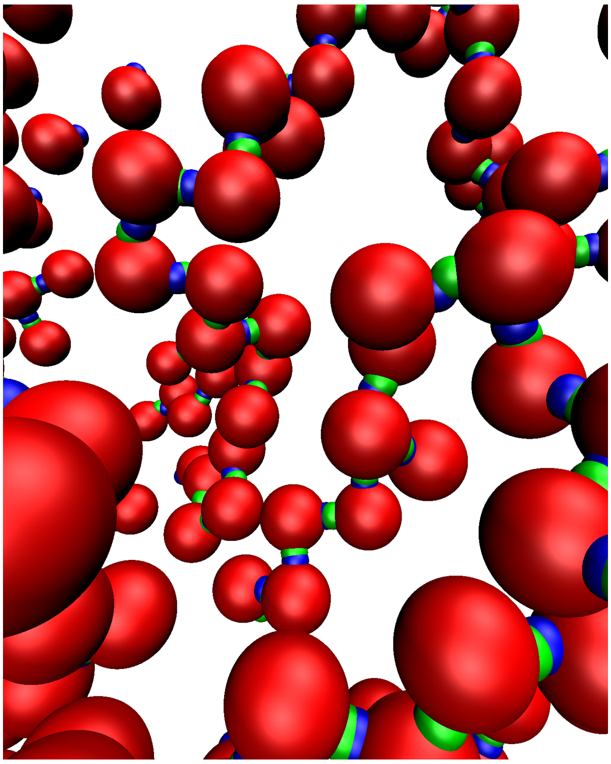

Figure 4.26: Aggregates formed by square-forming particles under certain conditions ($kT = 0.04\,\varepsilon$, $\sigma^{\mathrm{PW}} = 0.3$). Settings not specified are given in table 4.1.

so may have interacted weakly with themselves, but a chain which joined to itself was never observed in the simulations.

4.3 Conclusions

It is clear from the results that hierarchical assembly at one temperature is a feasible assembly mechanism. However, it usually gives similar maximum yields to non-hierarchical processes. In particular, it does not seem to offer a pathway to make more complex or more difficult-to-assemble targets (*e.g.* truncated octahedra) than using our previous approaches. As such, it is something of a null result. Although breaking the assembly into two simple steps seems potentially advantageous, the requirement that the first stage be particularly efficient to avoid attrition is quite stringent. Furthermore, the second stage involving the assembly of the intermediate clusters can be more challenging than one would imagine; *e.g.* the assembly of six squares into a truncated octahedron barely occurs. Hierarchical assembly does, however, permit some extension of the available parameter range for the formation of a target structure, disfavouring the formation of aggregates. The type of assembly mechanism seems not to be as important as the removal of other competing structures because of a more specified set of interactions (which acts to reduce the number of ways aggregates can form and hence their entropy). This can be seen for the IWC and IhO systems, and the truncated tetrahedral systems.

We have shown that the wider the range of conditions in which an intermediate can form successfully, the wider the range of hierarchical assembly. The assembly competence of intermediates is also a key factor. Our general rules are that assembly competent triangular intermediates and monomer particles with many patches give the widest parameter range for the formation of a target structure.

Chapter 5

Modelling pore formation

5.1 Introduction

This chapter illustrates how we can use our simple patchy particle systems to model biological self-assembly. This chapter looks at one such system, namely the formation of α-haemolysin pores from the *Staphylococcus aureus* bacterium on a membrane. The chapter examines the biological background of the formation of the α-haemolysin pores and then explains how our patchy particle model was augmented to allow self-assembly to occur on a surface. We compare our results with those seen in experiment, and also use reaction kinetics to help provide further insight into what might be happening experimentally.

5.2 Motivation

α-Haemolysin pores are heptameric protein complexes which create a channel through a membrane of many different target cells [211]. These pores are produced by *Staphylococcus aureus* and form part of its ability to damage neutrophils and thus evade part of the body's immune response to infection [211]. *In vitro*, they are found to form spontaneously on the surface of a membrane. Each monomer is identical and the crystal structure of the pore can be seen in figure 5.1 (each of the monomers is given a different colour to make the connectivity clear). *In vivo*, the barrel (lower region of the right image in figure 5.1) inserts into and spans the membrane, while the upper rim section is found on the outer surface of the membrane [212].

We were motivated to study the α-haemolysin system by the experimental research on the self-assembly of these pores by the Wallace group in Oxford [8]. In particular, the results showed intriguing features that were not straightforward to interpret mechanistically. We therefore decided to investigate a simple theoretical model of this system in order to try

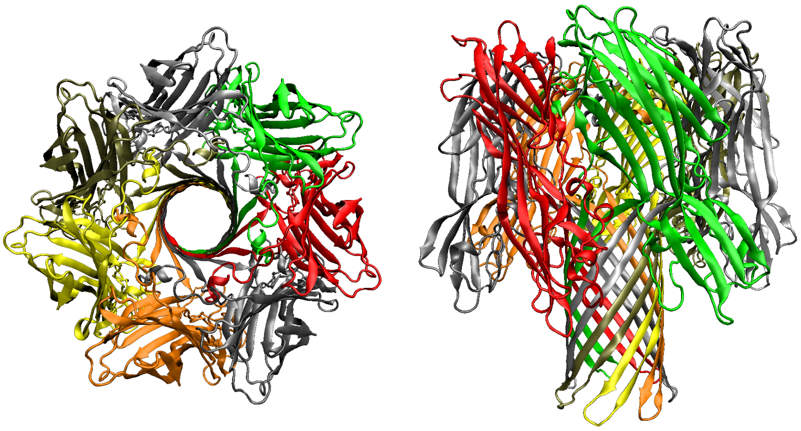

Figure 5.1: α-Haemolysin pore as viewed from the outside of the cell and from the plane defined by the membrane. The barrel region in the lower region of the right-hand diagram inserts into the membrane. Image created from the protein database entry for 7AHL [213, 214].

to obtain insights into the mechanism of self-assembly that would aid the interpretation of the experimental results.

5.2.1 Why is this experimentally interesting?

The Wallace group use fluorescent labelling to track the formation of proteins on a synthetically created phospholipid bilayer. Their results for the α-haemolysin pore only show the presence of monomeric and heptameric species: no intermediates are observed under biologically relevant conditions. This is not what is expected by intuition: a single step process in which 7 monomers must come together to form a pore seems unlikely on the basis of simple reaction dynamics (see section 5.4.5). The formation of pores appears to be a rare event, but one that happens rapidly. Pore self-assembly is also highly selective, with no hexameric or octameric products being observed, as shown by photo-bleaching experiments [215]. We look into how the experimental results compare with our simple model and, with some simple reaction pathway considerations, we attempt to explain the formation of the pore. Experiments on α-haemolysin mutants have shown that assembly is most likely to occur on the membrane and that a prepore is formed which then inserts through the membrane [216]. We consider this insertion step as a modification to the energy landscape given by simulations of our model. If it is indeed a single step insertion process, then this is a valid approximation. Other experiments have shown that the formation of the β-barrel stem domain is a co-operative process [217]. We look at this co-operativity by simulating an internal degree of freedom in our simulation and using a kinetic model using an activated pathway of assembly.

5.2.2 General importance of pore-forming toxins

Pore forming toxins are responsible for the death of many people worldwide. They are responsible for the toxic effects of diseases such as diphtheria [218] and *Staphylococcus aureus* infections [219] (including the methicillin and vancomycin resistant strains, commonly known as MRSA and VRSA).

A better understanding of pore formation may help with the design of possible antibiotic agents [220]. Bacteria often use peptides capable of forming transient pores to give them a competitive advantage over other bacteria for scarce food resources. Their pores are often bacteriocidal, and generally at least bacteriostatic. The pore does not affect the bacterial strain which produces it, and only prevents other bacteria from growing easily. One of these peptides, nisin, has been used in a purified form as a food preservative for many years [221].

5.3 Model details

In order to simulate this system, we had to make a number of changes to our model.

5.3.1 Surface potential

In order to model the membrane, we introduced a surface potential into our model. The requirements for the potential were that, firstly, the particles would bind more strongly to the surface than to each other (experimentally, monomers are found first to adsorb onto the surface and then to assemble [212]), and secondly, that the particles only gained the full interaction energy when they were suitably orientated with respect to the surface. The potential we chose was a Lennard-Jones potential in z, the distance of the patchy particle from the surface, with the attractive part modulated by a Gaussian curve in ϕ, the angle between the plane containing the ring forming patches and the "membrane surface" (figure 5.2). We used a well depth of $2\,\varepsilon$ (twice the well depth of the attractions between particles), and set the Gaussian to have a standard deviation of 0.1 (this is equivalent to a patch width of the same value). The full potential is given by

$$U^{\mathrm{SURF}}(z,\phi) = 8\varepsilon\left(\left(-\frac{\sigma^{\mathrm{LJ}}}{z}\right)^{6} \times \exp\left(-\frac{\phi^2}{0.2}\right) + \left(\frac{\sigma^{\mathrm{LJ}}}{z}\right)^{12}\right) \tag{5.1}$$

and its form is depicted in figure 5.2.

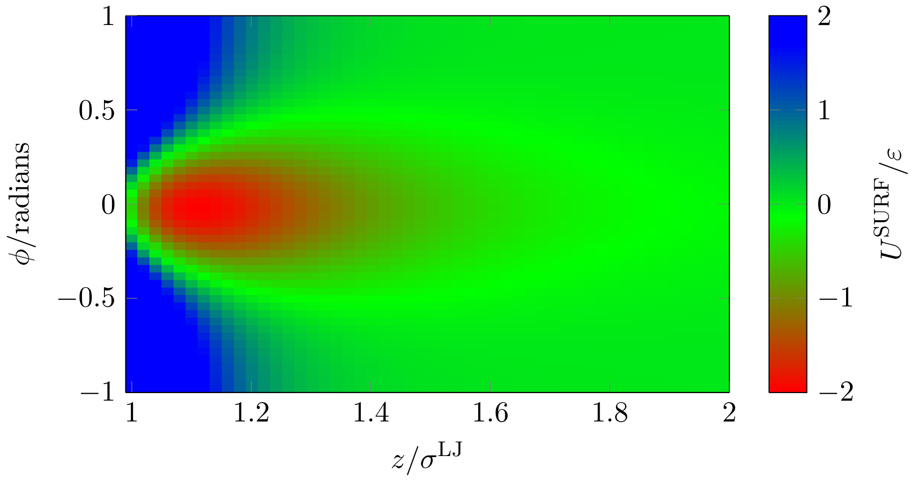

Figure 5.2: Plot of the potential used as a function of z position and rotation ϕ. $\phi = 0$ is the point at which the pore forming patches are coplanar with the "membrane". Note that only a small region is shown. Below $z = \sigma^{\mathrm{LJ}}$, the potential is steeply repulsive for all values of ϕ.

5.3.2 Flipping

As explained previously, proteins are able to change their tertiary structure (section 1.3.1), and experiments have shown α-haemolysin to do so [217]. In order to allow a simple simulation of this phenomenon, our model was modified to allow particles to change the position of their patches. This was implemented by the creation of a new Monte Carlo move type, which attempted to change a single particle from one type to another. This move used the standard Metropolis Monte Carlo acceptance criterion (equation 2.22). The probability of attempting any type of internal rearrangement move was set in the control file and was initially set such that the probability of attempting a *flip move* was such that the expectation was for one flip attempt to be made in every 100 move attempts.

5.4 Heptameric pore formation

The structure of the heptameric ring forming particles is shown in figure 5.3. Each monomer has two patches: A (blue) and B (green). Each of them can only interact with its opposite patch on any other particles (an A-B type system). As we were interested in modelling a biological protein system, we used torsional constraints with a torsional patch width set at 0.2 (though this had little effect for particles adsorbed on the surface, as an effective surface orientational constraint already existed owing to the potential).

Initially, we scanned the parameter space of our model to find conditions under which heptameric rings formed on the surface reliably and in a reasonable time. Interestingly, a rather narrow patch width, 0.1 radians, had to be used to ensure the product was

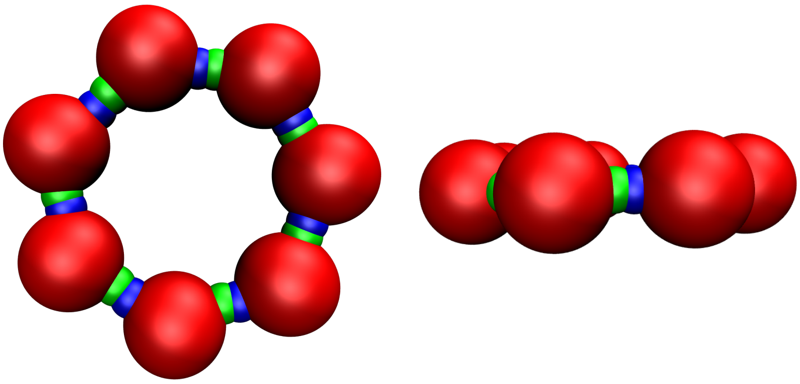

Figure 5.3: The representation of an α-haemolysin pore within our simple model (to be compared with figure 5.1).

monodisperse. For wider patches, significant numbers of hexamers and octamers were also found to form, particularly at temperatures close to where self-assembly first occurs. This is simply because of the relatively similar internal angles in these polygons (128.57° for a heptagon compared to 120° and 135° for hexagons and octagons, respectively). However, it illustrates that the interactions between the proteins in the α-haemolysin pore must be angularly very specific in order to give the monodisperse assembly seen by the Wallace group [8, 215].

5.4.1 Initial dynamic results

We used the potential described in section 5.3.1. After checking the speed at which the pore was formed under a wide range of conditions, a system with a volume number density of 0.1 $/(\sigma^{\mathrm{LJ}})^3$ and 49 particles was used initially (this is equivalent to a surface number density of 0.2 $/(\sigma^{\mathrm{LJ}})^2$, and the initial dynamic simulations were run for 10^6 Monte Carlo cycles. We used a cutoff of 3.5 σ^{LJ}.

We show the yields of different species in figure 5.4. It is clear that, as we intended, our adsorption temperature (T_{Ads}) is far above our heptamerisation temperature (T_{Hept}). At low temperature, we again are faced with the issue of trapped partially formed clusters (arrested assembly), as we have seen in other simulations before. The aggregates were somewhat different from those formed in a bulk solution and tended to consist of partially formed clusters which bonded to each other. We refer to these as *Pac-Man chains*, as each partially formed ring appears to "eat" the next in the chain (figure 5.5).

Figure 5.6 shows the yield of heptamers when these simulations are started from fully formed pores. The results clearly confirm that the partially formed clusters are merely a kinetic trap. The values of the heptamerisation temperature, T_{Hept}, defined as the point at which there are equal yields of heptameric rings and monomers, are nearly identical.

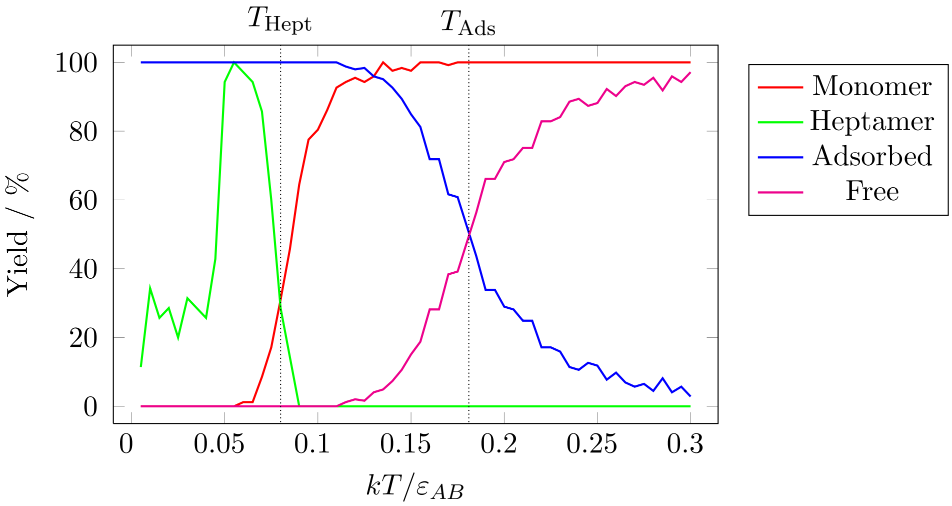

Figure 5.4: Yield of heptamers and number of particles adsorbed in our system when starting from monomers as a function of temperature after 10^6 Monte Carlo cycles, in a 3D box with a simulated surface, using a volume number density of $0.1\ /(\sigma^{\text{LJ}})^3$ and $\sigma^{\text{PW}} = 0.1$.

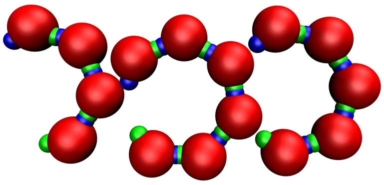

Figure 5.5: An example of the type of partially formed clusters that occur ($kT = 0.03\,\varepsilon$). Note that they are able to bond slightly to one another despite the low patch width, and at low temperature, these *Pac-Man chains* are common.

The small difference is in part due to the natural noise of the dynamic runs and is also the result of having finite temperature steps. At the heptamerisation temperature, there is a significant equilibrium population of clusters that are neither monomeric nor heptameric (these are mainly dimers and trimers).

The yield of the different structures as a function of time is shown in figure 5.7 for a regime where the self-assembly of heptameric rings occurs successfully. It can be clearly seen that assembly occurs by the addition of monomers to growing clusters with a population of intermediates that evolves to larger sizes as a function of time.

Our results show equilibrium populations of species other than monomers and heptamers, while experimental conditions only have monomers and heptamers present in the equilibrium populations. Dynamically, we see that the onset of assembly is rapid and simultaneous (there is a rapid decrease in monomer concentration early on in the simulations). We also see a relatively long-lived intermediates, which implies that they are stable with respect to monomers. By contrast, experimental studies show that the assembly of heptamers is sporadic, but rapid when it does occur [8]. The intermediates

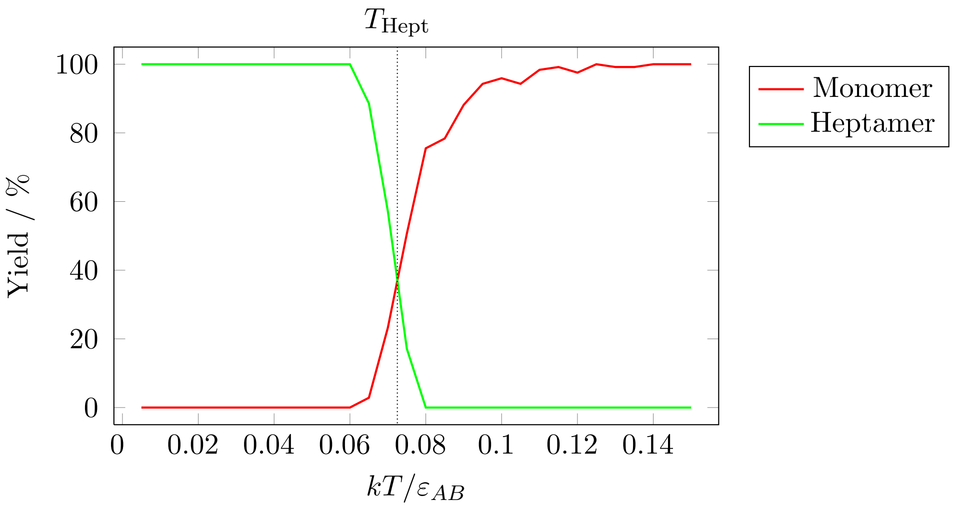

Figure 5.6: Yield of heptamers and number of particles adsorbed in our system when starting from formed rings as a function of temperature. Settings as in figure 5.4.

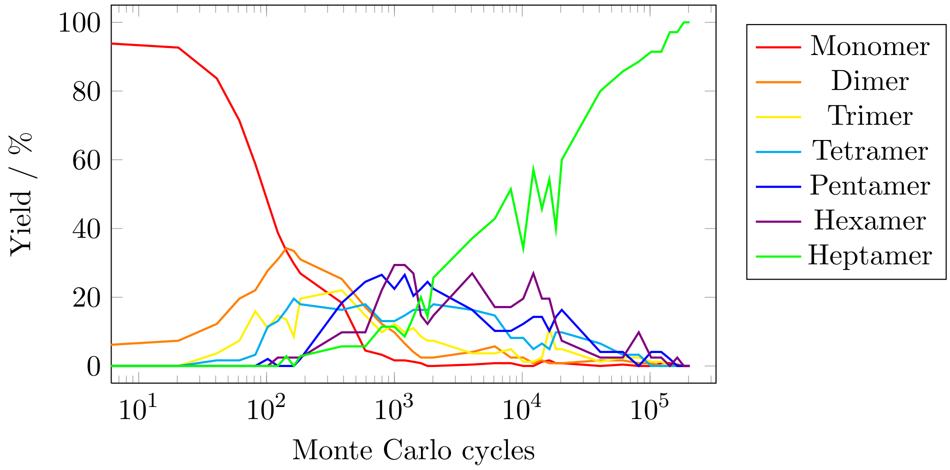

Figure 5.7: Time resolved yield of intermediates and heptameric rings at a temperature where a high yield of heptamers is obtained. $kT/\varepsilon = 0.055$.

are unresolvable on the experimental timescale. These differences imply that there is a free energy barrier that is absent in our model, but is present in the experimental system that regulates the rate of nucleation. We consider this further in section 5.4.4.

5.4.2 Relevant experimental conditions

The experimental surface number density is of the order of 10 $\mu\mathrm{m}^{-2}$ [8]. With the radius of the monomer being of the order of 1 nm, this is equivalent to a relative surface number density of 10^{-5} / $(\sigma^{\mathrm{LJ}})^2$. We are not able to simulate self-assembly at these densities because the low collision rate between particles makes it computationally unfeasible. However, we have considered the effect of lower densities to show that there is no change to the

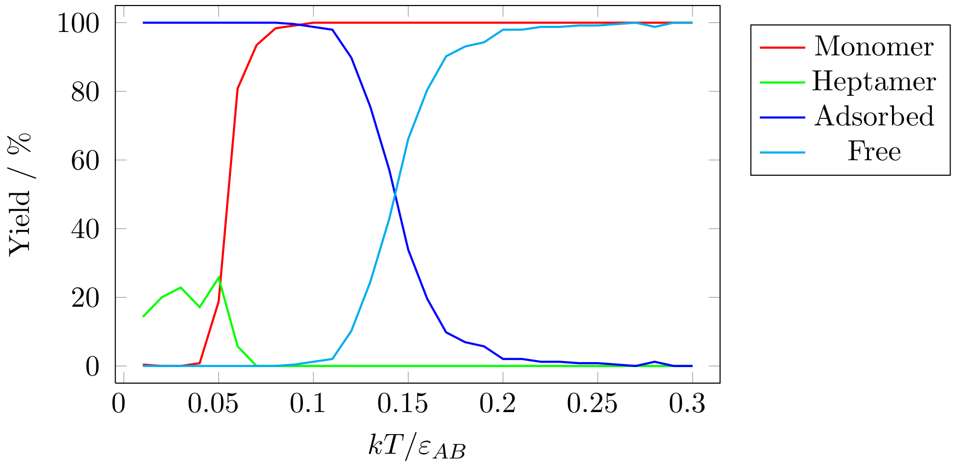

Figure 5.8: Yield of heptamers and number of particles adsorbed in our system when starting from monomers as a function of temperature after 10^8 Monte Carlo cycles, in a 3D box with a simulated surface, using a surface number density of 0.001 / $(\sigma^{\mathrm{LJ}})^2$ and $\sigma^{\mathrm{PW}} = 0.1$.

general patterns of thermodynamic stability. Note that lower densities also lead to much longer-lived intermediates in our simulations.

5.4.3 Dynamic results at lower densities

We ran simulations at a surface number density 200 times lower than our initial dynamic runs (*i.e.* we used $\sim$0.001 / $(\sigma^{\mathrm{LJ}})^2$). In order to get a reasonable yield, we used 10^8 Monte Carlo cycles (which required approximately 200 processor core hours per simulation) and five repeats per point. Figure 5.8 shows the yield when we started from monomers and it demonstrates a significantly lower yield of heptamers; this is simply due to the amount of time taken for part-formed heptamers to diffuse over the surface, and this is the reason most of our runs are carried out at higher densities. Figure 5.9 shows the yields when we start from formed heptamers and clearly demonstrates that the thermodynamic behaviour is similar for both densities, and that the limiting factor is simply the diffusion time required for the less dense system. Thus, we are justified in generalising thermodynamics from our higher density system, which we are able to simulate in a reasonable time.

5.4.4 Location of the barrier

In the comparison of our dynamic results with experiment, the experimental dynamics seem to indicate that there is a much larger free energy barrier present than in our simulations. To examine this issue further, we investigate the nature of the free energy profile for the heptamer formation in our model. These profiles are computed using umbrella sampling (see section 2.2.4).

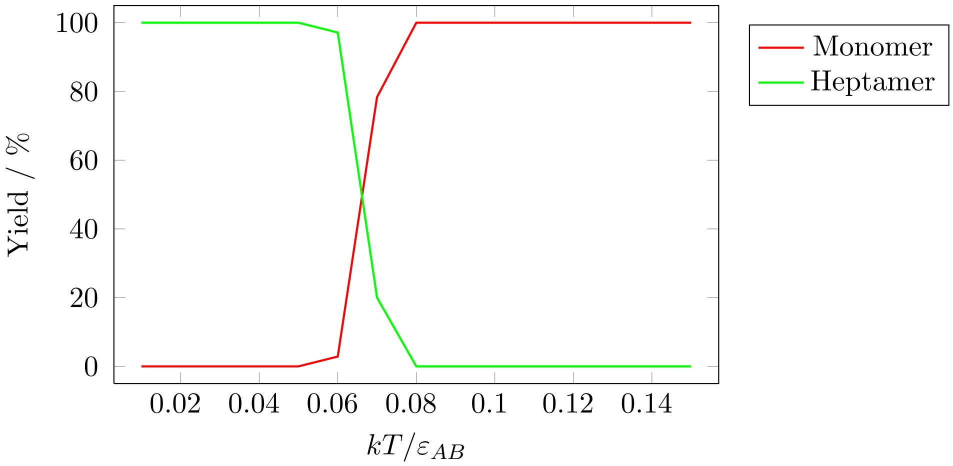

Figure 5.9: Yield of heptamers and number of particles adsorbed in our system when starting from formed heptamers as a function of temperature after 10^8 Monte Carlo cycles, in a 3D box with a simulated surface, using a a surface density of $0.001\ /(\sigma^{\text{LJ}})^2$, and $\sigma^{\text{PW}} = 0.1$.

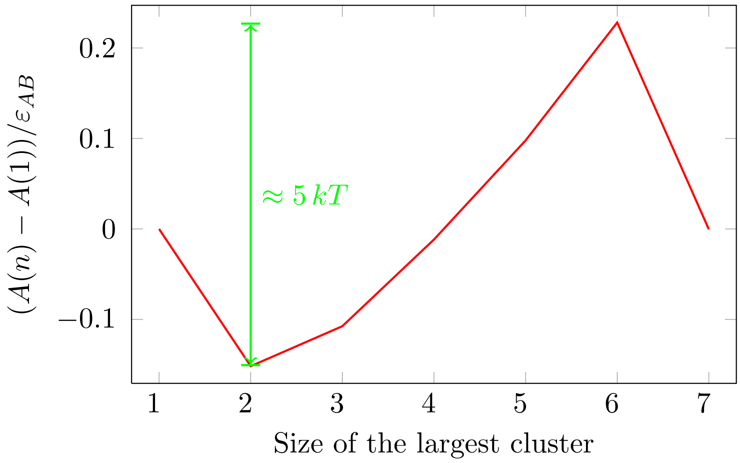

Figure 5.10: Free energy profile for the formation of a heptameric pore at the transition temperature of $0.078\,\varepsilon/k$.

As with previous implementations of umbrella sampling to calculate free energy profiles for monodisperse self-assembly in our group, we performed the umbrella sampling using a system with exactly the right number of particles to form a single target structure. We did this because it is hard to devise an order parameter that can drive the simultaneous formation of multiple target structures. As our order parameter, we used the size of the largest cluster. The smaller system size allowed us to use a slightly lower volume number density of $\sim 0.01\,/\,(\sigma^{\text{LJ}})^3$ (equivalent to a surface number density of $(\sim 0.1\,/\,(\sigma^{\text{LJ}})^2)$, whilst still being able to achieve equilibrium in a reasonable time.

From our previous yield plots (which were at slightly different densities), we expected the transition between heptamers and monomers to occur at around a temperature of

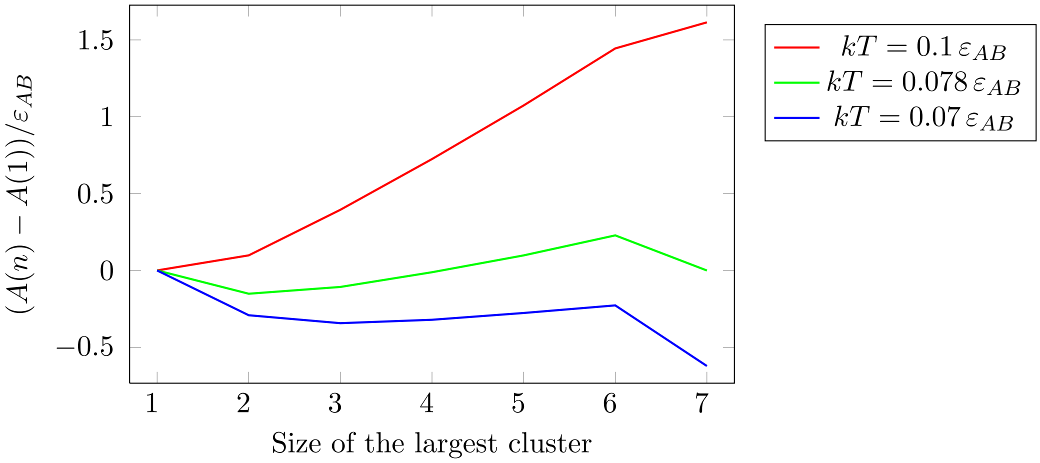

Figure 5.11: Free energy profiles for the formation of a heptameric pore, at temperatures above, at and below the transition temperature, showing how the barrier changes.

$0.075\,\varepsilon/k$, and our umbrella sampling runs have found the point of equal probability of monomer and heptamer to be at $0.078\,\varepsilon/k$. The barrier to formation at the transition temperature is shown in figure 5.10. It is clear that a system containing a dimer as the largest cluster has a lower free energy than a system of just monomers at this temperature. The size of the barrier ($\sim 5\,kT$) is too small to regulate nucleation, and we therefore see intermediates in our simulations.

Figure 5.11 shows how the free energy profile changes at temperatures above and below the transition temperature. Below the heptamerisation temperature, there is always stabilisation of intermediates with respect to monomers: this means that intermediates are long-lived.

With an insertion step

Our model of pore assembly is based on the simple formation of a cyclic prepore protein complex on the surface of the membrane and does not explicitly consider the process of insertion of the channel of the pore into the membrane. Here, we explore whether this might be the source of the discrepancies between our model and the experimental results. Clearly, pore insertion is a complex process involving not only a co-operative rearrangement of the protein complex, but also a rearrangement of the lipid bilayer. We take a very simple approach and assume that pore insertion occurs after a heptamer is formed and that insertion involves a decrease in free energy of magnitude $\Delta_{\text{ins}}A$. This approach allows us to use the data that we have already obtained from our umbrella sampling runs to describe the equilibrium between the non-inserted species. The relative probabilities of inserted and non-inserted heptamers are simply related to $\Delta_{\text{ins}}A$ (this type of pathway is demonstrated in the next section in figure 5.14).

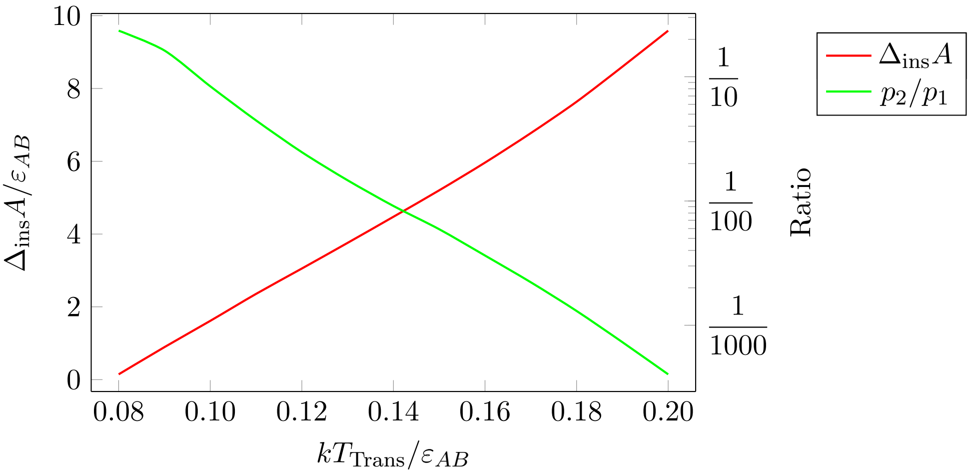

Figure 5.12: The required insertion energy and the ratio of probabilities of a randomly chosen particle being in a dimer or a monomer (p_2/p_1) when the transition temperature is varied for our system of heptamer-forming particles based on our umbrella sampling results of a single cluster at a volume number density of $\sim 0.01\,/\,(\sigma^{\text{LJ}})^3$.

We assume that the configuration corresponding to a given value of our order parameter, n, is a single cluster of size n with all other paticles being monomeric. We use our umbrella sampling data to find the insertion energy required for the inserted cluster and a purely monomeric system to have equal probability as a function of the transition temperature, T_{Trans}. The results are shown in figure 5.12. The temperature T_{Trans} increases with $\Delta_{\text{ins}}A$, and the relative probability of a particle being in a dimer or a monomer (p_2/p_1) decreases with increasing $\Delta_{\text{ins}}A$.

The introduction of an insertion free energy allows our model to be in better agreement with the experimental data in terms of the thermodynamics. Namely, we can tune $\Delta_{\text{ins}}A$ (and thus T_{Trans}) until dimers are negligible compared to monomers at the transition temperature. Using our data we can predict the proportions of dimers and monomers that would be expected experimentally. For example, if the gain in free energy of insertion is the same as the internal energy gain for prepore formation, *i.e.* $-7\,\varepsilon$, then $p_2/p_1 \approx 0.002$ and $kT \approx 0.17\varepsilon$, and thus our insertion energy is $\sim 40\,kT$.

Although we have not explicitly simulated the dynamics with this addition to the model, it is clear that it would also help to make the dynamics more consistent with experiment. The increased free energy barrier for heptamer formation at T_{Trans} would make pore formation a rarer event. Furthermore, the greater slope of the free-energy profile would make any intermediates that form short-lived and more likely to fall apart. Pore-formation events, when monomers happen to come together to form a heptamer that then inserts, would have to occur rapidly if the intermediate clusters are not simply to fall apart before the event can happen. This picture of pore-formation as an infrequent but rapid event is much more consistent with experiment.

5.4.5 Kinetic modelling

We are able to use the steady state approximation [222] to model the stepwise addition of monomers. We can initially assume that each step is at equilibrium except for the last. Note that all concentrations in this section refer to surface concentrations, as do rate constants, as all reactions take place on the membrane surface.

Our first simple model of the pathway is:

$$7\,\mathrm{A} \underset{k_{-1}}{\overset{k_1}{\rightleftharpoons}} \mathrm{A}_2 + 5\,\mathrm{A} \underset{k_{-2}}{\overset{k_2}{\rightleftharpoons}} \mathrm{A}_3 + 4\,\mathrm{A} \underset{k_{-3}}{\overset{k_3}{\rightleftharpoons}} \mathrm{A}_4 + 3\,\mathrm{A} \underset{k_{-4}}{\overset{k_4}{\rightleftharpoons}} \mathrm{A}_5 + 2\,\mathrm{A} \underset{k_{-5}}{\overset{k_5}{\rightleftharpoons}} \mathrm{A}_6 + \mathrm{A} \xrightarrow{k_6} \mathrm{A}_7 \tag{5.2}$$

The rate of the first step is:

$$\frac{\mathrm{d}[\mathrm{A}_2]}{\mathrm{d}t} = k_1[\mathrm{A}]^2 - k_2[\mathrm{A}_2][\mathrm{A}] - k_{-1}[\mathrm{A}_2] \tag{5.3}$$

For each subsequent step to form clusters of size i (where $i < 7$), the rate is:

$$\frac{\mathrm{d}[\mathrm{A}_i]}{\mathrm{d}t} = k_{i-1}[\mathrm{A}][\mathrm{A}_{i-1}] - k_i[\mathrm{A}_i][\mathrm{A}] - k_{-(i-1)}[\mathrm{A}_i] \tag{5.4}$$

Using the steady state approximation for all intermediates of size i (where, again, $i < 7$) gives:

$$[\mathrm{A}_i] = \frac{k_{i-1}[\mathrm{A}][\mathrm{A}_{i-1}]}{k_i[\mathrm{A}] - k_{-(i-1)}} \tag{5.5}$$

This gives the overall form for the rate of formation of our heptameric pore to be:

$$\frac{\mathrm{d}[\mathrm{A}_7]}{\mathrm{d}t} = \frac{k_1k_2k_3k_4k_5k_6[\mathrm{A}]^7}{(k_2[\mathrm{A}] + k_{-1})(k_3[\mathrm{A}] + k_{-2})(k_4[\mathrm{A}] + k_{-3})(k_5[\mathrm{A}] + k_{-4})(k_6[\mathrm{A}] + k_{-5})} \tag{5.6}$$

We look at a number of limiting possibilities which will allow us to simplify our rate equation. Note that, as our simulations are in the canonical ensemble, the relevant free energy is the Helmholtz energy (A). We plot schematics of the shape of the reaction co-ordinates in the following examples. Note that our reaction pathway consists of stepwise monomer addition and thus the reaction co-ordinate corresponds to the size of the growing cluster.

Steeply uphill reaction

If the forward reaction is steeply uphill (figure 5.13), with a hexameric transition state, we can state that $k_i[\mathrm{A}] \ll k_{i-1}$ (for $1 < i < 5$). As our final step may be considered to

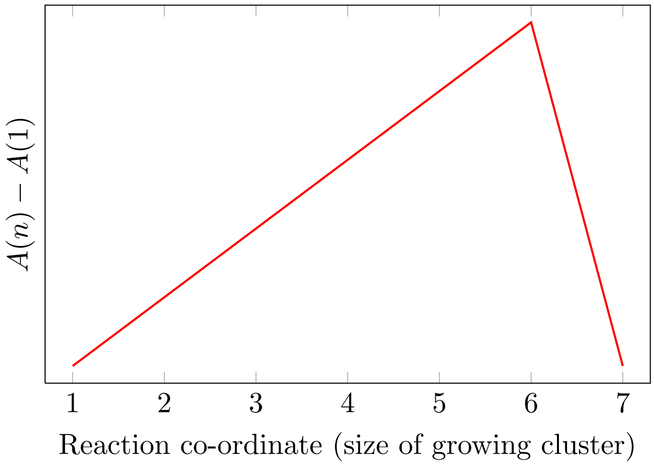

Figure 5.13: Schematic of the free energy of an uphill reaction pathway as a function of cluster size at the transition temperature.

be more steeply downhill than our reverse reactions, we can also state that $k_6[\mathrm{A}] \gg k_{-5}$. Thus we are able to simplify equation 5.6 to:

$$\begin{aligned}\frac{\mathrm{d}[\mathrm{A}_7]}{\mathrm{d}t} &= \frac{k_1k_2k_3k_4k_5k_6[\mathrm{A}]^7}{k_{-1}k_{-2}k_{-3}k_{-4}k_6[\mathrm{A}]} \\ &= \frac{k_1k_2k_3k_4k_5[\mathrm{A}]^6}{k_{-1}k_{-2}k_{-3}k_{-4}}\end{aligned} \tag{5.7}$$

This depends on [A] in the sixth order! If we set the first 5 forward reactions to all be equal ($k_i = k_+$ for $1 \leq i < 6$), and the reverse reactions similarly ($k_{-i} = k_-$ for $1 \leq i < 5$), our equation simplifies to:

$$\frac{\mathrm{d}[\mathrm{A}_7]}{\mathrm{d}t} = \frac{k_+^4k_5[\mathrm{A}]^6}{k_-^4} \tag{5.8}$$

Or, by setting $K = k_+/k_-$, we can write:

$$\frac{\mathrm{d}[\mathrm{A}_7]}{\mathrm{d}t} = k_5K^4[\mathrm{A}]^6 \tag{5.9}$$

We can use simple thermodynamics to estimate the energy profile. If we assume that translational entropy is the most relevant entropic term, and that this is simply proportional to the number of species present, then the loss in entropy as we increase the size of the growing cluster can be approximated as the translational entropy of a monomer, and thus be constant ($\approx \Delta_{\mathrm{add}}S$). We can also assume that the gain in internal energy is just the energy gained by bonds being formed ($-\varepsilon$ per bond: one bond is formed for the first 5 steps, and two are formed in the final step). Using these approximations, the initial slope would be expected to be $-\varepsilon - T\,\Delta_{\mathrm{add}}S$, and the final step would have a slope of $-2\varepsilon - T\,\Delta_{\mathrm{add}}S$. At the transition temperature, T_{Hept}, the free energies of the monomeric and heptameric systems are equal. Therefore, $A(7) - A(1) = 0$, and thus $6 \times T_{\mathrm{Hept}} \Delta_{\mathrm{add}} S = -7\varepsilon$, and $\Delta_{\mathrm{add}} S = -\dfrac{7\,\varepsilon}{6\,T_{\mathrm{Hept}}}$. The height of the barrier is expected to be $5\varepsilon - 5T_{\mathrm{Hept}} \Delta_{\mathrm{add}} S$, this gives a value of $0.8\,\varepsilon$, which, using the heptamerisation temperature for non-inserted pores found in our simulation ($kT_{\mathrm{Hept}} = 0.078\,\varepsilon$) is $\sim 10kT$. This simple barrier is not steep enough to discourage the formation of intermediates: for example, the free energy of a system with one dimer would only be $\sim 0.2\,\varepsilon$ higher than a system of monomers (which is $\sim 2\,kT$ at our heptamerisation temperature, and thus the probability of a dimer compared to a monomer is $\sim 3\,\%$). Other factors must therefore be important for this system.

Diffusion limited downhill reaction

If we assume that the reaction is entirely diffusion limited (the fastest we could assume, allowing bonding to occur rapidly), and that all sizes of clusters can diffuse at the same rate (this means that $k_i = k_{\mathrm{d}}$ for $1 \leq i \leq 6$), then we get the following rate equation:

$$\frac{\mathrm{d}[\mathrm{A}_7]}{\mathrm{d}t} = \frac{k_{\mathrm{d}}^6 [\mathrm{A}]^7}{(k_{\mathrm{d}}[\mathrm{A}] + k_{-1})(k_{\mathrm{d}}[\mathrm{A}] + k_{-2})(k_{\mathrm{d}}[\mathrm{A}] + k_{-3})(k_{\mathrm{d}}[\mathrm{A}] + k_{-4})(k_{\mathrm{d}}[\mathrm{A}] + k_{-5})} \tag{5.10}$$

This still has a strong dependence on concentration (experimentally, the surface concentration is of the order of 1×10^{13} particles/m^2, which is rather low).

Let us further simplify this by looking at the limit where we only see the forward reaction (*i.e.* set $k_{-i} = 0$ for $1 \leq i \leq 6$), and while still keeping the reaction diffusion limited. This gives us the rate equation:

$$\frac{\mathrm{d}[\mathrm{A}_7]}{\mathrm{d}t} = k_{\mathrm{d}}[\mathrm{A}]^2 \tag{5.11}$$

This would allow rapid formation of a pore, and it gives a significantly different rate dependence on concentration than the uphill model. However, such a kinetic pathway would not be expected to have monomers persisting for a long time, which is an important experimental feature of the assembly process [8].

With an insertion step

We modify our previous system to include an irreversible insertion step. A schematic of the free energy of the reaction coordinate is shown in figure 5.14. The inserted pore is referred to as ${\mathrm{A}_7}^{\star}$. Our model reaction pathway is

$$7\,\mathrm{A} \underset{k_{-1}}{\overset{k_1}{\rightleftharpoons}} \mathrm{A}_2 + 5\,\mathrm{A} \underset{k_{-2}}{\overset{k_2}{\rightleftharpoons}} \mathrm{A}_3 + 4\,\mathrm{A} \underset{k_{-3}}{\overset{k_3}{\rightleftharpoons}} \mathrm{A}_4 + 3\,\mathrm{A} \underset{k_{-4}}{\overset{k_4}{\rightleftharpoons}} \mathrm{A}_5 + 2\,\mathrm{A} \underset{k_{-5}}{\overset{k_5}{\rightleftharpoons}} \mathrm{A}_6 + \mathrm{A} \underset{k_{-6}}{\overset{k_6}{\rightleftharpoons}} \mathrm{A}_7 \xrightarrow{k_7} {\mathrm{A}_7}^{\star}. \tag{5.12}$$

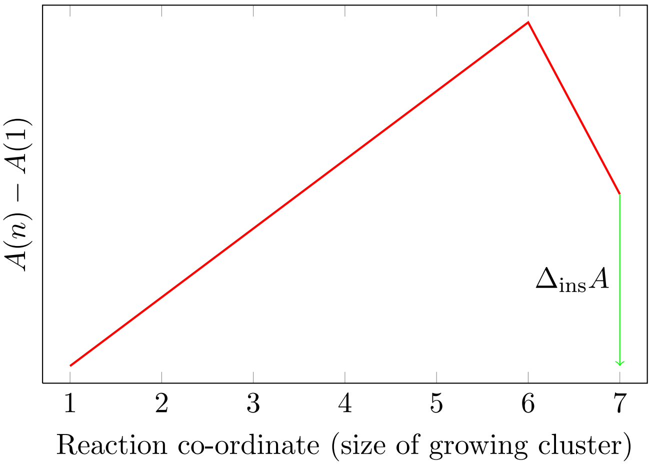

Figure 5.14: Schematic of the free energy of a reaction pathway including an insertion step as a function of cluster size at the transition temperature.

Using the steady state approximation, we are able to estimate the rate of reaction

$$\frac{\mathrm{d}[\mathrm{A}_7{}^\star]}{\mathrm{d}t} = \frac{k_1k_2k_3k_4k_5k_6k_7[\mathrm{A}]^7}{(k_2[\mathrm{A}]+k_{-1})(k_3[\mathrm{A}]+k_{-2})(k_4[\mathrm{A}]+k_{-3})(k_5[\mathrm{A}]+k_{-4})(k_6[\mathrm{A}]+k_{-5})(k_7+k_{-6})}. \tag{5.13}$$

We assume that the rate constant for insertion is fast compared to all other rates, thus $k_7 \gg k_{-6}$, which is a valid approximation, provided the membrane and pore rearrangement needed for insertion is minimal, and the insertion is steeply downhill. Our rate equation therefore simplifies to:

$$\frac{\mathrm{d}[\mathrm{A}_7{}^\star]}{\mathrm{d}t} = \frac{k_1k_2k_3k_4k_5k_6[\mathrm{A}]^7}{(k_2[\mathrm{A}]+k_{-1})(k_3[\mathrm{A}]+k_{-2})(k_4[\mathrm{A}]+k_{-3})(k_5[\mathrm{A}]+k_{-4})(k_6[\mathrm{A}]+k_{-5})} \tag{5.14}$$

By using similar approximations to those which we made for a steeply uphill reaction, we are able to estimate the overall rate of reaction.

We select the insertion free energy, $\Delta_{\text{ins}}A$, to be $-7\,\varepsilon$. We define the temperature of transition (heptamerisation followed by insertion), T_{Trans}, as the temperature at which the free energies of a system of only monomers and an inserted heptamer are equal. The free energy gained upon insertion must therefore be equal by definition to the free energy of the stepwise break down of the non-inserted heptameric cluster. This allows us to write: $\Delta_{\text{ins}}A = 7\,\varepsilon + 6\,T_{\text{Trans}}\,\Delta_{\text{add}}S$, and thus $\Delta_{\text{add}}S = -7\,\varepsilon/(3 \times T_{\text{Trans}})$.

At equilibrium, the forward and reverse rates are equal,

$$k_i[\mathrm{A}_i][\mathrm{A}] = k_{-i}[\mathrm{A}_{i+1}]. \tag{5.15}$$

The equilibrium constant, K, is defined in terms of the standard surface concentration,

$[\ominus]$, and the concentrations of the relevant species,

$$K_i = \frac{\dfrac{[\mathrm{A}_{i+1}]}{[\ominus]}}{\dfrac{[\mathrm{A}_i]}{[\ominus]}\dfrac{[\mathrm{A}]}{[\ominus]}}. \tag{5.16}$$

We can therefore rewrite our forward rate constant in terms of our equilibrium constant, reverse rate constant and standard concentration,

$$k_i = K_i \times \frac{k_{-i}}{[\ominus]}. \tag{5.17}$$

Simple thermodynamics states that the equilibrium constant can be calculated using

$$K_i = \exp\left(-\frac{\Delta A^{\ominus}(T)}{kT}\right). \tag{5.18}$$

In order to get a value for the forward rates, we need to define a value for the reverse rates, which we take to be equal for all reverse reactions, and we set this to k_r. We define our standard concentration as that of the initial monomer surface density in our simulation (surface density of $0.1\ /\ (\sigma^{\mathrm{LJ}})^2$), and then use the transition temperature, T_{Trans}, from our simulations that corresponds to the insertion energy being $7\,\varepsilon$ ($kT_{\mathrm{Trans}} \approx 0.17\,\varepsilon$). Our initial rate within this model is $\sim 10^{-20}\mathrm{particle}\,(\sigma^{\mathrm{LJ}})^{-2}k_\mathrm{r}$. Using the monomer radius of 1 nm [8], this would give an initial rate of $\sim 10^{-26}\mathrm{mol\,dm}^{-2}k_\mathrm{r}$. This would allow the reaction to occur on a reasonable timescale, but does not mean that we would not resolve intermediates over the proposed reaction pathway.

Activated pathway

An activated pathway means that we expect a change in free energy landscape during formation. An activated pathway may shift the transition state from a hexamer (as it is in the uphill pathway) or heptamer (as it may be in the insertion pathway) to any other sized cluster. An example reaction coordinate is shown in figure 5.15. As long as the downhill section is steep compared to the initial pathway, then the rate equation will depend solely on the formation of the n-sized transition state, and will be n-fold dependent on the concentration of the monomer. This is a possibility for the formation of the pore, and may give the dynamics observed experimentally. We look at an example of this in the next section.

The physical interpretation of an activated pathway could be one of two scenarios. Firstly, the activation could be a rearrangement of the internal structure of a protein

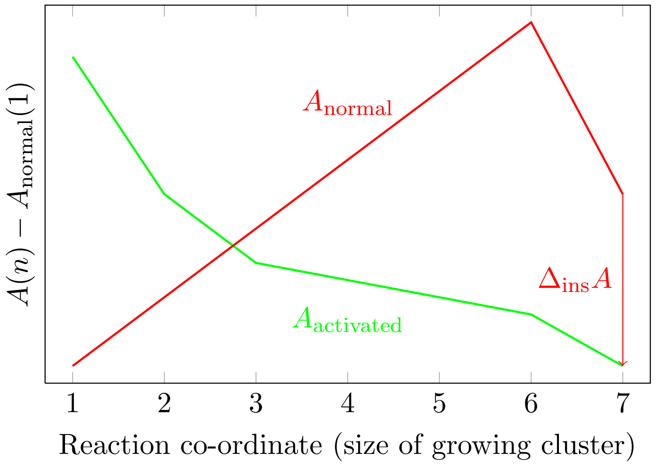

Figure 5.15: Schematic of the free energy of an activated reaction pathway as a function of cluster size at the transition temperature. The activated pathway is shown in green, and the insertion pathway in red.

to a higher energy state that allows binding to occur more strongly between monomers. Alternatively, the activation could occur by the pore inserting into the membrane at an earlier stage than we previously assumed (see section 5.6 for further details).

5.5 Including internal rearrangement in simulations

As mentioned in section 1.3.1, biological systems often have an internal degree of freedom which allows co-operativity to enhance the formation of certain clusters. We examine whether this may play a rôle in this system. As can be seen from the structure of the pore (figure 5.1), there is a strong possibility of internal degrees of freedom leading to co-operative binding when forming the heptamer, as the stem domains are clearly interlinked in the final structure. We therefore simulated an internal degree of freedom to test this theory. Using our simple Monte Carlo model with added flipping moves, we tried a number of different settings for internal degrees of freedom, most of which were unable to produce a yield of heptameric pores.

As any change in internal degrees of freedom means that the protein is no longer in its ground state conformation, we considered the effect of there being a change in the internal energy between the two states of our model. Our lower energy state had weak isotropic interactions of strength $0.1\,\varepsilon$ (following the example of the Wallace group [8]), while our narrow-patch state was $0.1\,\varepsilon$ higher in energy, with the same interactions as our original heptamer-forming model. The results are shown in figure 5.16. There is clearly a region where heptamers form, and this suggests that activated assembly via an internal degree of freedom is a possible mechanism in the formation of the heptameric pore. However, the

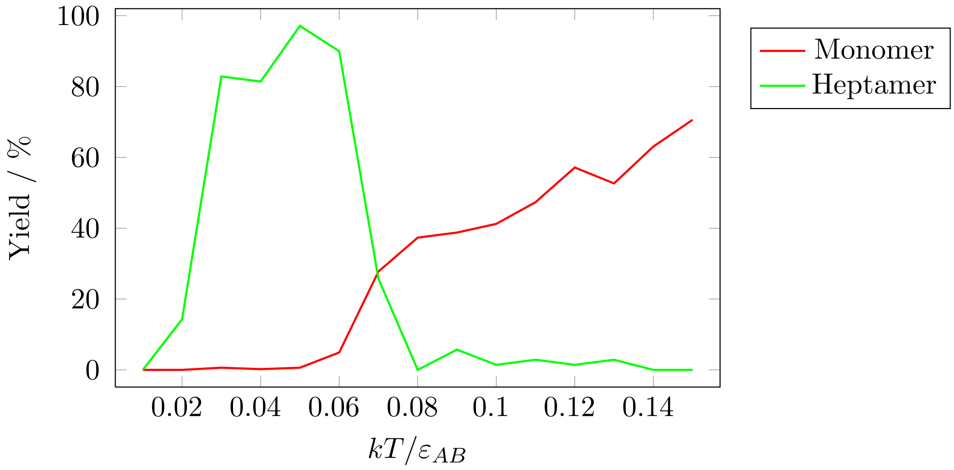

Figure 5.16: Yield of heptamers when starting from monomers in a system with an internal degree of freedom and a barrier to flipping. The isotropically attractive lower energy state had attractions of $0.1\,\varepsilon$, while the higher energy state (which was $0.1\,\varepsilon$ higher in energy than the isotropically attractive state) had the same patches and attractions as the standard model for this chapter.

yield of non-heptameric intermediates is more appreciable at the transition temperature than it was in our original model. There are many other possible ways in which this internal degree of freedom could manifest, though many that we tested did not give any yield of heptamers. In order to estimate the relevant states, some structural calculations or further biochemical data are needed.

5.6 Discussion of pore formation

We have shown using both simulations and models that there are many possible self-assembly energy profiles for pore formation. Further experimental data on the behaviour of a mutant α-haemolysin protein which is unable to form a pore that can insert into the membrane (for example, that considered in Parker's review [212]) may help the classification of pathways. If insertion is indeed the source of the activated behaviour, then the experiment would need to be performed at a much lower temperature to see the formation of the prepore, and intermediates should be observed on the simple uphill-type pathway we originally postulated.

When we consider pore formation in general, there are a few areas of interest:

The insertion of a pore into a membrane may occur at different stages during formation. For some pores, a single step insertion process of a completed prepore seems possible, though there would be a need for the phospholipid bilayer to undergo sudden rearrangement. While this may be possible for small rings, it seems unlikely for larger rings (as the channel of the central pore is, generally, free from any phospholipids, being hydrophilic in character

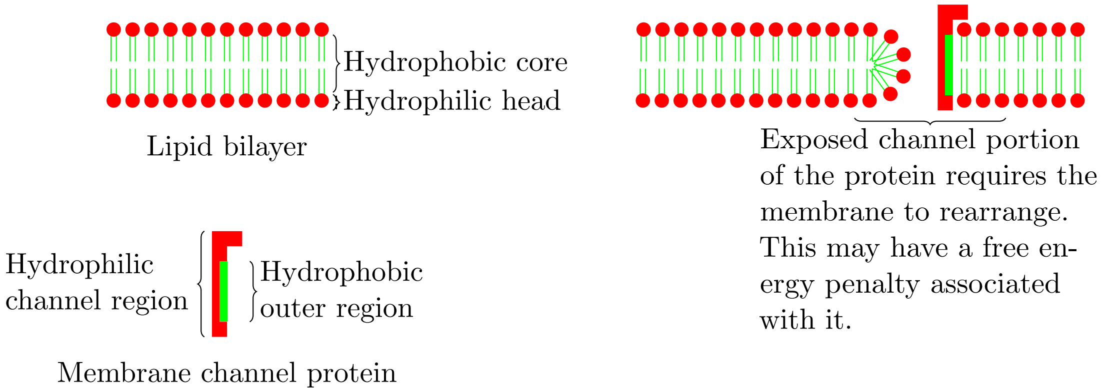

Lipid bilayer

Membrane channel protein

Exposed channel portion of the protein requires the membrane to rearrange. This may have a free energy penalty associated with it.

Figure 5.17: Two-dimensional schematic of hydrophobic and hydrophilic regions and strain induced by membrane rearrangement from early insertion.

[214]). The formation of partly-complete, inserted rings may require a pocket to be formed in the membrane, with the phospholipids bending back on themselves (demonstrated in figure 5.17). This would cost free energy. Alternatively, the part-formed channel could bend back on itself to expose only a hydrophobic outer surface to the membrane (this would have a strain energy of the ring associated with it, and may involve an internal degree of freedom).

While the α-haemolysin pore is only seen in heptameric form by photosensitive labelling and bleaching experiments [8, 215], a low yield of hexameric pre-pores has also been seen in certain mutants [212]. For larger pores, the formation of only one size of pore seems unlikely, especially as it would require a very narrow attractive region on a protein (equivalent to a small patch in our model), which may limit the rate of assembly greatly. Experiments into larger pores have shown that these pores are not monodisperse and may form stable structures with a range of sizes; for example, pneumolysin forms pores with 30-50 subunits [223].

Larger pores would also be expected to form via resolvable intermediates. If the barrier to formation is too high, then large pores are not expected to be able to form: a less steep slope means longer lifetimes for intermediates.

5.7 Conclusions

We have demonstrated that our model may be applied to more complicated biological systems, but that, without further experimental data, it is difficult to classify the behaviour seen in the current experiments.

Chapter 6

Re-entrant phase behaviour in self-assembling systems

6.1 Authorship

This chapter contains collaborative work. The section on the theoretical model (section 6.6) was completed in collaboration with my colleague, Aleks Reinhardt. The work was approximately equally shared between us, and was submitted for publication in a jointly authored paper [7].

6.2 Introduction

The characterisation of the self-assembly behaviour of the patchy particle systems studied in this thesis and previously in our group [46, 192] has raised some interesting questions about the fluid phase behaviour of these self-assembling systems. These questions can be most easily understood by thinking about the potential phase behaviour of these systems as a function of the patch width.

At wide patch widths, the potential is sufficiently close to isotropy that the phase behaviour is similar to the Lennard-Jones potential with liquid-vapour phase coexistence occurring at all temperatures below the critical point.

By contrast, at sufficiently narrow patch width, only a chemical equilibrium between a gas of monomers and a gas of self-assembled clusters is observed.

Of most interest in this chapter is the phase behaviour at patch widths between these two extremes, where self-assembly and liquid-vapour separation compete. Indications of what may occur come from results such as figure 6.1, where the nature of the product in dynamical simulation is shown as a function of patch width and temperature for a fixed

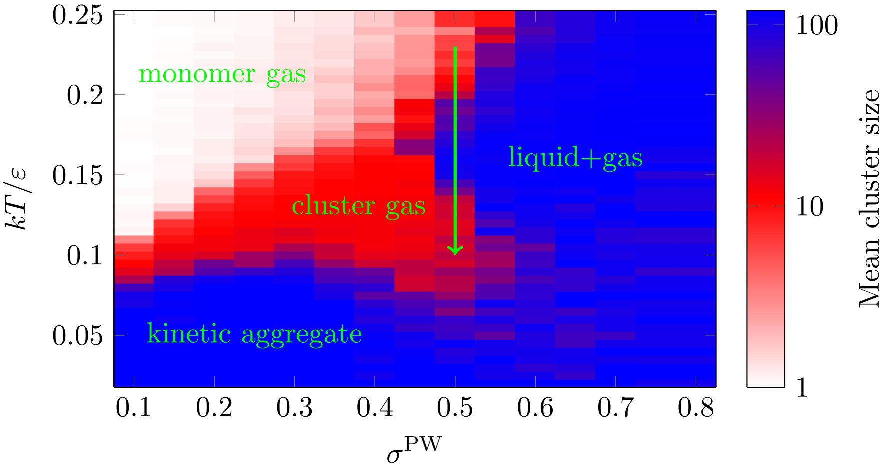

Figure 6.1: Plot of the mean cluster size for the uncentred icosahedral system showing simulation evidence of phase re-entrance. At the position of the green arrow it is clear that decreasing the temperature causes the mean cluster size to increase sharply (vapour of monomers to liquid), and then decrease (as a vapour of clusters is formed).

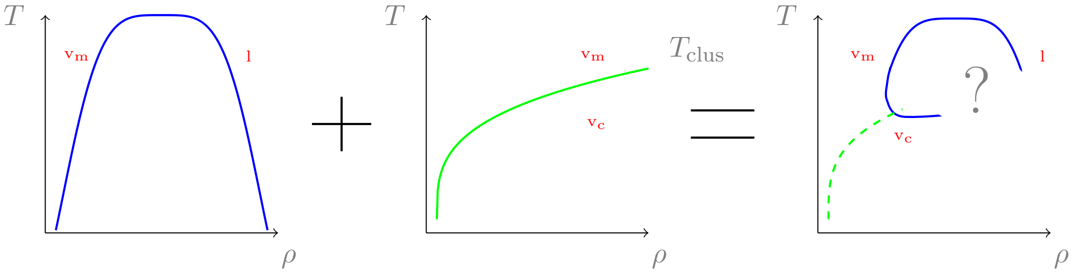

Figure 6.2: A schematic what we were trying to find. The combination of liquid-vapour coexistence and clustering will change the phase diagram, but we were unsure if the loop would close. T_{clus} is the clustering temperature, $\mathrm{v_m}$, $\mathrm{v_c}$ and l refer to a vapour of monomers, a vapour of clusters and a liquid, respectively.

density. At intermediate values of the patch width (*e.g.* $\sigma^{\mathrm{PW}} = 0.5$) the system passes from a monomer gas to a region of liquid-vapour coexistence to a gas of clusters as the temperature is decreased.

If the behaviour suggested in figure 6.1 is representative of the equilibrium behaviour, then the system exhibits re-entrant phase behaviour with the low temperature cluster gas phase presumably being stabilised by its lower internal energy compared to the higher entropy liquid seen at higher temperature. In this chapter we examine what happens at intermediate patch widths, *i.e.* when the clustering temperature is below but close to the monomer gas-liquid transition. Under these conditions we would expect a three way equilibrium to be set up (see figure 6.2), though we were unsure of what would happen at high density: would we see the re-entrance loop close off or other behaviour?

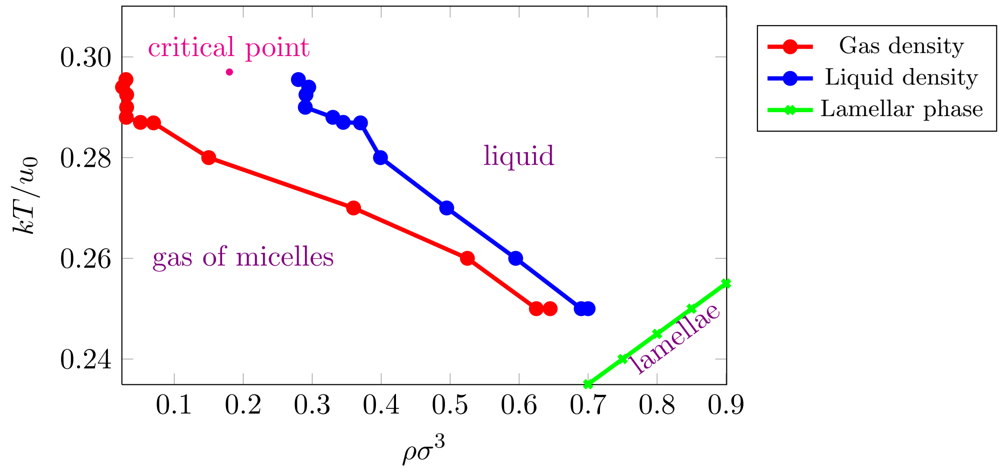

Figure 6.3: A phase-diagram for the Janus particle system as calculated by the Sciortino group [133]. Note the phase re-entrance.

6.2.1 Other research into phase re-entrance

Clear evidence of re-entrant phase behaviour has also been observed by Sciortino and co-workers in simulations of Janus particles with a single attractive patch. Certain parameter ranges lead to a competition between liquid-vapour phase separation and micelle formation [133, 134] (figure 6.3). In the phase diagrams that they computed, the standard behaviour of an ever-larger two-phase region in the T-ρ phase diagram as the temperature is decreased below the critical point is perturbed by the formation of micelles, shrinking the coexistence region and shifting it to higher density. However, what happens to this two-phase region at low temperature (does it end at a critical point, or persist to zero temperature?) is unclear, because the system instead formed a lamellar phase. In this chapter, our aim is to address such questions by attempting to calculate complete phase diagrams for the fluid using both simulations and a simple theoretical model of a system that can both self-assemble into clusters and exhibit liquid-vapour phase separation.

Re-entrant phase behaviour is of course not limited to patchy particle systems. Indeed, the classic example of such behaviour is the nicotine-water mixture, which exhibits a closed loop in its temperature-composition phase diagram with both an upper and a lower critical solution temperature [222]. Reducing the temperature whilst keeping the mole fraction of nicotine constant first leads to phase separation, and then to remixing driven by nicotine-water association. Liquid-liquid immiscibility leading to such phase diagrams is often attributed to hydrogen bonding [224, 225]. Complete mixing will occur if the temperature is sufficiently high so as to overcome the enthalpic incompatibility of mixing. However, at sufficiently low temperatures, miscibility results from the association of the components to form a complex. Similar considerations have been used to explain re-entrant phase behaviour seen in polymer aggregation in water [226], in non-ionic micellar solutions [227] and in nematic re-entrant phases in liquid crystals [228].

Although the systems of interest in this chapter are not two-component mixtures, unlike most systems where re-entrant phase behaviour has previously been observed, the behaviour can nevertheless in some ways be analogous to such systems. The similarities arise due to the mixing of two states of the same component (monomers and associated clusters), rather than two components proper. An example of re-entrant phase behaviour as a function of temperature in a system with a single effective density is ionic association [229, 230]. When dielectric constants are large, phase separation is driven by the usual enthalpic incompatibility. With low dielectric constants, coulombic-style phase separation between a low density (neutral cluster) structured gas phase and a high density (free ion) liquid phase occurs due to the long-range ionic interactions between the ionic components, rather than between the ions and the solvent. The formation of ionic clusters is somewhat similar to the cluster formation in the self-assembling patchy particle system that we focus upon, and so the re-entrant behaviour (in simple theories of these ionic systems, banana-shaped coexistence curves are found [229]) may have similar origins.

6.3 Computation of phase diagrams using the Wang-Landau algorithm

In this section we describe our attempts to compute the phase diagrams for some of the self-assembling patchy-particle systems that have been studied in our group. There are a plethora of methods to compute the liquid-vapour phase separation, such as the Gibbs ensemble and thermodynamic integration. We chose initially to use the Wang-Landau grand canonical approach described in section 2.2.6. However, applying it to our self-assembling systems was considerably more challenging than the isotropic Lennard-Jones test case considered there, because of the competition between self-assembly and phase separation. It is noteworthy that the Gibbs ensemble phase diagram calculations by the Sciortino group took an extremely long time to converge for their Janus particle systems [133].

6.3.1 Octahedra

The first system for which we attempted to compute the phase diagram was a system of octahedra without a torsional constraint in the interparticle potential. The system was run using the Wang-Landau algorithm in the grand canonical ensemble. We chose this system because it was the smallest cluster for which clear evidence of re-entrant phase behaviour

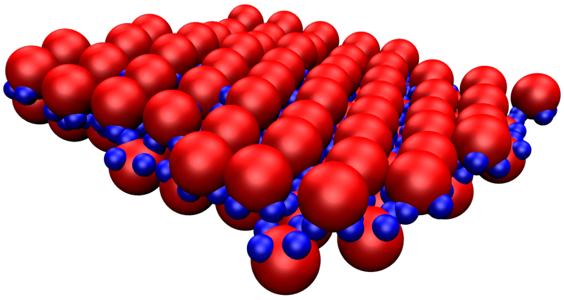

Figure 6.4: Example of the *egg-box* 2-D crystalline lamellar structure formed by octahedron-forming particles in simulations at low temperature (this is from $kT = 0.08\varepsilon$, $\sigma^{\mathrm{PW}} = 0.7$, with $N/V = 0.25(\sigma^{\mathrm{LJ}})^{-3}$, though it also occurred for many other patch widths and temperatures).

had been observed in yield plots equivalent to figure 6.1 [112], the hope being that this would facilitate equilibrium sampling, as the self-assembly of the small target clusters should be faster than for some of the larger clusters we have considered. The approach was successful in mapping out the liquid-vapour phase separation for parameter ranges where the critical point lies significantly above the clustering temperature. However, equilibrium was particularly difficult to obtain in regions where re-entrance was expected. This was partly due to the longer time scales associated with clustering, but was further exacerbated by the presence of other low-temperature phases that we had not envisaged from our initial investigations. In particular, as with the Janus particles studied by Sciortino [133], we observed lamellar structures at low temperature. To a first approximation, our particles also have an attractive side, and a non-attractive side. In the case of the the octahedron-forming particles, an ordered two-dimensional crystalline lamellar phase can form in which a particle's four patches point directly at a patch of its four neighbours on the opposite face of the lamella (figure 6.4). Also, at intermediate temperatures, the residual attraction between complete clusters can cause them to form linear aggregates (figure 6.5).

6.3.2 Icosahedra

In order to look at the phenomenon in a system less likely to form crystalline structures that spanned the box, we chose to study icosahedron-forming particles. Here, the particles have five patches, and thus they are unlikely to form ordered lamellar phases within our box. The required 5-fold two-dimensional symmetry would be incommensurate with our periodic (cubic) box. Though it would be possible to form a crystal based on square-triangle tiling, it is not feasible to use this lattice in a lamellar phase (unlike the square lattice of the octahedral system).

Again, we expected to find three phases: a gas of monomers, a gas of clusters, and a liquid. The simulations took several months to converge, and we report our findings here.

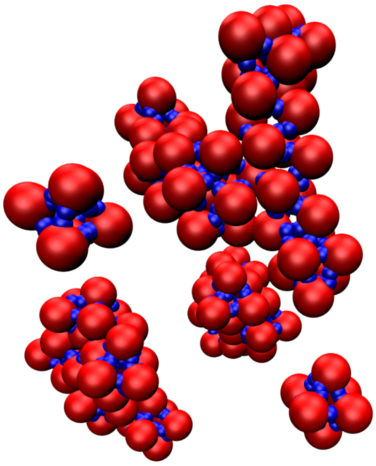

Figure 6.5: Columnar aggregates formed at narrow patch widths and low temperature (this is from $kT = 0.1\varepsilon$, $\sigma^{\mathrm{PW}} = 0.5$, with $N/V = 0.25(\sigma^{\mathrm{LJ}})^{-3}$).

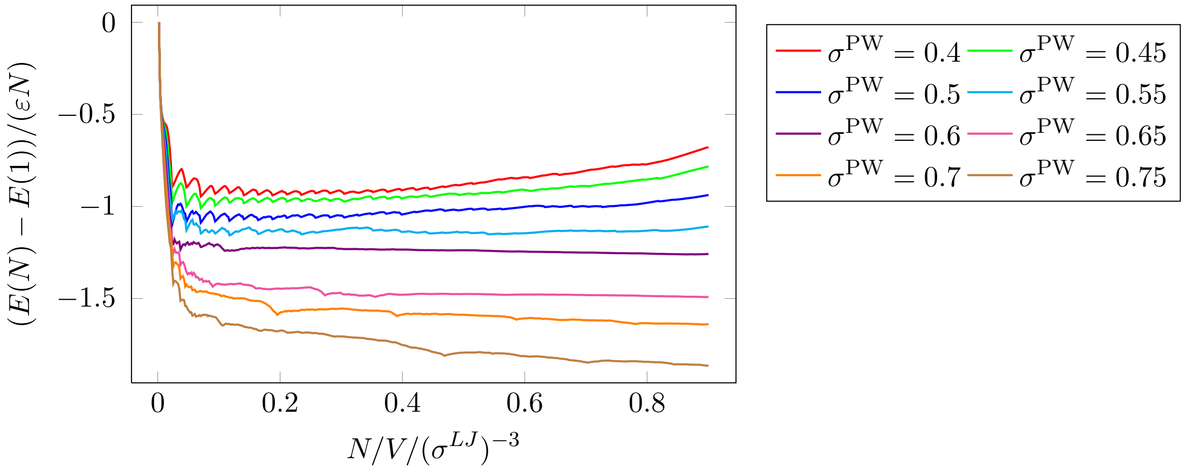

Figure 6.6: Effect of changing patch width on the energy per particle. Note the regular undulations at low patch widths, this is an effect of a non-infinite system size. $kT = 0.12\,\varepsilon$.

We have found several limitations with the use of the algorithm on this type of system (see section 6.3.3). The majority of issues were the result of clustering.

Looking in detail at the free energy per particle (figure 6.6) at a set temperature (here $kT = 0.12\,\varepsilon$) at different patch widths, the presence of clustering at narrow patch width can be clearly seen by the regular oscillations to the free energy per particle visible in the upper curves. The troughs correspond to an integer number of clusters being formed. In an infinite system, we would not expect these to be visible and, in fact, when running the system at a different box size and looking at the density profile, we find that the peaks move in density space. At wider patch widths, the lowest points in energy no longer correspond to icosahedral clusters, but are in fact the *magic number* sized clusters seen for Lennard-Jones particles [231] (the first two occur at densities corresponding to 13 and 19

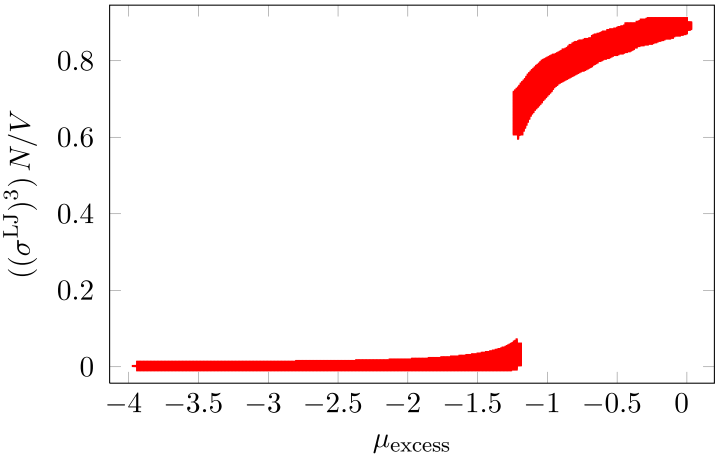

Figure 6.7: A plot of the densities which have a probability greater than 0.01. This is in a region with a clear phase transition between monomeric gas and liquid phases. Taken from a run at $kT = 0.3\,\varepsilon$, with $\sigma^{\mathrm{PW}} = 0.75$.

particles). This change of behaviour occurs at wide patch widths as the attractions tend towards an isotropic (thus Lennard-Jones) type of behaviour. We consider the issue of clustering and density in section 6.3.3.

The effects of clustering on a finite system mean that it is difficult to distinguish phase changes and clustering when looking at the probability of densities as a function of chemical potential. Figure 6.7 shows the densities which had a probability above 0.01 as a function of excess chemical potential (μ_{excess}) at high patch width ($\sigma^{\mathrm{PW}} = 0.75$) and high temperature ($kT = 0.3\,\varepsilon$) below the critical temperature and above the clustering temperature. This clearly shows a discontinuous change between a low and high density phase, with a coexistence region. On the other hand, figure 6.8 shows a similar plot but at low temperature ($kT = 0.1\,\varepsilon$) and patch width ($\sigma^{\mathrm{PW}} = 0.4$) in a region below the clustering temperature. The lower temperature system clusters and this leads to certain densities being more likely, even though there is no phase change. This makes it difficult to identify actual phase changes. For intermediate values of patch width we are able to assign liquid-vapour coexistence at higher temperatures, but when we decrease the temperature below the clustering temperature, we see complicated behaviour that is a simulation artefact due to our box size. When clustering is present in the system, the data also take significantly longer to converge (see section 6.3.3).

Grand canonical Wang-Landau results

We were able to find the coexistence points at high temperatures for all systems with a patch width above 0.5. However, at lower temperatures, the difficulty of differentiating equilibrium clustering from true phase transitions owing to clusters and box size effects

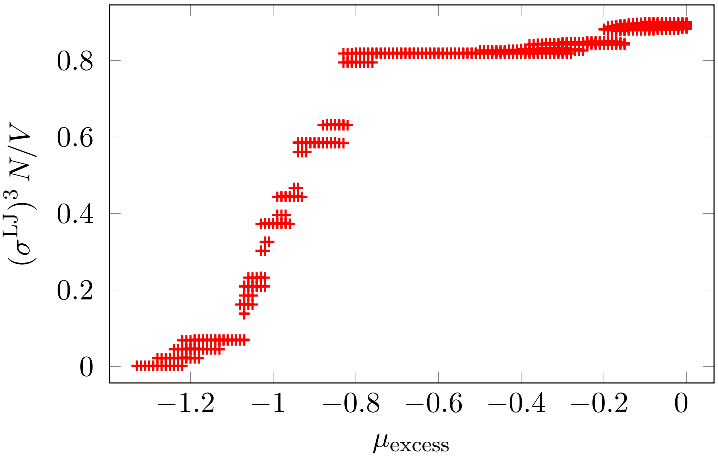

Figure 6.8: A plot of the densities which have a probability greater than 0.01. This shows the difficult in locating a phase change when clustering is present. The regions of stability occur around an integer number of clusters. Taken from a run at $kT = 0.1\,\varepsilon$, and with $\sigma^{\mathrm{PW}} = 0.4$.

(see section 6.3.3) meant that we had to find alternative methods.

The results obtained from the grand canonical Wang-Landau algorithm are shown later in section 6.4 (figure 6.10) along with the Gibbs ensemble results. These are clearly commensurate, though the Gibbs algorithm was able to predict to lower temperatures and was not as strongly affected as the grand canonical Wang-Landau algorithm by clustering.

6.3.3 Issues with the algorithm

There are several unsolved issues with our implementation of the algorithm. We look at these in this section, and propose some possible solutions which will require future work.

Clustering

The grand canonical Wang-Landau algorithm in its current form is very slow in conditions where clustering occurs. The main issue is that forming a cluster takes a significant time (up to 10^5 cycles, depending on the conditions used), especially when there are many kinetic aggregates that may trap the system; this produces noise in our sampling and slows convergence. It is possible that the use of two-dimensional Wang-Landau sampling across energy and number of particles in the grand canonical ensemble will give better results and sampling. Runs would need to be carried out at a number of different densities using canonical-ensemble Wang-Landau simulations, then these results could be used as a starting point for a two-dimensional energy-density Wang-Landau simulation: the system would be able to sample energy states at each density easily and thus would be far less likely to become trapped.

Discerning discontinuous phase changes and continuous changes of structure

When looking at our grand canonical Wang-Landau simulation data, it is necessary to be able to tell the difference between a true first-order phase transition (for example, from a gas to a liquid), and a continuous change (for example, from a monomer gas to a cluster gas). Generally, a first-order phase transition can be easily identified in chemical potential/density/probability space, as it shows a discontinuous jump from one density to another with a small change in chemical potential. This can be clearly seen in figure 6.7. However, when we have clustering, we often see apparent jumps that do not correspond to discontinuous phase changes (figure 6.8).

The simplest way to decide whether a phase change is continuous or discontinuous is to change the size of the simulated system. True phases are unaffected by the number of particles in the system, and depend only on the density, while clustering or other equilibrium transitions would depend on the number of particles present, so the densities of the transitions change when the volume is changed.

Small simulations and the effect of box size

As Ouldridge *et al.* have shown [232], when attempting to extract bulk properties from a small simulation, that there are a number of issues which must be taken into account. We consider these issues in more detail in appendix B.3. We also should note the issue of the limited size of the simulation box meaning that, within the region where clusters are formed, predictions of free energy show an unphysical undulating pattern (figure 6.6).

6.4 Gibbs ensemble approach

As there were some clear issues present in our Wang-Landau approach, we ran a number of Gibbs ensemble runs (see section 2.2.7 for further information) using the above results as a starting point to help distinguish the effects of clustering and to help locate the limits of phase transitions. We looked at the parameters for which we expected there to be a reasonable region for the stability of a liquid. We chose patch widths of 0.6 and 0.7 radians, and ran a series of Gibbs ensemble runs starting with the two boxes each having a number density of 0.25 / $(\sigma^{\mathrm{LJ}})^3$. The time evolution of density of an example run is demonstrated in figure 6.9.

The results obtained by the algorithm (figure 6.10) were only consistent within certain regions (each set of points is the average from 5 runs of 10^9 Monte Carlo steps). At low temperature the formation of large aggregates often caused trapping and thus we were unable to find the thermodynamic coexisting phases. The re-entrance is clear from

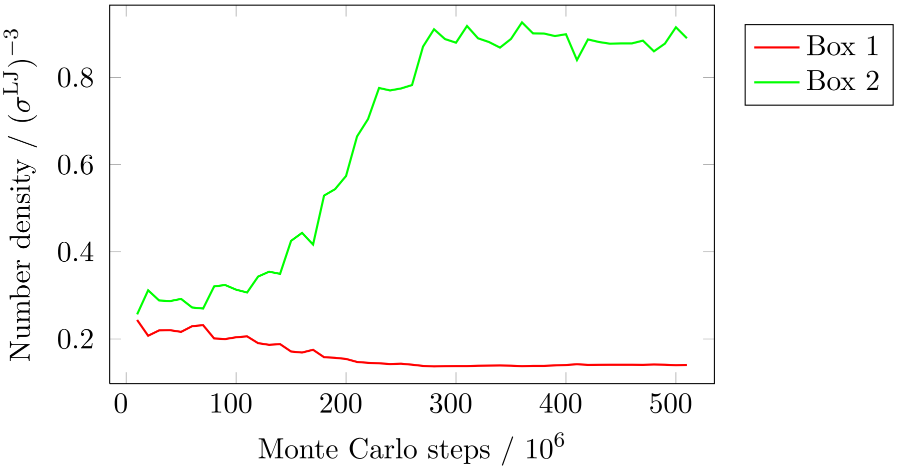

Figure 6.9: Evolution of density with respect to time in the two boxes of an example of a single Gibbs ensemble simulation. $kT/\varepsilon = 0.12$, $\sigma^{\mathrm{PW}} = 0.6$.

the diagram, and there is clear agreement between our grand canonical Wang-Landau algorithm and Gibbs ensemble results. For the patch width of 0.6 (figure 6.10(a)), at a number density of 0.15 σ^{-3}, the Gibbs ensemble results predict that the coexistence region occurs between $kT \approx 0.13\,\varepsilon$ and $kT \approx 0.25\,\varepsilon$. This is also in agreement with our simulation data (figure 6.11).

6.5 Dynamic simulation approach

To supplement our understanding and lend more support to the existence of phase re-entrance we looked at the results from dynamic Monte Carlo runs. We started simulations from a box full of formed icosahedra and looked at the proportion formed as a function of patch width. The results from this approach can be seen in figure 6.11. We will only see aggregates where they are thermodynamically stable. It is clear that we have correctly assigned the kinetic aggregate region in figure 6.1. and also the liquid. We also notice that we have to superheat clusters at low patch width to see dissociation, due to slow kinetics. The liquid-vapour/cluster boundary is slightly shifted in favour of clusters. These results clearly show phase re-entrance.

Plots run at different densities were strongly affected by the surface energy required in order to form a liquid / gas mixture. Hence our use of the Gibbs ensemble method (as explained above).

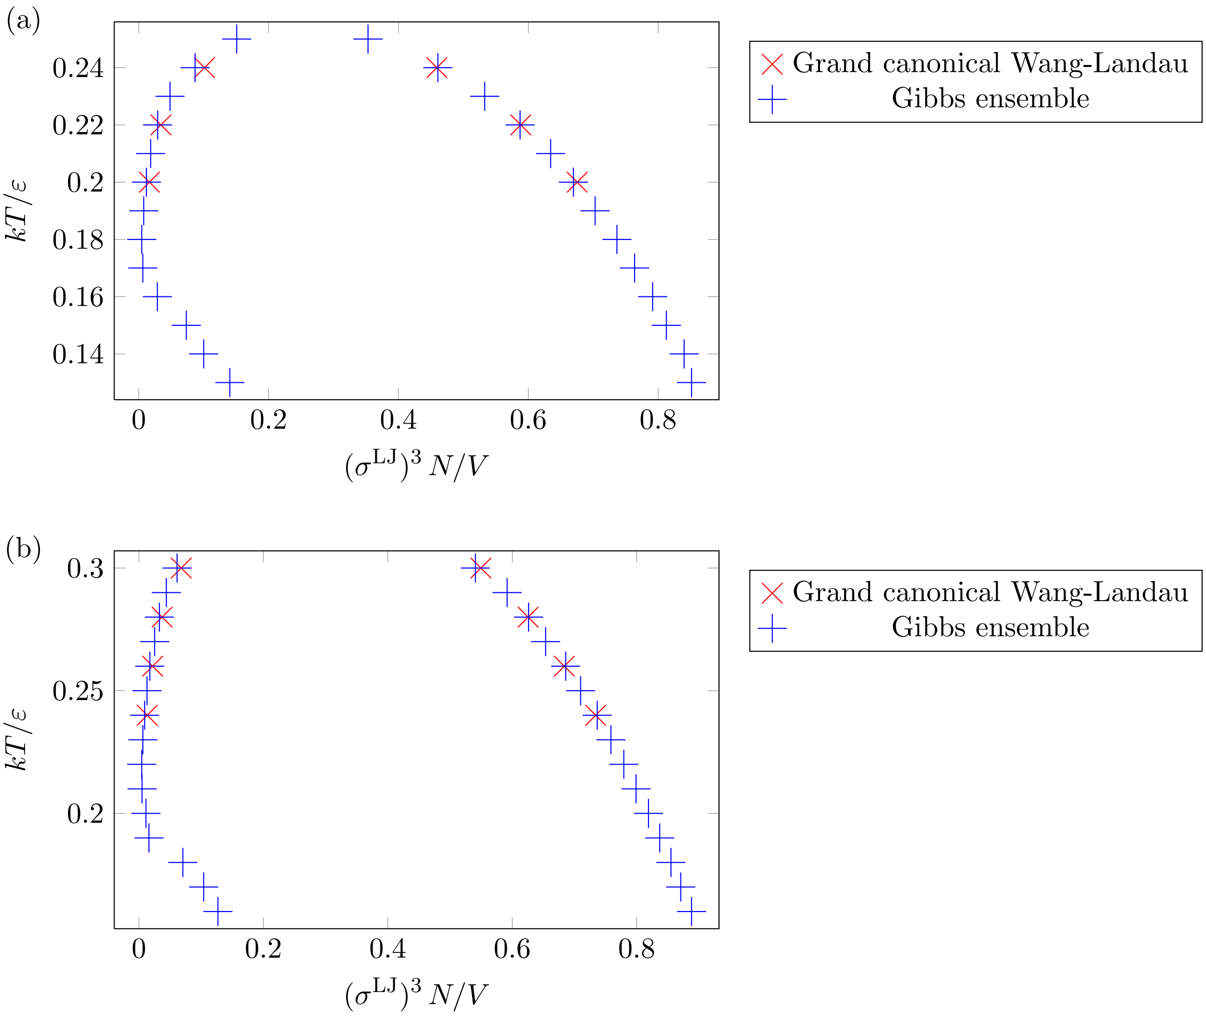

Figure 6.10: Temperature-density phase diagram for the icosahedral system for patch widths of (a) 0.6 radians and (b) 0.7 radians using the Gibbs ensemble (blue) and the Wang-Landau algorithm in the grand canonical ensemble (red).

6.6 *Ab initio* approach

This section of work was done in close collaboration with Aleks Reinhardt, Oxford. The work was completed jointly, with all parts being done equally by each of us. Mathematica was used for numerical calculations using routines written by both of us, and the mathematical derivations were carried out collaboratively.

Owing to the difficulties encountered in trying to reach equilibrium in our simulations in the regions of parameter space where we expected to see re-entrance, it was decided that it would be worthwhile creating a simple mathematical model which would be able to be used to delimit the stability of phases. As already mentioned, this has been accepted for publication [7], and an expanded version of the paper makes up this section.

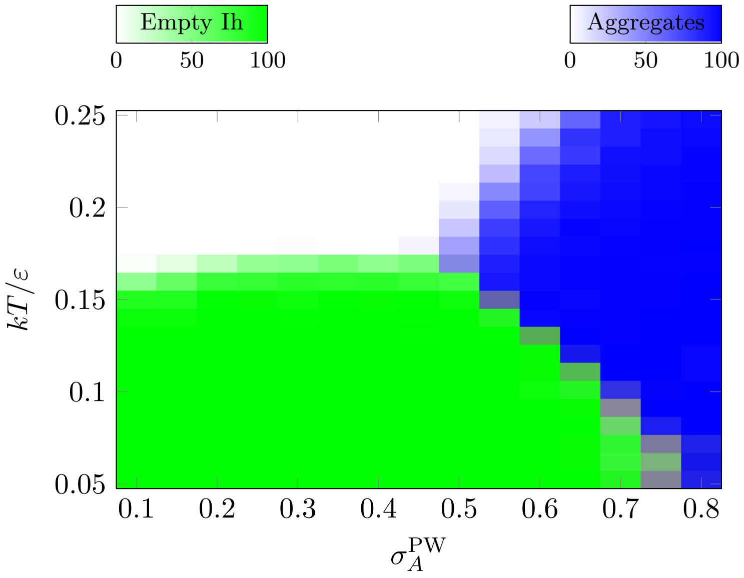

Figure 6.11: Yield of aggregates and clusters starting from formed clusters using different patch widths. The white regions corresponds to monomeric gas, the green to a gas of clusters and the blue to a liquid-vapour coexistence region. Other settings are given in table 3.1 for the IhO system.

6.6.1 Model

The aim of this research was to use a simple model system that would help to show the existence of phase re-entrance in our patchy particle system, and that could be used to estimate their phase diagram. The system can potentially form a vapour of monomers, a liquid of monomers and a gas of non-interacting monodisperse clusters. We initially modelled the uncentred icosahedral cluster forming system (as seen in section 3.5.1), but the approach can easily be generalised to other clusters, as in section 6.6.7.

We chose to model the chemical equilibrium between clusters and monomers using explicit molecular partition functions, and to treat the vapour-liquid equilibrium as a van der Waals fluid. Before investigating the three state system, we first look at just the monomer-cluster equilibrium in section 6.6.2 and the pure van der Waals fluid in section 6.6.4. We then consider our 'self-assembling van der Waals fluid' in section 6.6.5.

6.6.2 Monomer-cluster equilibrium

Partition functions

We first constructed the molecular partition function of an icosahedral cluster using statistical thermodynamics (in a manner similar to that done previously for a simple model of tetrameric protein complexes [4, 5]).

Using standard notation, the overall molecular partition function of a cluster is given by:

$$q_{\text{clus}} = q_{\text{trans, clus}} \times q_{\text{rot, clus}} \times q_{\text{vib, clus}} \times \exp\left[-E_{\text{clus}}/kT\right], \tag{6.1}$$

where E_{clus} is the ground state energy of the cluster.

The monomers have both rotational and translational degrees of freedom, but we did not consider their attractions initially. Therefore, the molecular partition function of the monomer is given by:

$$q_{\text{mon}} = q_{\text{trans, mon}} \times q_{\text{rot, mon}} \tag{6.2}$$

We use σ to denote the hard-core diameter of our particle, and, for systems that include our angular-modified Lennard-Jones potential, we set $\sigma^{\text{LJ}} = \sigma$.

Our model is designed to simulate patchy colloids, which are classical objects. Thus we ignore all small-scale quantum effects and use classical partition functions throughout this chapter.

Translational partition functions The translational partition function is given by:

$$q_{\text{trans}} = \frac{V_{\text{free}}}{\Lambda^3} \tag{6.3}$$

Where Λ is the de Broglie wavelength, and V_{free} is the free volume in which a particle has free movement. There are many ways to estimate the free volume term.

Estimation of free volume terms In order to represent the short-range repulsions between species in the system, we use an excluded volume term in the translational molecular partition functions of the free volume approximation [229]. This has the form:

$$V_{\text{total}} - V_{\text{free}} = V_{\text{exc}} = \sum_{\text{species, } i} B_i N_i \tag{6.4}$$

N_{mon} and N_{clus} refer to the number of monomers and the number of clusters in the system, respectively, while B_i is the volume exclusion coefficient of each species. We denote the monomer coefficient as $B_{\text{mon}} = b$ and the cluster coefficient as $B_{\text{clus}} = B$. Although there are more accurate functions which describe the effects of excluded volume better than the simple summation we have chosen [233], we deliberately wish to keep the partition functions as simple as possible, as we are interested in the *generic* properties and phase behaviour resulting from the competition between self-assembly and phase separation, rather than the detailed behaviour of a specific system. Moreover, we want a form consistent with the van der Waals description of the fluid (which is important for section 6.6.5).

We could match the volume exclusion coefficients (B and b) values to the hard-sphere second virial coefficients; however, this only works well at low densities where higher-order virial coefficients can be neglected (and actually leads to an unphysical range of possible densities). Instead, we chose B and b values to correspond to the reciprocals of the maximum densities of the respective species.

By using the free volume approximation, we chose not to have a cross term (in the case of the clusters this would be $B_{\text{mon-clus}} N_{\text{mon}}$). This term would be appropriate if we were matching to exact hard-sphere second virial coefficients, and in this case $B_{\text{mon-clus}}$ can be related to B and b by the expression:

$$B_{\text{mon-clus}} = [(B_{\text{clus}}^{1/3} + B_{\text{mon}}^{1/3})/2]^3 \tag{6.5}$$

This equation is known as the Lorentz expression [234]. However, this form again leads to unphysical behaviour at higher densities, where it overestimates the volume excluded.

As noted above, we choose the free volume coefficient b to be $1/\rho_{\text{max}}$ [229], allowing the system to span a suitable density range. Throughout this chapter, all densities ρ refer to the number densities, $\rho = (N/V)$, rather than mass densities. We set the maximum density achievable in the system to σ^{-3}, and hence $b = 1$. In our figures we also use the reduced number density $\rho^* = \rho \times \sigma^3$, and volume per particle $v = V/N$.

If we consider the maximum density of a system made purely of monomers to equal the maximum density of a random close-packed system, we can establish the packing fraction, ϕ, (the proportion of free space filled):

$$\begin{aligned} \phi &= \frac{N \cdot V_{\text{mon}}}{V_{\text{tot}}} \\ &= \rho \times V_{\text{mon}} \\ &= \rho \times \frac{4\pi}{3} \times \left(\frac{\sigma}{2}\right)^3 \\ &= \frac{\rho \times \sigma^3 \times \pi}{6} \end{aligned} \tag{6.6}$$

Using a randomly close-packed packing fraction, ϕ_{RCP} of ≈ 0.64 [235], we get:

$$\begin{aligned} b &= \frac{1}{\rho_{\text{max}}} = \frac{\sigma^3 \times \pi}{6\,\phi_{\text{RCP}}} \\ &= 0.82\,\sigma^3 \end{aligned} \tag{6.7}$$

This result suggests that our choice for b is a reasonable approximation. It also means that the van der Waals critical density (section 6.6.4) is in agreement with that of the Lennard-Jones fluid.

The relative values of B and b should be proportional to the cluster and monomer sizes. To determine the size of the cluster relative to that of the monomer, we assume that clusters are approximately spherical objects. The radius of a sphere which goes through all of the points of an icosahedron [236] with edge f is:

$$r = (f/4)\sqrt{10 + 2\sqrt{5}} \approx 0.95f \tag{6.8}$$

If the icosahedron is built of monomeric spheres of diameter σ (as it is in our case), then $f = \sigma$. However, we also need to take into account the fact that these monomers protrude out of the icosahedron itself, and so another hard-core radius of these constituent spheres, $\sigma/2$, must be added to the circumradius of the icosahedron to estimate the overall cluster radius R, giving $R = 1.45\,\sigma$ (see figure 6.12).

The volume occupied by this sphere relative to the volume occupied by a single particle is:

$$\begin{aligned}\frac{V_{\text{clus}}}{V_{\text{mon}}} &= \frac{(1.45\,\sigma)^3}{(\sigma/2)^3} \\ &= 24.4\end{aligned} \tag{6.9}$$

We therefore have a suitable measure of how much bigger the icosahedral cluster is compared to a single patchy particle, and therefore can estimate the relative sizes of the excluded volume coefficients as $B/b \approx 24$. We consider this case and also a few other values of b in order to characterise more fully the types of phase diagram that may arise. $B/b = 24$ is perhaps the most physical of all the values of B which we have used and it is specific for the icosahedron described. Other values would be relevant to different packing arrangements of monomers. This study mainly considers $B = 24b$, $B = 12b$ and, for the purpose of investigating the topology of curve changing, some values in between. In section 6.6.3, we also consider the ideal case ($B = b = 0$).

For the cluster, we can therefore use

$$\begin{aligned}q_{\text{trans, clus}} &= \frac{V - BN_{\text{clus}} - bN_{\text{mon}}}{\Lambda^3_{\text{clus}}} \\ &= \frac{48\sqrt{6}(\pi mkT)^{3/2}}{h^3}(V - BN_{\text{clus}} - bN_{\text{mon}}),\end{aligned} \tag{6.10}$$

for the monomer, we have

$$\begin{aligned}q_{\text{trans, mon}} &= \frac{V - BN_{\text{clus}} - bN_{\text{mon}}}{\Lambda^3_{\text{mon}}} \\ &= \frac{(2\pi mkT)^{3/2}}{h^3}(V - BN_{\text{clus}} - bN_{\text{mon}})\end{aligned} \tag{6.11}$$

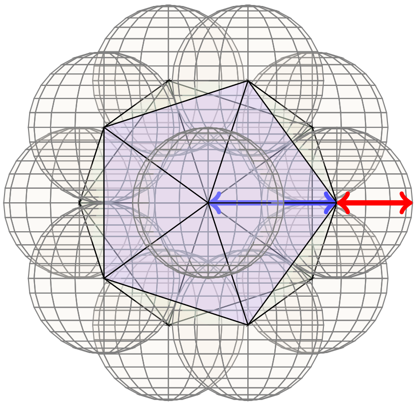

Figure 6.12: In an icosahedron, provided the system is incompressible, we calculate the radius of the circumscribed sphere and then add a further radius of the particle, as shown in red in the figure. This gives a physically relevant value of B/b to be approximately 24, as used throughout most of this work.

where Λ is the de Broglie thermal wavelength, the mass of the icosahedron is $12m$ (m being the mass of a monomer), and N_{mon} and N_{clus} refer to the number of monomers and the number of clusters in the system, respectively.

Rotational partition function The rotational partition function for a spherical rotor is given by [222]

$$q_{\text{rot, spherical}} = \frac{\sqrt{I^3\pi}}{s} \times \left(\frac{8\pi^2 kT}{h^2}\right)^{3/2}, \tag{6.12}$$

where I is the moment of inertia about an axis, and s is the symmetry number. Using point masses we found the moment of inertia about the principal axes of an icosahedron to be

$$I_{\text{clus}} = \frac{80m\sigma^2}{5-\sqrt{5}}, \tag{6.13}$$

where, as before, m is the monomer mass and σ is the diameter of the monomers. The symmetry number (s) of an icosahedral cluster is 60 (it has 60 different equivalent rotamers). We treat the three rotations as equal, and can therefore calculate the rotational partition function for the icosahedron

$$\begin{aligned} q_{\text{rot, clus}} &= \frac{\sqrt{I_{\text{clus}}^3\pi}}{60} \times \left(\frac{8\pi^2 kT}{h^2}\right)^{3/2} \\ &= \frac{128\pi^{7/2}}{3}\sqrt{1+\frac{2}{\sqrt{5}}}\left(\frac{mkT\sigma^2}{h^2}\right)^{3/2}. \end{aligned} \tag{6.14}$$

For a monomer the moment of inertia is:

$$I_{\text{mon}} = \frac{2}{5}m\sigma^2 \tag{6.15}$$

Thus, overall for the rotational partition function of the monomer, we have

$$\begin{aligned} q_{\text{rot, mon}} &= \frac{16\sqrt{2}\pi^{7/2}\left(I_{\text{mon}}kT\right)^{3/2}}{sh^3} \\ &= \frac{64\pi^{7/2}\sigma^3(mkT)^{3/2}}{25\sqrt{5}h^3}, \end{aligned} \tag{6.16}$$

note that the symmetry number, s, of a monomer is 5, since there are five identical patches on each monomer.

Vibrational partition function The monomers in the cluster each have 6 degrees of freedom (3 translational and 3 rotational). Thus, we have 6×12 degrees of freedom overall. We have already used 3 degrees of freedom in our translational partition function, and 3 in our rotational partition function. This leaves us with $(6 \times 12 - 6)$ overall vibrational modes of the cluster. To keep the model simple, we decided that $(3 \times 12 - 6)$ of these degrees of freedom probe the radial part of the potential, and 3×12 probe the orientational degrees of freedom, and that all the modes of each type have the same spring constant, and therefore the same angular frequency (ω), in a similar fashion to the Einstein model of solids [237]. This level of simplification is reasonable, given some of the other approximations already used. Hence, we can write the vibrational partition function as:

$$q_{\text{vib}} = \prod_i q_{\text{vib},\,i} \approx \left(\frac{kT}{\hbar\omega_{\text{RAD}}}\right)^{30}\left(\frac{kT}{\hbar\omega_{\text{ANG}}}\right)^{36}. \tag{6.17}$$

To obtain expressions for ω_{RAD} and ω_{ANG}, we consider the same interparticle potential as we use for simulations, which is given in section 2.1, and assume that these two frequencies are simply related to the second derivatives of this pair potential with respect to the interparticle distance and particle orientation at the equilibrium geometry of the pair interaction (for further information see appendix C.2).

In our system we therefore have:

$$\omega^2_{\text{ANG}} = -\frac{\varepsilon}{I_{\text{mon}}}\left(\frac{\partial^2 U^{\text{ANG}}}{\partial\theta^2_{kij}}\right)_{\left.\begin{matrix}\theta_{kij}\\ \theta_{lji}\end{matrix}\right\}=0} = \frac{\varepsilon}{I_{\text{mon}}(\sigma^{\text{PW}})^2} \tag{6.18}$$

and:

$$\begin{aligned} \omega^2_{\text{RAD}} &= \frac{2}{m}\left(\frac{\partial^2 U^{\text{LJ}}}{\partial r^2}\right)_{r=2^{1/6}\sigma} \\ &= 72 \times 2^{2/3} \times \frac{\varepsilon}{m\sigma^2} \end{aligned} \tag{6.19}$$

We used a patch width parameter σ^{PW} of 0.5, as this is the lowest value where we see re-entrance in our dynamical simulations.

Energy factor Finally, the Boltzmann factor in equation 6.1 reflects the additional energy obtained as a result of clustering (*i.e.* the energy of bonding between the patches of the monomers when arranged in an icosahedron). Ignoring next-neighbour interactions, the ground state energy of the icosahedron is $E_{\mathrm{clus}} = -30\varepsilon$, since each of the 12 monomers in an icosahedron is bonded to five other monomers with bond energy ε.

Overall, the icosahedral cluster molecular partition function is thus given by:

$$q_{\mathrm{clus}} = \left[\frac{k^{69}m^{36}T^{69}\sigma^{69}(\sigma^{\mathrm{PW}})^{36}\exp\left[30\varepsilon/kT\right]}{h^{72}\varepsilon^{33}}\right] \times \left[V - BN_{\mathrm{clus}} - bN_{\mathrm{mon}}\right] \times 9.4 \times 10^{20} \quad (6.20)$$

The overall monomer partition function is:

$$q_{\mathrm{mon}} = \frac{128\sqrt{\frac{2}{5}}\pi^5 k^3 m^3 \sigma^3 T^3}{25h^6} \times \left[V - BN_{\mathrm{clus}} - bN_{\mathrm{mon}}\right] \quad (6.21)$$

6.6.3 Clustering transition

When looking at the cluster-monomer transition, we make the assumption that only monomers and icosahedra occur in the system, and that the only reaction was $12\mathrm{A} \rightleftharpoons \mathrm{A}_{12}$ (*i.e.* we assume that there are no partially formed clusters). Previous work within the group [111] has shown that this is a fair assumption, given that the free energy of an icosahedron is considerably lower than any other non-monomeric cluster. This also ignores the free energy gains possible when several clusters are in contact (this is a secondary effect for all patch widths).

At equilibrium, the chemical potentials of the individual species, appropriately weighted by their stoichiometric coefficients, will be equal; in this case, $12\mu_{\mathrm{mon}} = \mu_{\mathrm{clus}}$, where $\mu_i = -kT\ln(q_i/N_i)$. We define the centre of this equilibrium as the point at which the probabilities of a particle being in a cluster or a monomer are equal, which imposes the condition that the size-weighted number densities are equal ($\rho_{\mathrm{mon}}^{\mathrm{eq}} = 12\rho_{\mathrm{clus}}^{\mathrm{eq}}$). The temperature at which both conditions are fulfilled is denoted as T_{clus}.

These equations are insoluble analytically, and we used Mathematica to minimise numerically $12\mu_{\mathrm{mon}} - \mu_{\mathrm{clus}}$ (subject to the constraint $\rho_{\mathrm{mon}} = 12\rho_{\mathrm{clus}}$) to obtain T_{clus}, as a function of the total number density ρ. The results are plotted in figure 6.13. Note that the transition between the two states is continuous (see figure 6.14), thus T_{clus} is not a phase boundary for a first order phase transition, but rather the centre of the chemical equilibrium where the clusters and monomers are equally probably states.

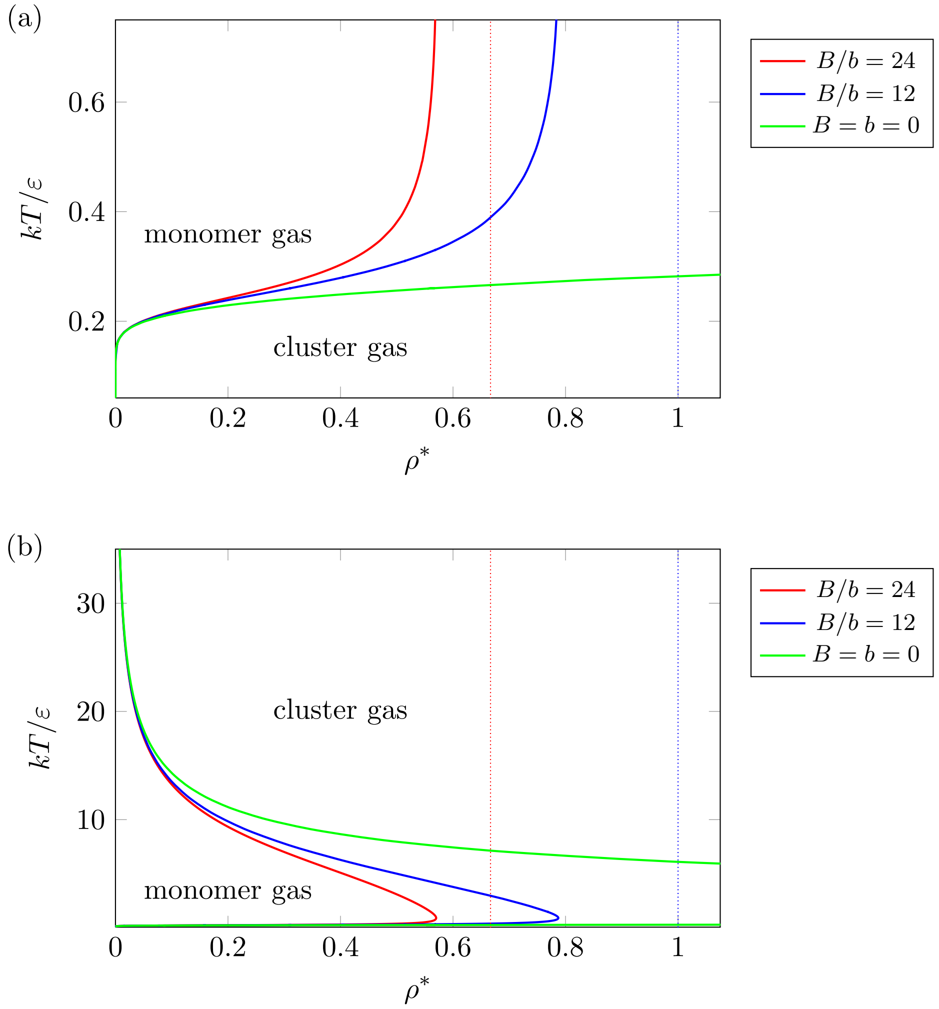

Figure 6.13: $T_{\text{clus}}(\rho^*)$, at (a) low and (b) high temperatures, for the transition from monomers to icosahedral clusters for different ratios of the free volume coefficients, B/b. For the non-ideal systems, $b\sigma^{-3} = 1$, and dotted lines show the maximum possible density for which particles can have an equal probability of being monomeric or in a cluster. $\sigma^{\text{PW}} = 0.5$.

In the ideal limit (where particles have no volume, and thus $b = 0$ and $B = 0$), the cluster gas is always more stable than the monomer at low temperature (figure 6.13). This stability occurs as, at low thermal energy, the internal energy of the clusters resulting from their bonding is more important than the loss of entropy due to fewer translational degrees of freedom. As the overall density increases, the clustering temperature (T_{clus}) increases sharply initially before becoming relatively flat. When we introduce excluded volumes into the monomer-cluster equilibrium, we change the behaviour of the transition line. At low densities, the excluded volume interactions have little effect and the $T_{\text{clus}}(\rho^*)$ lines initially follow that for the ideal case. However, as the density increases, the lines are displaced upwards with respect to the ideal case because the volume term $V - BN_{\text{clus}} - bN_{\text{mon}}$ in the translational partition functions of both monomers and clusters decreases, and this destabilises the monomeric state more because it has a greater number of translational degrees of freedom.

At sufficiently high temperatures, the cluster gas is also more stable in our model. This feature is an unphysical artefact resulting from our treatment of the bonds as harmonic oscillators with effectively infinite depth wells (when, in reality, the bonds would be dissociated at these temperatures). This problem is not an issue for the application of the model, as we are only interested in significantly lower temperatures where physical cluster formation and liquid-vapour phase separation occur.

On introducing excluded volumes, we also create a maximum density beyond which our definition of equilibrium is no longer possible, thus $\rho_{\text{mon}} + B\rho_{\text{clus}} \leq 1$. We have to keep our equilibrium condition of $\rho^{\text{eq}}_{\text{mon}} = 12\rho^{\text{eq}}_{\text{clus}}$, which allows us to write $12\rho^{\text{eq}}_{\text{clus}} + B\rho^{\text{eq}}_{\text{clus}} < 1$, and therefore the maximum cluster density obtainable which maintains equilibrium ($\rho^{\text{eq}}_{\text{clus, max}}$), must follow $(12 + B)\rho^{\text{eq}}_{\text{clus, max}} = 1$. For the $B/b = 24$ case, we therefore have $\rho^{\text{eq}}_{\text{clus}} \leq 1/36$. This gives us a maximum overall density ($\rho = 12\rho_{\text{clus}} + \rho_{\text{mon}}$) of two thirds ($\rho^{\text{eq}}_{\text{max}} = 12\rho^{\text{eq}}_{\text{clus, max}} + \rho^{\text{eq}}_{\text{mon, max}} = 24\rho^{\text{eq}}_{\text{clus, max}} = 24/36 = 2/3$). For the $B/b = 12$ system, we get $\rho^{\text{eq}}_{\text{max}} = 1$. These lines are plotted as dotted lines in figure 6.13, and thus denote the point at which the system would again prefer to be a monomeric rather than a cluster gas. As these points involve regions of practically no free volume, they become numerically unstable, and thus difficult to solve numerically using Mathematica (though we can extrapolate the relative proportions by assuming the free volume is equal to zero). We use our calculations and an extrapolation to show the changes in the relative proportions of clusters as a function of overall monomer density in figure 6.14. When we consider the self-assembling van der Waals fluid in section 6.6.5, we do not need to worry about the issues that occur when $\rho > 12\rho_{\text{clus, max}}$, because phase separation intervenes before this density is reached.

6.6.4 Standard van der Waals fluid

The canonical partition function for a van der Waals (vdW) monatomic fluid is given by [151] (where a and b are positive constants):

$$Q = \frac{1}{N!}\left(\frac{2\pi mkT}{h^2}\right)^{(3N/2)} (V - Nb)^N \exp\left[\frac{a\rho N}{kT}\right] \tag{6.22}$$

This expression is analogous to that for an ideal gas, but with an excluded volume term (Nb) and an attractive exponential term. From the canonical partition function, we can

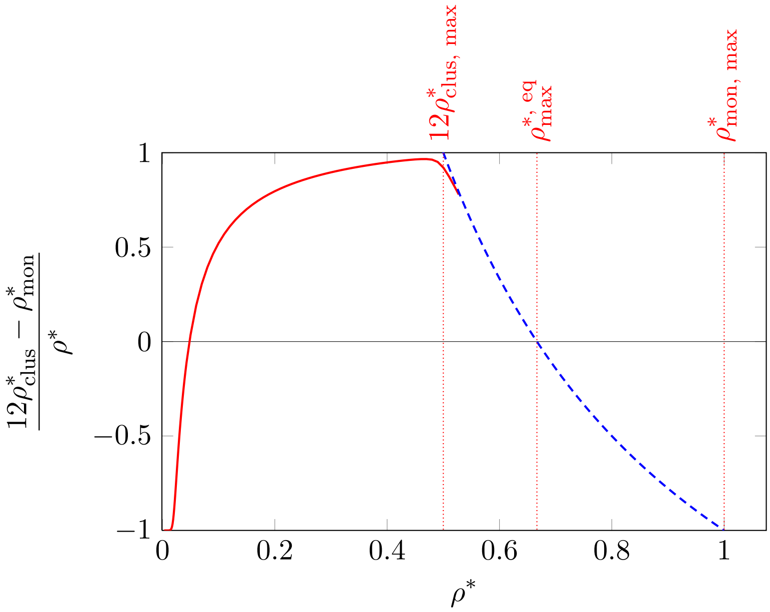

Figure 6.14: The relative proportion of monomers and clusters in the equilibrium mixture is shown as a function of the density at $kT/\varepsilon = 0.2$ for $B/b = 24$. The function $(12\rho^*_{\text{clus}} - \rho^*_{\text{mon}})/\rho^*$ takes the value -1 when only monomers are present, $+1$ when only clusters are in the system, and 0 at the centre of the clustering transition (*i.e.* on the $T_{\text{clus}}(\rho^*)$ line). Vertical dotted lines indicate ρ^* at the maximum cluster density $12\rho^*_{\text{clus, max}}$, the maximum equilibrium density $\rho^{*,\text{ eq}}_{\text{max}}$ and the maximum monomer density $\rho^*_{\text{mon, max}}$, as labelled. The solid line is calculated from the partition functions, whilst the dashed line is the extrapolation for the regime where the clusters fill as much volume as they can.

find the Helmholtz energy (A), pressure (p) and chemical potential (μ) using:

$$A = -kT \ln Q \tag{6.23}$$

$$p = -(\partial A/\partial V)_{T,N} \tag{6.24}$$

$$\mu = (\partial A/\partial N)_{V,T} \tag{6.25}$$

Our molecular partition function can be derived from the canonical partition function (needed for the system in section 6.6.5, where we include an additional rotational term) by using the identity for indistinguishable identical particles [158]:

$$Q = \frac{q^N}{N!} \tag{6.26}$$

Thus:

$$q_{\text{vdW}} = \left(\frac{2\pi mkT}{h^2}\right)^{(3/2)} (V - Nb) \exp\left[\frac{a\rho}{kT}\right] \tag{6.27}$$

We are interested in the liquid and gas phases of our system, and their coexistence points. We wish to find the binodals (points at which the thermodynamically preferred phase changes), as well as the spinodals (the points at which one phase ceases to be mechanically stable (metastable), and thus the system would be unstable with respect to density fluctuations).

When we are trying to find the coexistence points between the liquid and gas phases (*i.e.* the phase transition line), we could use the common method of Maxwell construction. This scheme involves plotting the pressure against volume per particle for a temperature below the critical point, and finding a line, parallel to the x-axis, which encloses equal areas of the two van der Waals loops on either side of the line (figure 6.15(a)). Unfortunately, van der Waals loops do not always exist in more complicated systems [229, 230], and the calculations of areas is computationally expensive. So, instead, we calculate the chemical potential (μ) and the pressure (p), and parametrically plot these two quantities against each other to determine the binodal and spinodal points. At coexistence, the two phases have the same chemical potential and, whenever a given chemical potential has two (or more) different pressures associated with it, the stable phase is the one with the *largest pressure* (maximum entropy and therefore favoured by the second law of thermodynamics) [229, 238]. This method is demonstrated in figure 6.15(b).

The resulting phase diagram for the van der Waals fluid is shown in figure 6.16. Unsurprisingly the results from the Maxwell construction and the equality of the chemical potentials are in full agreement.

6.6.5 Self-assembling van der Waals fluid

We now wish to calculate the phase diagram of a ‘self-assembling’ van der Waals fluid, where monomers can either assemble into clusters (which we model as having no attractive interactions) or condense to form a liquid. We use the partition functions of the cluster (equation 6.20) and a modified version of the partition function for the van der Waals fluid. In our self-assembling system, we add a rotational degree of freedom to the partition function (as we did for monomers in the previous section), adjust the value of a to give a similar behaviour to that which we find in our model from simulations (which is $a = \varepsilon$, though we vary this later), and change the volume term to exclude volume taken up by the cluster (using the free volume approximation [229]). We chose a with the aim of ensuring that the critical van der Waals temperature $kT_\mathrm{c} = 8a/27b$ is near to, but greater than, the clustering temperature (at the critical density). Our choice gives a reduced critical temperature ($kT_\mathrm{crit}/\varepsilon$) of approximately 0.30 for the pure van der Waals fluid and our clustering temperature at the critical density ($kT_\mathrm{clus}(\rho_\mathrm{crit})/\varepsilon$) is approximately 0.23 (in the absence of liquid formation (*i.e.* $a = 0$)). This gives us the following as our molecular

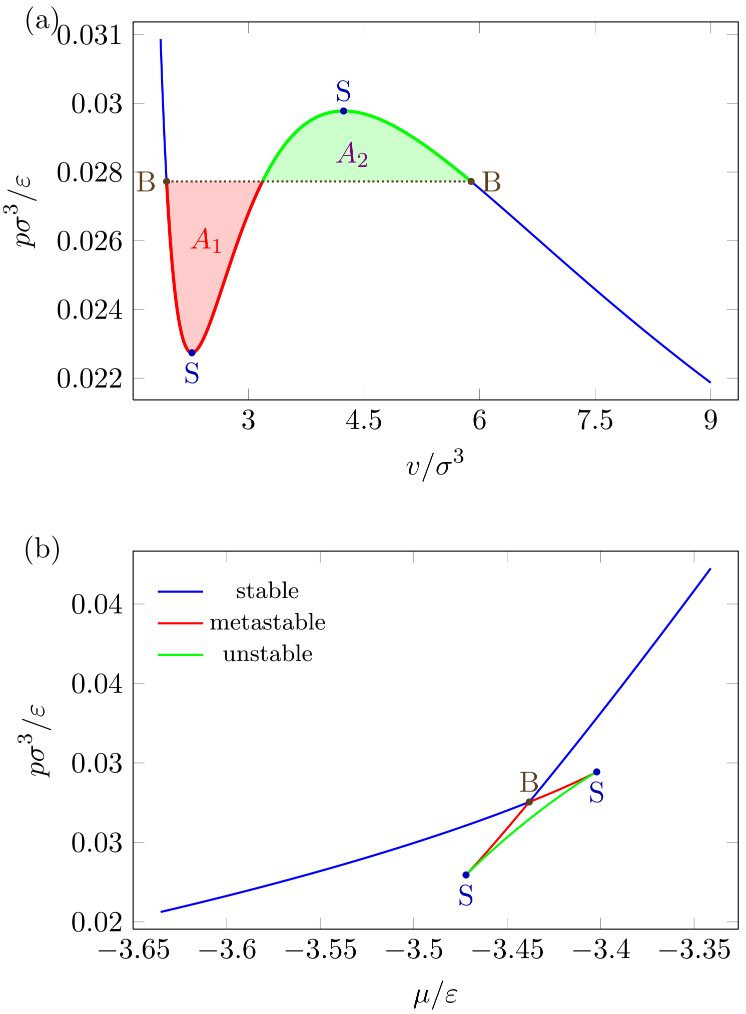

Figure 6.15: Methods to locate coexistence. (a) Maxwell equal areas construction, and (b) a plot of pressure as a function of chemical potential. B denotes the binodals and S the spinodals.

partition function for the van der Waals fluid:

$$q_{\text{vdW}} = \left(\frac{2\pi mkT}{h^2}\right)^{(3/2)} (V - bN_{\text{mon}} - BN_{\text{clus}}) \exp\left[\frac{\varepsilon\rho}{kT}\right] \times \frac{64\pi^{7/2}\sigma^3(mkT)^{3/2}}{25\sqrt{5}h^3} \tag{6.28}$$

Coexistence curve

To calculate coexistence we use the chemical potential equivalence approach (demonstrated in figure 6.15(b)). As before (section 6.6.2), we assume that the only relevant reaction in the system is $12\text{A} \rightleftharpoons \text{A}_{12}$. We then use our overall partition functions to find the coexistence curves for the combined system (with van der Waals monomers and clusters in equilibrium). For equilibrium, we again set the chemical potentials, weighted by the stoichiometric coefficients, to be equal ($\mu_{\text{clus}} = 12\mu_{\text{mon}}$). The system of equations must

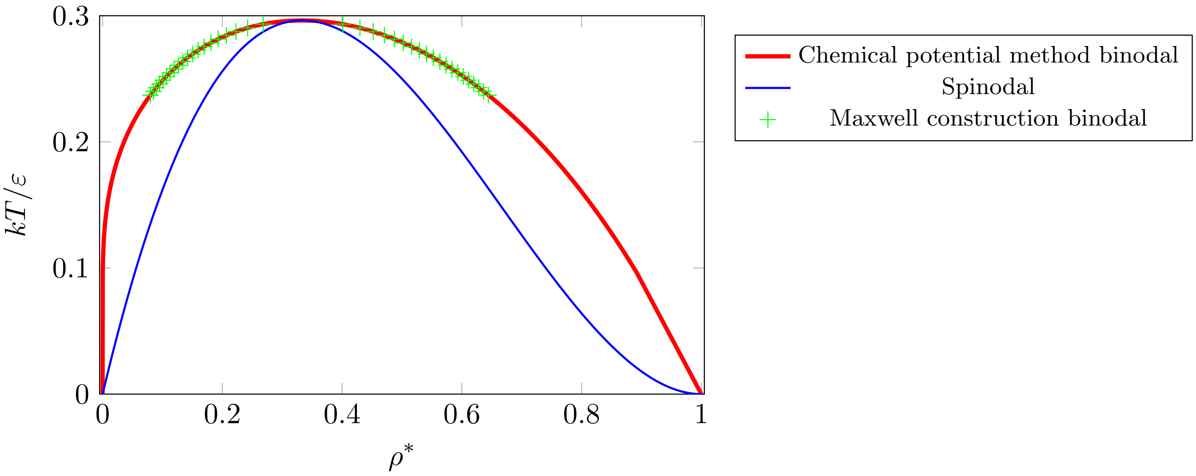

Figure 6.16: The binodals of a van der Waals fluid by chemical potential equivalence method, the binodals from the Maxwell construction method, and the spinodals from setting $\frac{\partial p}{\partial v} = 0$, using $b = 1\,\sigma^3$ and $a = 1\,\varepsilon$.

be solved numerically (as they are analytically insoluble owing to the interdependence between the cluster and monomer density). When finding the pressure, p, as a function of the chemical potential, we chose a chemical potential μ and used Mathematica to minimise the absolute differences of the variables we wanted to be equal (*i.e.* we minimise both $|\mu_{\text{clus}} - 12\mu_{\text{mon}}|$ and $|\mu_{\text{clus}}/12 + \mu_{\text{mon}} - 2\mu|$) by varying both ρ_{mon} and ρ_{clus}). We used the constraints that the free volume must not be negative ($V - BN_{\text{clus}} - bN_{\text{clus}} \geq 0$), and that each density must be positive ($\rho_{\text{mon}} \geq 0$ and $\rho_{\text{clus}} \geq 0$). These constraints are equivalent to requiring that we have physical densities, that $\mu_{\text{clus}} \approx 12\mu_{\text{mon}}$ (the system is at equilibrium) and that $\mu_{\text{clus}} \approx 12\mu$ (the system has the desired chemical potential). We only recorded values which gave an error term (which is the amount by which the two equations above are not equal) of less than 10^{-6}, though generally the error was significantly less. With values for the densities calculated, the pressure was calculated using the equations 6.23 and 6.24, and noting that the total Helmholtz energy is the sum of the free energies of each component. By performing these calculations at different temperatures, we were able to construct a phase diagram (figure 6.18). This part of the project required a large number of workarounds to deal with numerical instabilities and multiple solutions of the equations, and took approximately 6 months to complete.

6.6.6 Stability criteria

As mentioned previously, metastable states exist where a phase may persist and be kinetically stable within a region where they are thermodynamically unstable. We can delimit

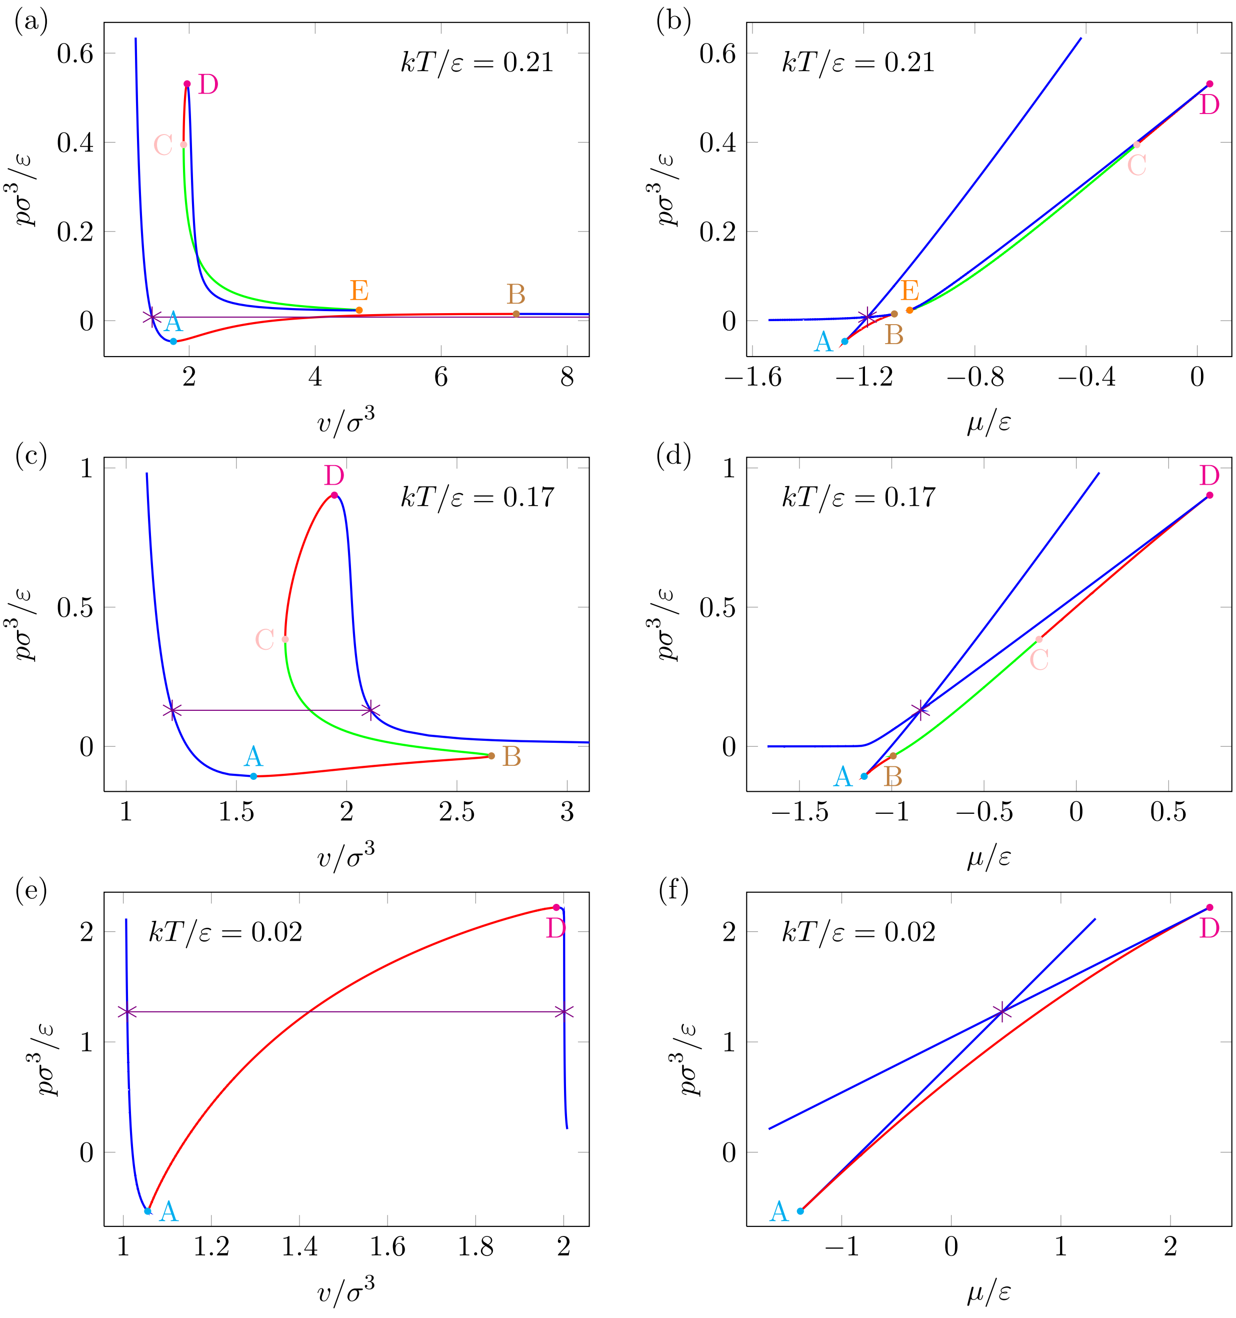

Figure 6.17: The p-v ((a), (c) and (e)) (v is the volume per particle, V/N) and p-μ ((b), (d) and (f)) curves for the self-assembling van der Waals fluid at reduced temperatures of (a) and (b) $kT/\varepsilon = 0.21$, (c) and (d) $kT/\varepsilon = 0.17$ and (e) and (f) $kT/\varepsilon = 0.02$. Regions of mechanical stability are coloured blue, those of mechanical instability red, and the mechanically stable, but compositionally unstable, region of back-bending is coloured green. Binodal points are marked by violet asterisks connected by a tie line. Spinodal points, where the derivative of the pressure with respect to the volume is either zero or infinity, are also shown with the labelling and colour-coding matching that used in figure 6.18. $a = \varepsilon$, $b = \sigma^3$, $\sigma^{\mathrm{PW}} = 0.5$, $B/b = 24$.

the regions of mechanical stability on a plot of pressure and volume. Where the gradient of the p-v curve is positive, we have mechanical instability. Figure 6.17(c) shows that in our model at some temperatures there are two regions of mechanical instability (between A and B, and C and D), with an unusual central region of mechanical stability (between B and C) and two regions that are mechanically stable and thermodynamically unstable between the higher-density binodal and point A, and between point D and the lower-density binodal.

Unlike the standard van der Waals model, we also have to check for compositional stability, *i.e.* whether fluctuations away from the equilibrium composition of monomers and clusters are stable. To derive the limits of compositional stability, we start from the fundamental equation for the Helmholtz free energy, which may be written as:

$$\mathrm{d}A = -p\,\mathrm{d}V - S\,\mathrm{d}T + \mu_{\mathrm{mon}}\,\mathrm{d}N_{\mathrm{mon}} + \mu_{\mathrm{clus}}\,\mathrm{d}N_{\mathrm{clus}} \tag{6.29}$$

The number of monomers and cluster are not independent, since $N = 12N_{\mathrm{clus}} + N_{\mathrm{mon}}$ therefore $\mathrm{d}N = 12\,\mathrm{d}N_{\mathrm{clus}} + \mathrm{d}N_{\mathrm{mon}}$. This allows us to rewrite equation 6.29 as:

$$\mathrm{d}A = -p\,\mathrm{d}V - S\,\mathrm{d}T + \mu_{\mathrm{mon}}\,\mathrm{d}N + (\mu_{\mathrm{clus}} - 12\mu_{\mathrm{mon}})\,\mathrm{d}N_{\mathrm{clus}} \tag{6.30}$$

So overall:

$$\left(\frac{\partial A}{\partial N_{\mathrm{clus}}}\right)_{N,V,T} = \mu_{\mathrm{clus}} - 12\mu_{\mathrm{mon}}, \tag{6.31}$$

which must be zero at equilibrium, and the second derivative is our condition for compositional stability:

$$\left(\frac{\partial^2 A}{\partial N_{\mathrm{clus}}^2}\right) = \left(\frac{\partial(\mu_{\mathrm{clus}} - 12\mu_{\mathrm{mon}})}{\partial N_{\mathrm{clus}}}\right)_{N,V,T} > 0. \tag{6.32}$$

We have evaluated this stability criterion at all of the relevant equilibrium points of the combined system. The system is compositionally unstable along the line BC, and thus the mechanical stability in this region is not relevant.

6.6.7 Phase diagrams

The phase diagram we obtained is shown in figure 6.18. Re-entrant phase behaviour is clearly demonstrated for densities $0.05\,\sigma^{-3} \lessapprox \rho \lessapprox 0.5\,\sigma^{-3}$. As the temperature decreases at constant density within this range, the system first undergoes a transition from a monomeric fluid to a two-phase liquid-vapour coexistence region, and then as the temperature is further decreased, to a gas of icosahedral clusters.

As in section 6.6.3, we also calculate the clustering transition temperature, $T_{\mathrm{clus}}(\rho)$. With liquid formation being favourable at high densities, however, the clustering transition

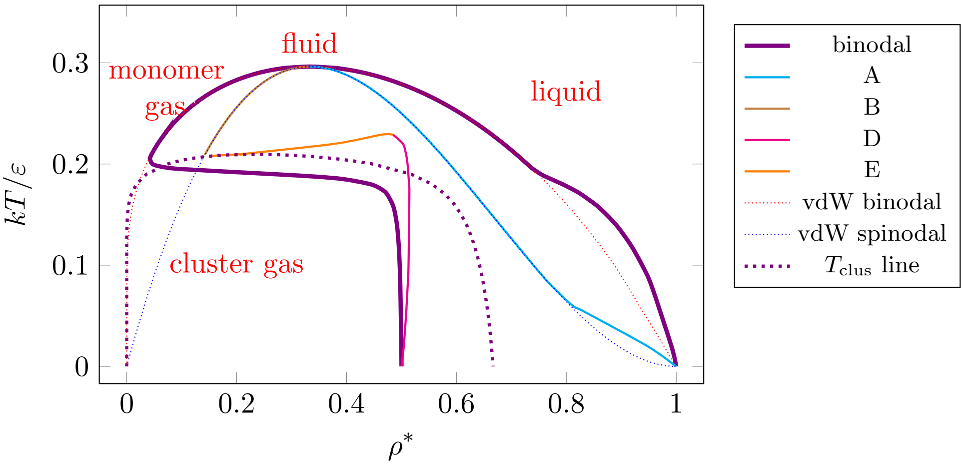

Figure 6.18: The T-ρ^* phase diagram for the self-assembling van der Waals fluid. Lettered lines correspond to points A, B D and E in the p-v plots in figure 6.17, and are limits of stability.

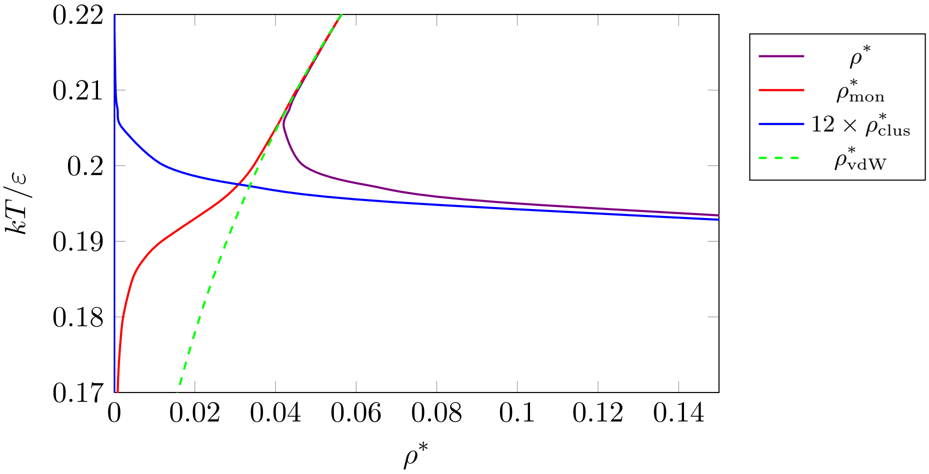

Figure 6.19: The region of T-ρ phase diagram for the self-assembling van der Waals fluid where we see a change in the regime. The binodals are plotted for the overall density, as well as for clusters and monomers. For comparison the curve for the van der Waals binodal is also shown as a dotted line.

curve does not monotonically increase in temperature as the density increases, but instead goes through a maximum and then gradually decreases, such that it reaches the maximum permissible density $\rho^{\text{eq}}_{\text{max}}$ at $kT/\varepsilon = 0$, as shown in figure 6.18.

The densities of the monomer and cluster species at equilibrium for the binodal in the re-entrance region are shown in figure 6.19, along with the binodal for the pure van der Waals system for comparison. The figure clearly shows that the divergence from van der Waals-like behaviour occurs when the clusters begin to be formed and that the system quickly changes from being mainly monomer to being dominated by clusters.

It should be noted that, at high temperatures, we see the self-assembling van der Waals fluid exhibiting liquid-vapour coexistence behaviour exactly as we do for the pure van der Waals fluid, and the binodals of the self-assembling fluid follow those of the van der Waals fluid (the cluster density is very low as we are far above the clustering temperature). This liquid-vapour phase separation is driven by the attractions between monomers. In this region, the phase boundary exhibits a positive gradient in the p-T phase diagram (figure 6.20). The slope is given by the Clapeyron equation [222]:

$$\frac{\mathrm{d}p}{\mathrm{d}T} = \frac{\Delta_{\mathrm{trs}} H}{T_{\mathrm{trs}} \Delta_{\mathrm{trs}} V} = \frac{\Delta_{\mathrm{trs}} S}{\Delta_{\mathrm{trs}} V} \tag{6.33}$$

Since the entropy, enthalpy and volume changes are all negative for the vapour-liquid transition, the resulting slope is naturally positive.

On the other hand, at sufficiently low temperature, we have an appreciable yield of clusters. Since $T_{\mathrm{clus}}(\rho)$ lies at lower density than the van der Waals binodal, at a given temperature, as we increase the density of the system, the vapour forms clusters before it reaches the point at which it would otherwise have become phase separated. As we have no inter-cluster attractions in our model, the vapour is now stable and does not phase separate.

If we further increase the density of this system, we approach the maximum density of clusters, and thus packing constraints become important. If the density is to move beyond $12\rho_{\mathrm{clus,\ max}}$ and remain homogeneous, the system must form a monomer-cluster mixture with the proportion of monomers increasing with ρ. However, unlike in section 6.6.3, there is a strong tendency for the system to de-mix (which is preferred for monomers as there is a lack of attractions between clusters and monomers). This de-mixing is a density-driven transition, and the position of the lower-density binodal varies little with temperature in this regime. The vapour phase is now lower in internal energy than the liquid due to the strong bonding within the clusters, so the enthalpy change associated with vapour-liquid transition is now positive. The volume change is, of course, still negative. Using the Clapeyron equation (equation 6.33), we can therefore explain why the cluster-liquid transition has a negative slope in the p-T diagram (figure 6.20).

Note that both the solubility of monomers in the cluster fluid and of clusters in the monomeric liquid is very low. For example, at $kT = 0.08\,\varepsilon$, $\rho_{\mathrm{mon}} = 3.44 \times 10^{-7}\,\sigma^{-3}$ at the lower-density binodal ($\rho = 0.498\,\sigma^{-3}$), and $\rho_{\mathrm{clus}} = 7.28 \times 10^{-120}\,\sigma^{-3}$ at the higher-density binodal ($\rho = 0.955\,\sigma^{-3}$). Such small numbers lead to serious numeric difficulties unless using a package with the ability to deal with this range of magnitudes, and is one of the reasons for our use of Mathematica over Fortran.

The crossover between regimes of van der Waals-like behaviour and density-driven

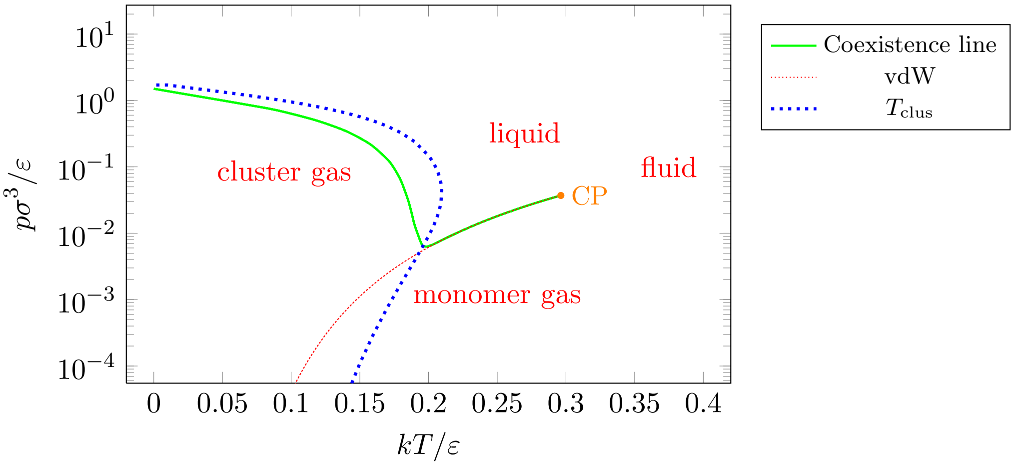

Figure 6.20: The p-T phase diagram for the self-assembling van der Waals fluid. The critical point is labelled 'CP'. The centre of the equilibrium between monomers and clusters (T_{clus}) is also plotted (thick dotted line).

de-mixing occurs at temperatures near to the crossing of $T_{\mathrm{clus}}(\rho)$ and the van der Waals binodal. This crossover leads to a rapid increase in the density of the low-density binodal with decreasing temperature until $12\rho_{\mathrm{clus,\ max}}$ is approached. The width of this crossover regime reflects the width of the monomer-cluster equilibrium. The change in slope occurs slightly above T_{clus}, as the vapour first begins to cluster. Initially, when monomers are still in the majority at the binodal, the transition is still driven by the attractions between the monomers, and the binodal occurs when the monomer density reaches the density of the binodal for the pure van der Waals system. However, as clusters begin to predominate in the coexisting mixture, the phase transition is driven more by the cluster-monomer de-mixing, and occurs when the monomer density reaches a critical (but low) solubility. This is demonstrated in figure 6.19.

These changes in the phase diagram are also reflected in the p-v and p-μ diagrams (figure 6.17). For temperatures below the critical point, but above the clustering temperature ($0.23\,\varepsilon/k < T < T_{\mathrm{crit}}$), we observe van der Waals-like p-v and p-μ curves. However, below this temperature, even though not yet in the region of re-entrance, we can observe additional metastable closed cluster-rich curves in both p-v and p-μ plots (figure 6.17(a) and (b)). As the temperature decreases, the cluster-rich loop merges with the van der Waals curve causing back-bending in the p-v plane, and thus exhibiting two additional spinodal points compared to the situation in the pure van der Waals case.

Unlike in the pure van der Waals system, spinodals are not as straightforward to determine as binodals. We numerically found derivatives of pressure with respect to volume using p-v phase diagrams as a starting point. The compositional stability criterion (section 6.6.6) was also calculated. Unusual spinodal lines that occur outside of the binodals can be seen in figure 6.17(c) (point B). Using the compositional stability criterion from equation (6.32), we can see that these unusual regions, although mechanically stable, are kinetically compositionally unstable, and therefore are not true limits of kinetic stability. The physically relevant spinodals are therefore those labelled A and B in the van der Waals-like region of the phase diagram, and then D instead of B in lower regions (labels from figure 6.20); these are plotted in figure 6.18.

Changing the relative size of clusters and monomers

For our first test system, where $B/b = 24$, we have re-entrance, but, at all temperatures below T_{crit}, there is always a density range for which phase separation occurs. We now consider how the topology of the phase diagram changes when we vary the parameters of the model. In particular we want to investigate whether the phase separation at low temperature can be removed, and, if so, what is the form of the phase diagram: would the system exhibit a closed loop in the p-T plane, similar to the nicotine-water mixture? We looked at lower ratios of B/b; as this determines $\rho_{\text{clus, max}}$, and hence where the onset of density-driven phase separation occurs. Initially, as B/b decreases, the clusters are able to pack more efficiently and the lower-density binodal moves to higher density. This decrease of B/b does not lead to any major change in the topology of the phase diagram. However, when the binodal associated with the cluster fluid meets and crosses the binodal associated with the monomeric liquid, we see a different topology at lower temperature entirely. Figure 6.22(a) illustrates the type of phase diagram that results and is for $B/b = 12.5$. We note that this scenario does not lead to a lower critical point; rather, at the point marked A in the phase diagram, there are two coexisting fluids with the same density, but which differ significantly in other characteristics. We have a monomeric liquid with high entropy, and a dense cluster fluid with low internal energy. These two fluids do not evolve continuously into each other as point A is approached (as would need to be the case for there to be a critical point), but instead retain their separate identities. Similarly, point B corresponds to a point with two coexisting fluids with the same density, but it also has phase separation between a cluster gas and a monomeric liquid.

For $B/b = 12.5$, at both point A and point B, the change in volume (ΔV) is zero. Thus, by the Clapeyron equation (equation 6.33), $\mathrm{d}p/\mathrm{d}T \to \infty$. We therefore have two points at which the tangent to the coexistence curve in the p-T phase diagram is vertical (figure 6.21). The lower pressure one corresponds to point A, and the higher one to point B. After point A, the coexisting cluster fluid is more dense than the monomer liquid, and so $\mathrm{d}p/\mathrm{d}T$ is positive. However, as $B/b > 12$, the clusters still pack less efficiently, and so, at low temperatures, and densities beyond $12\rho_{\text{clus, max}}$ there must again be a region of

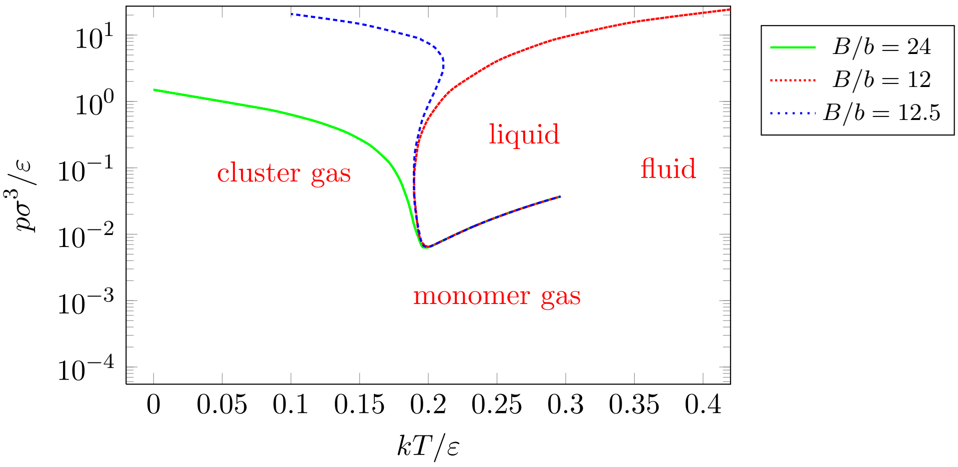

Figure 6.21: The p-T phase diagram for the self-assembling van der Waals fluid at different values of B/b.

de-mixing, because of the lack of attraction between clusters and monomers. The slope of the p-T phase boundary corresponding to this de-mixing region is negative. Interestingly, for this phase diagram topology, there is a temperature range ($0.19 \lessapprox kT/\varepsilon \lessapprox 0.21$ for $B/b = 12.5$) where there are three fluid-fluid phase transitions as the system pressure changes (figure 6.21).

Finally, at $B/b = 12$, there is a further change in the phase diagram topology (figure 6.22(b)), because at this point monomers can no longer pack more efficiently than clusters, so we will not see a density driven phase separation. As B/b is decreased from 12.5 to 12, point B moves to higher temperatures, and the density range for the low temperature cluster-fluid/monomeric liquid de-mixing becomes smaller before eventually disappearing at $B/b = 12$.

Changing the liquid-vapour critical point

We then considered how the phase behaviour of the system changes as the position of the critical point is varied. For the pure van der Waals fluid $T_{\text{crit}} = 8a/27b$, and so, for a given B, T_{crit} is determined by the value of a. In our model of the monomer-icosahedron equilibrium, T_{clus} is determined by σ^{PW}. In contrast to our theoretical model, simulated patchy particle systems do not have independent a and σ^{PW} variables; rather, the position of the critical point changes with the patch width (and the number of patches) [239, 240]. For example, it can be seen from figure 6.1 that the temperature range associated with liquid-vapour coexistence decreases as σ^{PW} decreases. Here, we choose to vary one of the parameters (a) whilst keeping the other one (σ^{PW}) fixed.

The phase diagrams for a variety of values of a are shown in figure 6.23. For $a = \varepsilon\sigma^3$ and $a = 0.8\,\varepsilon\sigma^3$, the van der Waals critical temperatures ($kT_{\text{crit}} = 0.296\varepsilon$ and

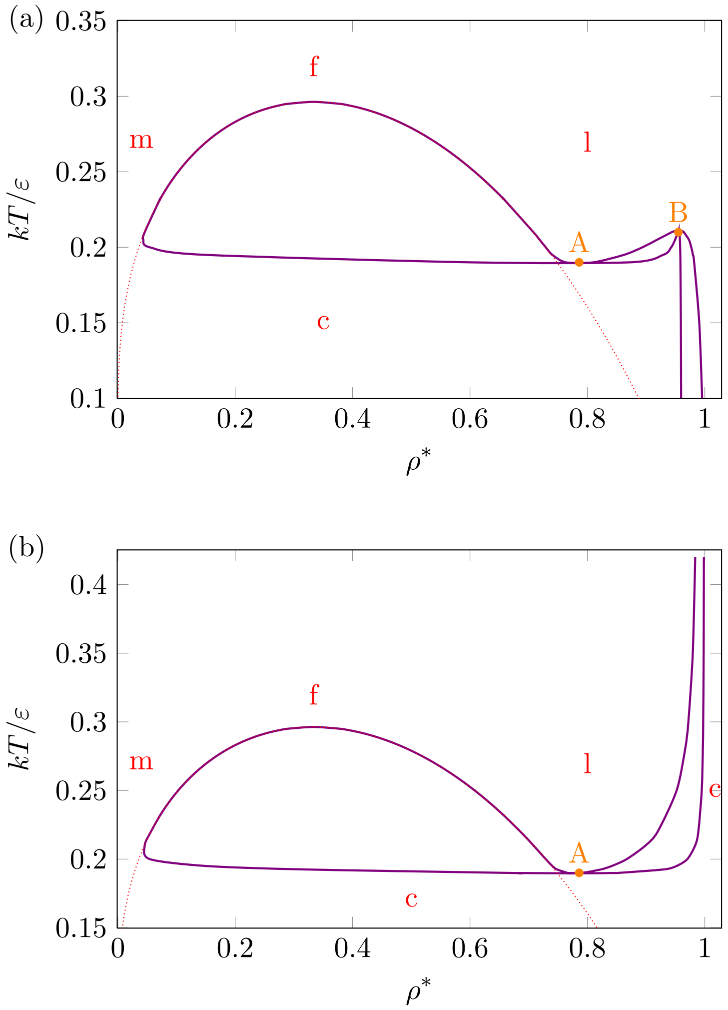

Figure 6.22: Binodal lines for (a) $B/b = 12.5$ and (b) $B/b = 12$ are shown in violet, compared with the dotted red line representing the pure van der Waals binodal. The labels 'c', 'm', 'l' and 'f' refer to cluster gas, monomer gas, liquid and super-critical fluid states, respectively.

$kT_{\text{crit}} = 0.237\varepsilon$, respectively) are above the clustering temperature at the critical density ($kT_{\text{clus}}(\rho_{\text{crit}}) \approx 0.23\varepsilon$), and van der Waals behaviour is observed at high temperatures. When the critical temperature descends below the clustering temperature, we nevertheless continue to observe a bending of the binodal curve to lower densities in the region where liquid-vapour coexistence previously took place. However, this feature is not some remnant of the liquid-vapour transition, but is instead associated with de-mixing in the vicinity of $T_{\text{clus}}(\rho)$. At T_{clus}, there is an equal probability of particles being monomeric or as part of clusters; however, if ρ and a are sufficiently large, the system will not form a homogeneous mixture, but will instead de-mix. The narrowness in temperature of this feature is associated with the sharpness with which the chemical equilibrium shifts between being monomer dominated to being cluster dominated as the temperature decreases. As a decreases, there

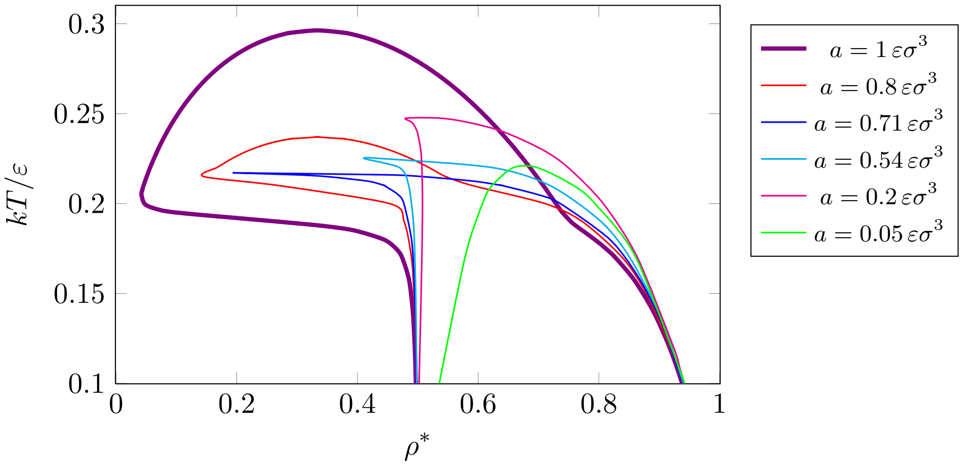

Figure 6.23: A T-ρ^* phase diagram for the self-assembling van der Waals fluid with a selection of values for the van der Waals attractive parameter.

is a progressively smaller de-mixing driving force and so the size of this feature decreases. Note that these changes lead to a non-monotonic dependence of T_{crit} on a, because T_{crit} initially increases as ρ_{crit} increases, mirroring the dependence of T_{clus} on ρ (figure 6.13). As a becomes smaller still, the clusters and monomers begin to be able to mix at high temperatures because less entropy of mixing is required to overcome the enthalpic mixing incompatibility. Hence, T_{crit} again begins to decrease. However, even at a small value of a, *e.g.* $a = 0.05\,\varepsilon\sigma^3$ in figure 6.23, there is still a substantial region of de-mixing. As a tends to zero, the de-mixing regime disappears, recovering the monomer-cluster equilibrium studied in section 6.6.2.

Other shapes of clusters

Obviously, icosahedra are not the only clusters possible, and our approach can be generalised to many other forms of cluster. We looked at tetrahedral cluster forming particles to demonstrate the flexibility of our theoretical model, and to illustrate how the phase diagram changes.

The equations should be modified as follows.

1. The number of particles in the cluster is 4.
2. The moment of inertia of a tetrahedron about the principal axes is $I = 4m\sigma^2$.
3. The mass of the cluster is $4m$.
4. The symmetry number of a tetrahedron is 12.
5. The symmetry number of a monomer is 3.

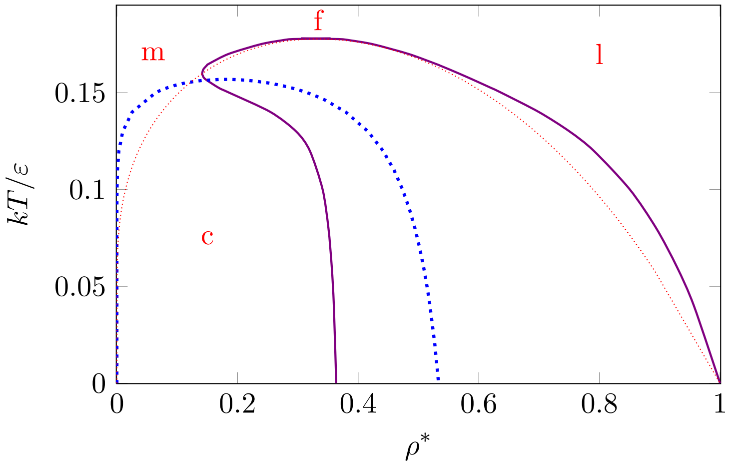

Figure 6.24: A T-ρ^* phase diagram for the self-assembling van der Waals fluid forming tetrahedral clusters. The solid line is the binodal curve, whilst the thin dotted line represents the pure van der Waals system binodal. The thick dashed line is the clustering temperature, $T_{\text{clus}}(\rho^*)$, for this system. The labels 'c', 'm', 'l' and 'f' refer to cluster gas, monomer gas, liquid and super-critical fluid states, respectively. $a/\varepsilon = 0.6$, $B/b = 11$.

6. There are 6 radial, and 12 angular vibrational modes.
7. The ground state energy of the tetrahedron is -6ε.
8. The circumradius of a tetrahedron with edge length σ is $r = \sqrt{3/8}\sigma$ [236], thus we use $B/b = 11$.

This B/b ratio gives the maximum particle density associated with clusters of $4\rho_{\text{clus, max}}\sigma^3 = 4/11 \approx 0.36$.

As part of our derivation, we assume that the tetrahedra pack approximately like spheres. This may sound far-fetched, but recent research in fact confirms that the maximum packing fractions of tetrahedra and spheres are in fact very similar [241]. In order to take into account that we only have three patches per particle instead of five, we also scale the value of the interaction parameter a to 3/5 of that used for the icosahedron-forming system.

The phase diagram for the tetrahedral system is depicted in figure 6.24. As was expected, the behaviour of the system is similar to the icosahedral system. However, the clustering temperature is significantly lower due to the lower energetic driving force to form clusters. Similarly, the crossover between the liquid-vapour and de-mixing regimes occurs over a wider temperature range due to the greater width of the monomer-cluster equilibrium, which in turn arises from the smaller size of the clusters.

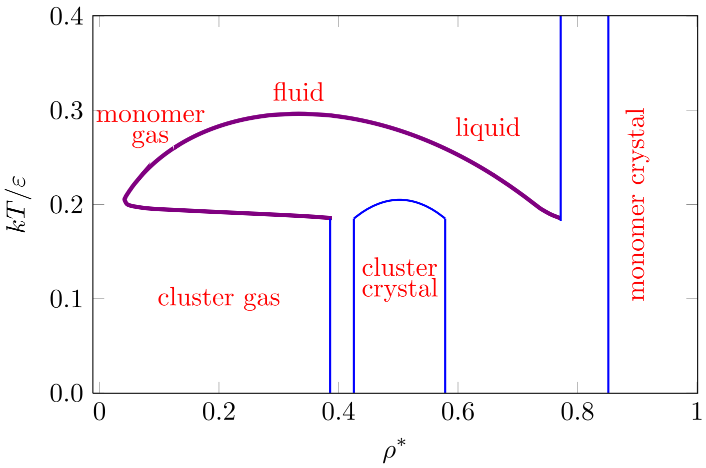

Figure 6.25: A schematic T-ρ^* phase diagram to illustrate the potential effects of crystallisation on the phase diagram in figure 6.18.

Inclusion of crystallisation

Our theoretical model completely neglects the formation of any crystalline phase. This choice was made because we were most fundamentally interested in the fluid behaviour of our system, and how re-entrance would be resolved at higher densities. However, in reality crystallisation (of both monomers and clusters) will have a major effect of the phase diagrams. There are precedents for cluster crystals; for example, icosahedral virus particles are often observed to form crystals when they are densely packed in cells [242]. We can construct a schematic phase diagram illustrating the potential effects of crystallisation by assuming hard-sphere crystallisation of both components. This approximation is expected to be quite reasonable for the clusters, as all the attractive patches are involved in the cluster bonds, though for monomers it is not so clear that this is a reasonable approximation, because of the potential effects of the patches on crystal stability and form.

As we used a value for the B and b parameters that was just the inverse of the maximum density of our fluid, when we now look at crystallisation, we need to find a hard sphere effective radius. If we assume that the maximum density of the cluster fluid is the random close packed density, we can scale the fluid-crystal coexistence packing fractions as follows:

$$\rho_{\mathrm{X}} = \frac{\phi_{\mathrm{X}}^{\mathrm{HS}}}{\phi_{\mathrm{RCP}}^{\mathrm{HS}}} \times \rho_{\mathrm{clus,\ max}} \tag{6.34}$$

where the subscript X refers to either the higher or lower density fluid-coexistence, the superscript HS refers to hard spheres, and RCP refers to randomly close packed.

For hard spheres, fluid-crystal coexistence occurs at packing fractions of $\phi = 0.494$ and $\phi = 0.545$ [243] and the random close packing density is $\phi_{\mathrm{RCP}}^{\mathrm{HS}} = 0.64$ [235]. For our system our maximum cluster fluid density, $\rho_{\mathrm{clus,\ max}}$, is $1/24\,\sigma^{-3}$. We can use these with our above equation to find the densities of fluid-crystal coexistence. The maximum hard-sphere crystalline packing fraction, $\phi = 0.74$ [244], allows us to estimate an upper density limit for the stability of the cluster crystal. We use the clustering temperature to estimate an upper temperature stability limit for the cluster crystal. The resulting schematic phase diagram is shown in figure 6.25. Note that this phase diagram has a second triple point associated with the disappearance of the cluster crystal phase, as well as the usual one associated with the disappearance of the liquid phase. We have not attempted to consider the effect of other possible phases on the phase diagram, such as the lamellar phase observed for Janus particles [133, 134] and the two-dimensional crystalline lamellae seen in simulations of our octahedron-forming patchy particles (figure 6.4).

6.7 Comparison of theoretical and simulation results

For most of the parameterisations of our model that we have considered, there is a range of densities that, on decreasing the temperature at constant density there is a series of transitions from a monomer fluid, to a liquid-gas phase-separated mixture, and then to a gas of clusters. Thus, re-entrant phase behaviour is a robust feature of our model, in agreement with our group's simulation results for patchy colloidal particles [111, 112], and the Janus particle simulations of Sciortino and co-workers [133, 134], both of which showed evidence of re-entrant phase behaviour.

There are, however, significant differences in the shape of the re-entrant region, which can be seen in figure 6.26. In the theoretical model, the lower-density binodal is predicted to increase in density at the region of re-entrance more quickly than the simulations predict. The simulation data also does not predict as sharp a change in the higher-density binodal as the theoretical model does, though we have not been able to simulate the higher-density binodal at sufficiently low temperatures to be able to see if there is a significant divergence between the two methods. These differences are partly due to limits in the theoretical description of the system. Firstly, we have assumed that clusters form either completely or not at all, *i.e.* only a single size of clusters can be formed. While the probability of observing partly-formed clusters along our monomer-cluster equilibrium line is small [111], at temperatures above the clustering transition, we do observe some small clusters. These small clusters mean that the width of the transition between monomers and clusters is wider than predicted by theory; this increased width may be partly responsible for the less steep changes in the lower-density binodal in our simulated systems compared to our theoretical model.

Our theoretical model is also limited in the use of the free volume approximation [229].

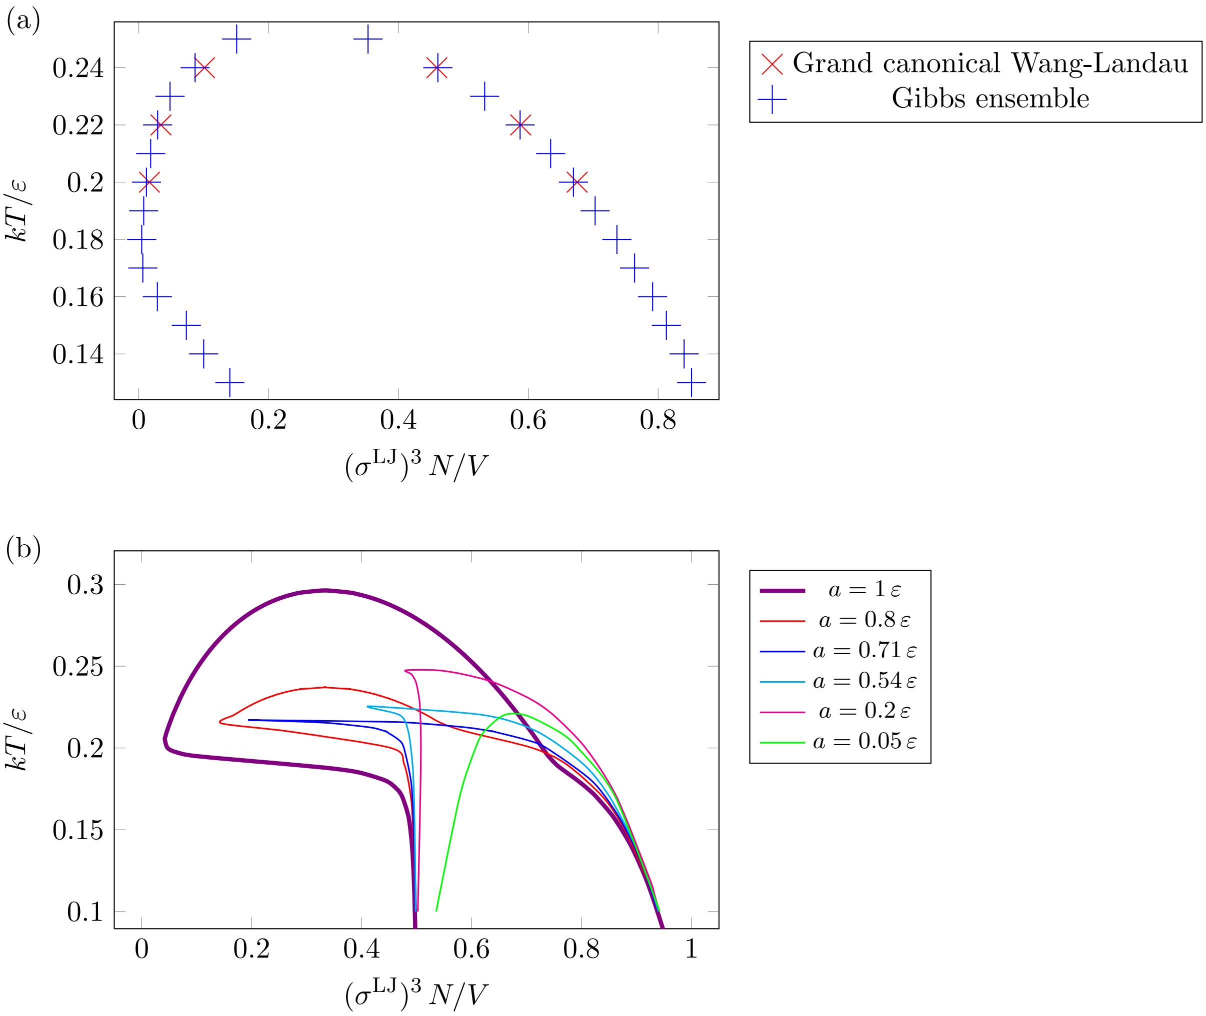

Figure 6.26: Phase diagrams for the icosahedral system for patch widths of 0.6 (above) using simulation based methods and 0.5 (below) using the theoretical model. Note the differences in behaviour when phase re-entrance occurs.

This approximation is not flexible enough to capture fully both the high- and the low-density behaviour (we choose B and b values appropriate to the high-density limit). There are also issues when the density of clusters approaches the high-density limit. The free volume approximation predicts that the translational partition function of both clusters and monomers should tend to zero as the free volume tends to zero. However, because of the large differences in size between monomers and icosahedral clusters, even when the clusters are close-packed, there would still be space available for the monomers in the gaps between the packed clusters. In other words, the translational partition function of the monomers should not tend to zero as the translational partition function of the clusters does. However, such correlation effects associated with the positions of the particles are not simple to take into account.

In our theoretical model, we also assume that the clusters pack like hard spheres. This assumption, which is responsible for the near vertical slope of the lower-density binodal in the low-temperature de-mixing regime, may be reasonable for monodisperse clusters that have a well-defined shape and are held together by specific bonds (although for our patchy particle potential, the clusters will be compressible). However, this assumed cluster incompressibility may be a significant factor in the mismatch in the detailed shape of our phase diagrams compared to those obtained in the simulations by Sciortino and co-workers [133, 134], where clusters are much more deformable, as well as not being monodisperse. In their computed phase diagrams, they observe the coexistence curve moving to higher densities with decreasing temperature considerably more gradually than we do.

Finally, by assuming we can model the monomeric liquid and vapour phases as a van der Waals fluid, we introduce another limit to our model. In this model, the critical density has a fixed value (depending only on the free volume parameter b), whereas we know that in patchy particle models, the critical point varies with both the patch width and the patch number [239, 240]. This deficiency is probably not so important for the qualitative shape of the phase diagram, especially as we are most concerned in this chapter with the behaviour at lower temperature associated with the onset of clustering.

These are perhaps some of the reasons for the lack of quantitative agreement in the shape of the diagrams. Nevertheless, we have demonstrated that our simple model allows efficient probing of phase diagrams which are in *general* agreement with both our and others' simulation results. It is easily applicable to other systems and, with the current advancements in materials chemistry, will be a useful method by which synthetic chemists are able to work out the behaviour of their patchy particle systems.

6.8 Conclusions

The utility of methods such as the Gibbs ensemble method has been shown, along with the grand canonical ensemble Wang-Landau algorithm. The limitations of the grand canonical Wang-Landau have been demonstrated, and the ways in which they can be improved has been highlighted. We have also shown that it is possible to predict the phase behaviour of patchy particles using a simple theoretical model, with good general agreement.

The existence of phase re-entrance in our system has been shown, along with some evidence of the topology of the loop. We note that the loop does not appear to close.

Chapter 7

Conclusions

7.1 Summary of aims

Our aims for this research were to find generally applicable rules for self-assembling systems. The sub-aims are summarised below:

- To explore simple models of biologically interesting systems.
- To find whether templating is an effective strategy to aid the self-assembly of monodisperse shells.
- To examine the effect of hierarchical assembly pathways, particularly whether they lead to an enhancement of self-assembly yields.
- To investigate the phase behaviour of self-assembling systems, particularly in regions of parameter space where self-assembly competes with phase separation.

7.2 Conclusions

We have clearly shown that there are a number of generally applicable rules to self-assembly. Our conclusions are summarised below for each of our chapters.

7.2.1 Encapsulation

The use of templating particles has a number of obvious advantages over other self-assembly routes:

Kinetic traps: Aggregation and other traps affect encapsulated systems less than their uncentred counterparts, and this allows the formation of otherwise inaccessible clusters. In particular, tuning the strength of the interaction with the template creates conditions where aggregation does not compete with assembly.

Robustness: Encapsulated systems are more robust to noise in the positioning of shell patches when compared to their uncentred counterparts, and are also able to form when the central particle is not isotropically attractive.

We have also shown that **simple isotropic templating** is sufficient within our model. More complicated templates with specific interactions did not lead to a better yield than globular isotropic templates. Thus, particle design is easier than might intuitively be expected.

7.2.2 Hierarchical assembly

Hierarchical assembly, although a feasible assembly mechanism, only usually gives similar maximum yields to non-hierarchical processes and has not allowed us to make significantly more complex targets. Hierarchical assembly does, however, permit some extension of the available temperature range for formation of target structure, disfavouring the formation of liquids and aggregates, although this is probably more due to the removal of other competing structures by setting more complicated interactions and patch types than the hierarchical nature of the assembly mechanism. The assembly competence of intermediates is a key factor. Our results suggest that the hierarchical nature of many biological self-assembled structures, such as protein complexes and virus capsids, may be less to do with enabling efficient assembly and more to do with evolutionary history.

7.2.3 Pore formation

We have demonstrated that our model may be applied to more complicated biological systems. However, further experimental data is needed to pin down fully the nature of the assembly mechanism, such as the self-assembly of protein complexes on a membrane. Our results showed that pore formation cannot be viewed simply as the assembly of a cyclic complex, and that the effect on the thermodynamics of the insertion of the complex into the membrane needs to be accounted for to reconcile our results with experiments.

7.2.4 Thermodynamic limits of self-assembly

The utility of methods such as the Gibbs ensemble method has been shown, along with the grand canonical ensemble Wang-Landau algorithm. However, computing the phase diagram in the presence of re-entrance is extremely challenging due to the long time scales needed to reach equilibrium, and preliminary results have only so far been obtained.

We have also shown that it is possible to predict the phase behaviour of patchy particles using a theoretical model, with good general agreement. The existence of phase re-entrance in the temperature-density phase diagram in our system has been demonstrated along with some scenarios for the possible topology of the phase diagram in the presence of re-entrance. We note that the re-entrance does not lead to a closed loop in the phase diagram.

7.3 Future work

This thesis sets up the basis for possible future work in a number of areas, some of which are detailed below.

7.3.1 Phase re-entrance

The work on phase re-entrance in clustering systems has shown the presence of interesting and complex behaviour. Further work needs to be done on finding methods of improving the convergence of the grand canonical Wang-Landau algorithm, and/or applying new approaches to these systems in order to compute the complete phase diagram. The thermodynamic behaviour of the re-entrance at low temperature is especially interesting. Further theoretical modelling with intermediate-sized species being included in the equilibrium would also be an interesting area to expand this work.

7.3.2 Internal degrees of freedom

The group of Chengde Mao have pioneered the use of DNA star motifs [245–247] to create two-dimensional honeycomb, square and hexagonal lattices [103–105] and also three-dimensional finite clusters [106] (including tetrahedra, cubes, dodecahedra, icosahedra and buckyballs). These motifs may be considered to be patchy particles with a torsional constraint in their interactions. Each arm of DNA has a sticky end, so the motifs have a defined valence but, unlike simple patchy particles, the patch positions are not fixed and the motifs are somewhat flexible. This would be an ideal area to apply the internal degree of freedom module of the program, and, with the correct coarse-graining, would allow much information to be obtained about these systems.

Bibliography

[1] A. E. Handel, A. J. Williamson, G. Disanto, L. Handunnetthi, G. Giovannoni, S. V. Ramagopalan and S. Jacobson, 'An Updated Meta-Analysis of Risk of Multiple Sclerosis Following Infectious Mononucleosis', *PLoS ONE*, **5**, e12496, 2010.

[2] A. E. Handel, A. J. Williamson, G. Disanto, R. Dobson, G. Giovannoni and S. V. Ramagopalan, 'Smoking and Multiple Sclerosis: An Updated Meta-Analysis', *PLoS ONE*, **6**, e16149, 2011.

[3] A. C. Priest, A. J. Williamson and H. M. Cartwright, 'The Applications of Artificial Neural Networks in the Identification of Quantitative Structure-Activity Relationships for Chemotherapeutic Drug Carcinogenicity', *Advances in Intelligent Data Analysis IX*, 137–146, 2010.

[4] G. Villar, A. W. Wilber, A. J. Williamson, P. Thiara, J. P. K. Doye, A. A. Louis, M. N. Jochum, A. C. F. Lewis and E. D. Levy, 'The Self-Assembly and Evolution of Homomeric Protein Complexes', *Phys. Rev. Lett.*, **102**, 118106, 2009.

[5] G. Villar, A. W. Wilber, A. J. Williamson, J. Andreani, J. P. K. Doye and A. A. Louis, 'Modelling the Self-Assembly of Tetrameric Protein Complexes', *in preparation.*

[6] A. J. Williamson, A. W. Wilber, J. P. K. Doye and A. A. Louis, 'Templated Self-Assembly of Patchy Particles', *Soft Matter*, **7**, 3423–3431, 2011.

[7] A. Reinhardt, A. J. Williamson, J. P. K. Doye, J. Carrete, L. M. Varela and A. A. Louis, 'Re-entrant Phase Behavior for Systems with Competition between Phase Separation and Self-Assembly', *J. Chem. Phys.*, **134**, 104905, 2011.

[8] J. R. Thompson, B. Cronin, H. Bayley and M. I. Wallace, 'Rapid Cooperative Assembly of a Multimeric Membrane Protein Complex', *in preparation.*

[9] A. B. Pawar and I. Kretzschmar, 'Fabrication, Assembly and Application of Patchy Particles', *Macromol. Rapid Commun.*, **31**, 150–168, 2010.

[10] T. Baba, ‘Slow Light in Photonic Crystals’, *Nat. Photonics*, **2**, 465–473, 2008.

[11] Y. Xia, B. Gates, Y. Yin and Y. Lu, ‘Monodispersed Colloidal Spheres: Old Materials with New Applications’, *Adv. Mater.*, **12**, 693–713, 2000.

[12] S. A. Khan, A. Günther, M. A. Schmidt and K. F. Jensen, ‘Microfluidic Synthesis of Colloidal Silica’, *Langmuir*, **20**, 8604–8611, 2004.

[13] S. C. Glotzer and M. J. Solomon, ‘Anisotropy of Building Blocks and Their Assembly into Complex Structures’, *Nat. Mater.*, **6**, 557–562, 2007.

[14] C. Casagrande, P. Fabre, E. Raphaël and M. Veyssié, ‘≪Janus Beads≫: Realization and Behaviour at Water/Oil Interfaces’, *Europhys. Lett.*, **9**, 251–255, 1989.

[15] O. Cayre, V. N. Paunov and O. D. Velev, ‘Fabrication of Asymmetrically Coated Colloid Particles by Microcontact Printing Techniques’, *J. Mater. Chem.*, **13**, 2445–2450, 2003.

[16] N. B. Bowden, M. Weck, I. S. Chol and G. M. Whitesides, ‘Molecule-Mimetic Chemistry and Mesoscale Self-Assembly’, *Acc. Chem. Res*, **34**, 231–238, 2001.

[17] S. H. Chen, Z. L. Wang, J. Ballato, S. H. Foulger and D. L. Carroll, ‘Monopod, Bipod, Tripod, and Tetrapod Gold Nanocrystals’, *J. Am. Chem. Soc.*, **125**, 16186–16187, 2003.

[18] S. M. Lee, Y. W. Jun, S. N. Cho and J. Cheon, ‘Single-Crystalline Star-Shaped Nanocrystals and Their Evolution: Programming the Geometry of Nano-Building Blocks’, *J. Am. Chem. Soc.*, **124**, 11244–11245, 2002.

[19] L. Manna, E. C. Scher and A. P. Alivisatos, ‘Synthesis of Soluble and Processable Rod-, Arrow-, Teardrop-, and Tetrapod-Shaped CdSe Nanocrystals’, *J. Am. Chem. Soc.*, **122**, 12700–12706, 2000.

[20] Z. Y. Tang, Y. Wang, S. Shanbhag, M. Giersig and N. A. Kotov, ‘Spontaneous Transformation of CdTe Nanoparticles into Angled Te Nanocrystals: From Particles and Rods to Checkmarks, X-marks, and Other Unusual Shapes’, *J. Am. Chem. Soc.*, **128**, 6730–6736, 2006.

[21] T. S. Ahmadi, Z. L. Wang, T. C. Green, a. Henglein and M. A. El-Sayed, ‘Shape-Controlled Synthesis of Colloidal Platinum Nanoparticles’, *Science*, **272**, 1924–1926, 1996.

[22] L. F. Gou and C. J. Murphy, 'Solution-Phase Synthesis of Cu_2O Nanocubes', *Nano Lett.*, **3**, 231–234, 2003.

[23] E. C. Greyson, J. E. Barton and T. W. Odom, 'Tetrahedral Zinc Blende Tin Sulfide Nano- and Microcrystals', *Small*, **2**, 368–371, 2006.

[24] B. R. Martin, D. J. Dermody, B. D. Reiss, M. M. Fang, L. A. Lyon, M. J. Natan and T. E. Mallouk, 'Orthogonal Self-Assembly on Colloidal Gold-Platinum Nanorods', *Adv. Mater.*, **11**, 1021–1025, 1999.

[25] X. G. Peng, L. Manna, W. D. Yang, J. Wickham, E. Scher, A. Kadavanich and A. P. Alivisatos, 'Shape Control of CdSe Nanocrystals', *Nature*, **404**, 59–61, 2000.

[26] D. Dendukuri, D. C. Pregibon, J. Collins, T. A. Hatton and P. S. Doyle, 'Continuous-flow Lithography for High-Throughput Microparticle Synthesis', *Nat. Mater.*, **5**, 365–369, 2006.

[27] D. J. Kraft, J. Groenewold and W. K. Kegel, 'Colloidal Molecules with Well-Controlled Bond Angles', *Soft Matter*, **5**, 3823–3826, 2009.

[28] Y.-S. Cho, G.-R. Yi, J.-M. Lim, S.-H. Kim, V. N. Manoharan, D. J. Pine and S.-M. Yang, 'Self-Organization of Bidisperse Colloids in Water Droplets', *J. Am. Chem. Soc.*, **127**, 15968–15975, 2005.

[29] Y.-S. Cho, G.-R. Yi, S.-H. Kim, S.-J. Jeon, M. T. Elsesser, H. K. Yu, S.-M. Yang and D. J. Pine, 'Particles with Coordinated Patches or Windows from Oil-in-Water Emulsions', *Chem. Mater.*, **19**, 3183–3193, 2007.

[30] V. N. Manoharan, M. T. Elsesser and D. J. Pine, 'Dense Packing and Symmetry in Small Clusters of Microspheres', *Science*, **301**, 483–487, 2003.

[31] S. Sacanna, W. T. M. Irvine, P. M. Chaikin and D. J. Pine, 'Lock and Key Colloids', *Nature*, **464**, 575–578, 2010.

[32] L. Wang, L. Xia, G. Li, S. Ravaine and X. S. Zhao, 'Patterning the Surfaces of Colloidal Microspheres and Fabrication of Non-Spherical Particles', *Angew. Chem. Int. Ed.*, **47**, 4725–4728, 2008.

[33] S. John, 'Strong Localization of Photons in Certain Disordered Dielectric Superlattices', *Phys. Rev. Lett.*, **58**, 2486–2489, 1987.

[34] E. Yablonovitch, 'Inhibited Spontaneous Emission in Solid-State Physics and Electronics', *Phys. Rev. Lett.*, **58**, 2059–2062, 1987.

[35] A.-P. Hynninen, J. H. J. Thijssen, E. C. M. Vermolen, M. Dijkstra and A. van Blaaderen, 'Self-Assembly Route for Photonic Crystals with a Bandgap in the Visible Region', *Nat. Mater.*, **6**, 202–205, 2007.

[36] Z. Zhang, A. S. Keys, T. Chen and S. C. Glotzer, 'Self-Assembly of Patchy Particles into Diamond Structures Through Molecular Mimicry', *Langmuir*, **21**, 11547–11551, 2005.

[37] S. Smith, 'Optical Bistability, Photonic Logic, and Optical Computation', *Appl. Opt.*, **25**, 1550–1564, 1986.

[38] X. Peng, 'An Essay on Synthetic Chemistry of Colloidal Nanocrystals', *Nano Res.*, **2**, 425–447, 2009.

[39] M. Yin, C. Wu, Y. Lou, C. Burda, J. T. Koberstein, Y. Zhu and S. O'Brien, 'Copper Oxide Nanocrystals', *J. Am. Chem. Soc.*, **127**, 9506–9511, 2005.

[40] R. P. Sear, 'Phase Behavior of a Simple Model of Globular Proteins', *J. Chem. Phys.*, **111**, 4800, 1999.

[41] J. E. Andréani. *Modelling the Evolution of Homomeric Protein Complexes.* Master's thesis, University of Oxford & École Polytechnique, Paris, UK & France, 2009.

[42] M. Llewleyn-Jones. *Positive Design for Self-Assembly: The Formation of Helical Polymers.* Master's thesis, University of Oxford, UK, 2008.

[43] J. P. K. Doye, A. A. Louis, I. C. Lin, L. R. Allen, E. G. Noya, A. W. Wilber, H. C. Kok and R. Lyus, 'Controlling Crystallization and its Absence: Proteins, Colloids and Patchy Models', *Phys. Chem. Chem. Phys.*, **9**, 2197–2205, 2007.

[44] L. Stryer. *Biochemistry. 4th Edition.* W. H. Freeman and Company, New York, 1998. ISBN 0716720094.

[45] F. Eisenhaber, B. Persson and P. Argos, 'Protein Structure Prediction: Recognition of Primary, Secondary, and Tertiary Structural Features from Amino Acid Sequence', *Crit. Rev. Biochem. Mol. Biol.*, **30**, 1–94, 1995.

[46] A. W. Wilber. *Simulation Studies of Monodisperse Self-Assembly.* Doctoral thesis, University of Oxford, UK, 2009.

[47] R. Jaenicke, 'Folding and Association versus Misfolding and Aggregation of Proteins', *Phil. Trans. R. Soc. Lond.*, **348**, 97–105, 1995.

[48] M. G. Mateu, M. M. S. D. Pino and A. R. Fersht, 'Mechanism of Folding and Assembly of a Small Tetrameric Protein Domain from Tumor Suppressor P53', *Nat. Struct. Biol.*, **6**, 191–198, 1999.

[49] M. Jones and G. Jones. *Advanced Biology.* Cambridge University Press, 1997. ISBN 0521484731.

[50] S. Jones and J. M. Thornton, 'Protein-Protein Interactions: A Review of Protein Dimer Structures.', *Prog. Biophys. Mol. Biol.*, **63**, 31, 1995.

[51] N. T. Southall, K. A. Dill and A. D. J. Haymet, 'A View of the Hydrophobic Effect', *J. Phys. Chem. B*, **106**, 521–533, 2002.

[52] M. Schaefer, M. Sommer and M. Karplus, 'pH-Dependence of Protein Stability: Absolute Electrostatic Free Energy Differences between Conformations', *J. Phys. Chem. B*, **101**, 1663–1683, 1997.

[53] R. H. Abeles and D. Dolphin, 'The Vitamin B12 Coenzyme', *Acc. Chem. Res.*, **9**, 114–120, 1976.

[54] W. Eaton, E. Henry, J. Hofrichter and A. Mozzarelli, 'Is Cooperative Oxygen Binding by Hemoglobin Really Understood?', *Rendiconti Lincei*, **17**, 147–162, 2006.

[55] A. Lwoff, 'The Concept of Virus', *J. Gen. Microbiol.*, **17**, 239–253, 1957.

[56] C. A. Suttle, 'Viruses in the Sea', *Nature*, **437**, 356–361, 2005.

[57] G. Rice, L. Tang, K. Stedman, F. Roberto, J. Spuhler, E. Gillitzer, J. E. Johnson, T. Douglas and M. Young, 'The Structure of a Thermophilic Archaeal Virus Shows a Double-Stranded DNA Viral Capsid Type that Spans All Domains of Life', *Proc. Natl. Acad. Sci. USA*, **101**, 7716–7720, 2004.

[58] B. Böttcher and R. A. Crowther, 'Difference Imaging Reveals Ordered Regions of RNA in Turnip Yellow Mosaic Virus', *Structure*, **4**, 387–394, 1996.

[59] J. C. Venter, M. D. Adams, E. W. Myers, P. W. Li and R. J. Mural *et al*, 'The Sequence of the Human Genome', *Science*, **291**, 1304–1351, 2001.

[60] V. C. Wasinger, J. D. Pollack and I. Humphery-Smith, 'The Proteome of *Mycoplasma genitalium*', *Eur. J. Biochem.*, **267**, 1571–1582, 2000.

[61] R. S. Thomas, 'The Chemical Composition and Particle Weight of *Tipula* Iridescent Virus', *Virology*, **14**, 240–252, 1961.

[62] D. L. Caspar and A. Klug, 'Physical Principles in the Construction of Regular Viruses', *Cold Spring Harbor Symp. Quant. Biol.*, **27**, 1–24, 1962.

[63] L.-F. He, D. Alling, T. Popkin, M. Shapiro, H. J. Alter and R. H. Purcell, 'Determining the Size of Non-A, Non-B Hepatitis Virus by Filtration', *J. Infect. Dis.*, **156**, 636–640, 1987.

[64] F. J. Trueba and C. L. Woldringh, 'Changes in Cell Diameter During the Division Cycle of *Escherichia coli*', *J. Bacteriol.*, **142**, 869–878, 1980.

[65] U. Smith, L. Sjöström and P. Bjöntorp, 'Comparison of Two Methods for Determining Human Adipose Cell Size', *J. Lipid Res.*, **13**, 822–824, 1972.

[66] D. A. Steinhauer and J. J. Holland, 'Rapid Evolution of RNA Viruses', *Annu. Rev. Microbiol.*, **41**, 409–431, 1987.

[67] A. Klug, 'The Tobacco Mosaic Virus Particle: Structure and Assembly', *Philos. T. Roy. Soc. B.*, **354**, 531–535, 1999.

[68] H. R. Gelderblom, 'Assembly and Morphology of HIV: Potential Effect of Structure on Viral Function', *AIDS*, **5**, 617, 1991.

[69] T. Loeb and N. D. Zinder, 'A Bacteriophage Containing RNA', *Proc. Natl. Acad. Sci. USA*, **47**, 282, 1961.

[70] 'VIrus Particle ExploreR Database'. URL `http://viperdb.scripps.edu/`.

[71] S. Casjens and J. King, 'Virus Assembly', *Annu. Rev. Biochem.*, **44**, 555–611, 1975.

[72] C. Sachse, J. Z. Chen, P.-D. Coureux, M. E. Stroupe, M. Fändrich and N. Grigorieff, 'High-Resolution Electron Microscopy of Helical Specimens: A Fresh Look at Tobacco Mosaic Virus', *J. Mol. Biol.*, **371**, 812–835, 2007.

[73] P. L. Freddolino, A. S. Arkhipov, S. B. Larson, A. McPherson and K. Schulten, 'Molecular Dynamics Simulations of the Complete Satellite Tobacco Mosaic Virus', *Structure*, **14**, 437–449, 2006.

[74] X. D. Yan, P. R. Chipman, T. Castberg, G. Bratbak and T. S. Baker, 'The Marine Algal Virus PpV01 has an Icosahedral Capsid with T=219 Quasisymmetry', *J. Virol.*, **79**, 9236–9243, 2005.

[75] D. G. Angelescu and P. Linse, 'Viruses as Supramolecular Self-Assemblies: Modelling of Capsid Formation and Genome Packaging', *Soft Matter*, **4**, 1981–1990, 2008.

[76] S. D. Moore and P. E. Prevelige, Jr., ‘DNA Packaging: A New Class of Molecular Motors’, *Curr. Biol.*, **12**, R96–R98, 2002.

[77] M. G. Weinbauer, I. Brettar and M. G. Höfle, ‘Lysogeny and Virus-Induced Mortality of Bacterioplankton in Surface, Deep, and Anoxic Marine Waters’, *Limnol. Oceanogr.*, **48**, 1457–1465, 2003.

[78] P. E. Prevelige, ‘Inhibiting Virus-Capsid Assembly by Altering the Polymerisation Pathway’, *Trends Biotechnol.*, **16**, 61–65, 1998.

[79] S. J. Stray, J. M. Johnson, B. G. Kopek and A. Zlotnick, ‘An *in Vitro* Fluorescence Screen to Identify Antivirals that Disrupt Hepatitis B Virus Capsid Assembly’, *Nat. Biotechnol.*, **24**, 358–362, 2006.

[80] A. Zlotnick, A. Lee, C. R. Bourne, J. M. Johnson, P. L. Domanico and S. J. Stray, ‘*In Vitro* Screening for Molecules that Affect Virus Capsid Assembly (and Other Protein Association Reactions)’, *Nat. Protoc.*, **2**, 490–498, 2007.

[81] L. A. Koutsky, K. A. Ault, C. M. Wheeler, D. R. Brown, E. Barr, F. B. Alvarez, L. M. Chiacchierini and K. U. Jansen, ‘A Controlled Trial of a Human Papillomavirus Type 16 Vaccine’, *New Engl. J. Med.*, **347**, 1645–1651, 2002.

[82] R. Kirnbauer, F. Booy, N. Cheng, D. R. Lowy and J. T. Schiller, ‘Papillomavirus L1 Major Capsid Protein Self-Assembles into Virus-Like Particles that are Highly Immunogenic’, *Proc. Natl. Acad. Sci. USA*, **89**, 12180–12184, 1992.

[83] M. Shank-Retzlaff, F. Wang, T. Morley, C. Anderson, M. Hamm, M. Brown, K. Rowland, G. Pancari, J. Zorman, R. Lowe, L. Schultz, J. Teyral, R. Capen, C. B. Oswald, Y. Wang, M. Washabaugh, K. Jansen and R. Sitrin, ‘Correlation between Mouse Potency and *in Vitro* Relative Potency for Human Papillomavirus Type 16 Virus-Like Particles and Gardasil Vaccine Samples’, *Hum. Vaccin.*, **1**, 191, 2005.

[84] T. Douglas and M. Young, ‘Viruses: Making Friends with Old Foes’, *Science*, **312**, 873–875, 2006.

[85] P. Singh, G. Destito, A. Schneemann and M. Manchester, ‘Canine Parvovirus-Like Particles, a Novel Nanomaterial for Tumor Targeting’, *J. Nanobiotechnol.*, **4**, 2, 2006.

[86] P. Singh, M. J. Gonzalez and M. Manchester, ‘Viruses and Their Uses in Nanotechnology’, *Drug Dev. Res.*, **67**, 23–41, 2006.

[87] M. Young, D. Willits, M. Uchida and T. Douglas, ‘Plant Viruses as Biotemplates for Materials and Their Use in Nanotechnology’, *Annu. Rev. Phytopathol.*, **46**, 361–384, 2008.

[88] K. E. Sapsford, C. M. Soto, A. S. Blum, A. Chatterji, T. Lin, J. E. Johnson, F. S. Ligler and B. R. Ratna, ‘A Cowpea Mosaic Virus Nanoscaffold for Multiplexed Antibody Conjugation: Application as an Immunoassay Tracer’, *Biosens. Bioelectron.*, **21**, 1668–1673, 2006.

[89] Z. Boldogköi, A. Sík, Ádám Dénes, A. Reichart, J. Toldi, I. Gerendai, K. J. Kovács and M. Palkovits, ‘Novel Tracing Paradigms–Genetically Engineered Herpesviruses as Tools for Mapping Functional Circuits Within the CNS: Present Status and Future Prospects’, *Prog. Neurobiol.*, **72**, 417–445, 2004.

[90] World Health Organization. ‘Fact Sheet 310: The Top Ten Causes of Death’. 2008. URL `http://www.who.int/entity/mediacentre/factsheets/fs310_2008.pdf`.

[91] M. A. Nowak, S. Bonhoeffer, G. M. Shaw and R. M. May, ‘Anti-Viral Drug Treatment: Dynamics of Resistance in Free Virus and Infected Cell Populations’, *J. Theor. Biol.*, **184**, 203–217, 1997.

[92] W. M. Stanley, ‘Isolation of a Crystalline Protein Assessing the Properties of Tobacco-Mosaic Virus’, *Science*, **81**, 644–645, 1935.

[93] P. K. Sorger, P. G. Stockley and S. C. Harrison, ‘Structure and Assembly of Turnip Crinkle Virus : II. Mechanism of Reassembly in Vitro’, *J. Mol. Biol.*, **191**, 639–658, 1986.

[94] K. N. Johnson, L. Tang, J. E. Johnson and L. A. Ball, ‘Heterologous RNA Encapsidated in Pariacoto Virus-Like Particles Forms a Dodecahedral Cage Similar to Genomic RNA in Wild-Type Virions’, *J. Virol.*, **78**, 11371–11378, 2004.

[95] J. M. Johnson, D. A. Willits, M. J. Young and A. Zlotnick, ‘Interaction with Capsid Protein Alters RNA Structure and the Pathway for *in Vitro* Assembly of Cowpea Chlorotic Mottle Virus’, *J. Mol. Biol.*, **335**, 455–464, 2004.

[96] F. L. Homa and J. C. Brown, ‘Capsid Assembly and DNA Packaging in Herpes Simplex Virus’, *Rev. Med. Virol.*, **7**, 107–122, 1997.

[97] J. M. Hogle, M. Chow and D. J. Filman, ‘Three-Dimensional Structure of Poliovirus at 2.9 Å Resolution’, *Science*, **229**, 1358–1365, 1985.

[98] D. S. Goodsell and A. J. Olson, ‘Structural Symmetry and Protein Function’, *Annu. Rev. Biophys. Biomol. Struct.*, **29**, 105–153, 2000.

[99] F. F. V. Chevance and K. T. Hughes, ‘Coordinating Assembly of a Bacterial Macromolecular Machine’, *Nat. Rev. Microbiol.*, **6**, 455–465, 2008.

[100] D. S. Goodsell. *Bionanotechnology.* Wiley-Liss, Hoboken, 2004.

[101] K. Namba. ‘Self-Assembly of a Bacterial Flagellum Motor’. Accessed 10-1-2011, URL `http://www.fbs.osaka-u.ac.jp/labs/namba/npn/movie5.html`.

[102] R. P. Goodman, R. M. Berry and A. J. Turberfield, ‘The Single-Step Synthesis of a DNA Tetrahedron’, *Chem. Commun.*, 1372–1373, 2004.

[103] Y. He, Y. Chen, H. Liu, A. E. Ribbe and C. Mao, ‘Self-Assembly of Hexagonal DNA Two-Dimensional (2D) Arrays’, *J. Am. Chem. Soc.*, **127**, 12202–12203, 2005.

[104] Y. He, Y. Tian, A. E. Ribbe and C. Mao, ‘Highly Connected Two-Dimensional Crystals of DNA Six-Point-Stars’, *J. Am. Chem. Soc.*, **128**, 15978–15979, 2006.

[105] Y. He, Y. Tian, Y. Chen, A. E. Ribbe and C. Mao, ‘Geometric self-sorting in DNA self-assembly’, *Chem. Commun.*, 165–167, 2007.

[106] Y. He, T. Ye, C. Zhang, A. E. Ribbe, W. Jiang and C. Mao, ‘Hierarchical Self-Assembly of DNA into Symmetric Supramolecular Polyhedra’, *Nature*, **452**, 198–201, 2008.

[107] K. Wade, ‘Structural and Bonding Patterns in Cluster Chemistry’, *Advances in Inorganic Chemistry*, **18**, 1–66, 1976.

[108] S. R. Seidel and P. J. Stang, ‘High-Symmetry Coordination Cages via Self Assembly’, *Acc. Chem. Res*, **35**, 972–983, 2002.

[109] B. T. Burnley and J. P. L. Cox, ‘An Efficient Biomimetic Assembly of a Macroscopic Polyhedral Shell from Identical Subunits’, *Mat. Sci. Eng. C*, **25**, 529–540, 2005.

[110] J. E. Andréani and J. P. K. Doye, ‘Simulating the Evolution of Homomeric Protein Complexes’, *in preparation.*

[111] A. W. Wilber, J. P. K. Doye, A. A. Louis, E. G. Noya, M. A. Miller and P. Wong, ‘Reversible Self-Assembly of Patchy Particles into Monodisperse Icosahedral Clusters’, *J. Chem. Phys.*, **127**, 085106, 2007.

[112] A. W. Wilber, J. P. K. Doye and A. A. Louis, 'Self-Assembly of Monodisperse Clusters: Dependence on Target Geometry', *J. Chem. Phys.*, **131**, 175101, 2009.

[113] A. Reinhardt. *Computer Simulation of Homogeneous Ice Nucleation.* Master's thesis, University of Oxford, UK, 2009.

[114] C. Valeriani, R. J. Allen, M. J. Morelli, D. Frenkel and P. R. ten Wolde, 'Computing Stationary Distributions in Equilibrium and Nonequilibrium Systems with Forward Flux Sampling', *J. Chem. Phys.*, **127**, 114109, 2007.

[115] M. F. Hagan and D. Chandler, 'Dynamic Pathways for Viral Capsid Assembly', *Biophys. J.*, **91**, 42–54, 2006.

[116] H. D. Nguyen, V. S. Reddy and C. L. Brooks III, 'Deciphering the Kinetic Mechanism of Spontaneous Self-Assembly of Icosahedral Capsids', *Nano Lett.*, **7**, 338–344, 2007.

[117] H. D. Nguyen and C. L. Brooks III, 'Generalized Structural Polymorphism in Self-Assembled Viral Particles', *Nano Lett.*, **8**, 4574–4581, 2008.

[118] H. D. Nguyen, V. S. Reddy and C. L. Brooks III, 'Invariant Polymorhpism in Viral Capsid Assembly', *J. Am. Chem. Soc.*, **131**, 2606–2614, 2009.

[119] T. Zhang and R. Schwartz, 'Simulation Study of the Contribution of Oligomer/Oligomer Binding to Capsid Assembly Kinetics', *Biophys. J.*, **90**, 57–64, 2006.

[120] B. Sweeney, T. Zhang and R. Schwartz, 'Exploring the Parameter Space of Complex Self-Assembly Through Virus Capsid Models', *Biophys. J.*, **94**, 772–783, 2008.

[121] D. C. Rapaport, 'Self-Assembly of Polyhedral Shells: A Molecular Dynamics Study', *Phys. Rev. E*, **70**, 051905, 2004.

[122] D. C. Rapaport, 'Role of Reversibility in Viral Capsid Growth: A Paradigm for Self-Assembly', *Phys. Rev. Lett.*, **101**, 186101, 2008.

[123] I. G. Johnston, A. A. Louis and J. P. K. Doye, 'Modelling the Self-Assembly of Virus Capsids', *J. Phys.: Condens. Matter*, **22**, 104101, 2010.

[124] D. J. Wales, 'The Energy Landscape as a Unifying Theme in Molecular Science', *Phil. Trans. R. Soc. A*, **363**, 357–377, 2005.

[125] S. N. Fejer, T. R. James, J. Hernandez-Rojas and D. J. Wales, 'Energy Landscapes for Shells Assembled from Pentagonal and Hexagonal Pyramids', *Phys. Chem. Chem. Phys.*, **11**, 2098–2104, 2009.

[126] S. N. Fejer, D. Chakrabarti and D. J. Wales, 'Emergent Complexity from Simple Anisotropic Building Blocks: Shells, Tubes, and Spirals', *ACS Nano*, **4**, 219–228, 2010.

[127] M. F. Hagan, 'Controlling Viral Capsid Assembly with Templating', *Phys. Rev. E*, **77**, 051904, 2008.

[128] O. M. Elrad and M. F. Hagan, 'Mechanisms of Size Control and Polymorphism in Viral Capsid Assembly', *Nano Lett.*, **8**, 3850–3857, 2008.

[129] A. Kivenson and M. F. Hagan, 'Mechanisms of Capsid Assembly Around a Polymer', *Biophys. J.*, **99**, 619, 2010.

[130] O. M. Elrad and M. F. Hagan, 'Encapsulation of a polymer by an icosahedral virus', *Physical Biology*, **7**, 045003, 2010.

[131] N. Ghofraniha, P. Andreozzi, J. Russo, C. La Mesa and F. Sciortino, 'Assembly Kinetics in Binary Mixtures of Strongly Attractive Colloids', *J. Phys. Chem. B*, **113**, 6775–6781, 2009.

[132] F. Romano, E. Sanz and F. Sciortino, 'Role of the Range in the Fluid-Crystal Coexistence for a Patchy Particle Model', *J. Phys. Chem. B*, **113**, 15133, 2009.

[133] F. Sciortino, A. Giacometti and G. Pastore, 'Phase Diagram of Janus Particles', *Phys. Rev. Lett.*, **103**, 237801, 2009.

[134] F. Sciortino, A. Giacometti and G. Pastore, 'A Numerical Study of One-Patch Colloidal Particles: From Square-Well to Janus', *Phys. Chem. Chem. Phys.*, **12**, 11869–11877, 2010.

[135] D. Endres and A. Zlotnick, 'Model-Based Analysis of Assembly Kinetics for Virus Capsids or Other Spherical Polymers', *Biophys. J.*, **83**, 1217–1230, 2002.

[136] D. Endres, M. Miyahara, P. Moisant and A. Zlotnick, 'A Reaction Landscape Identifies the Intermediates Critical for Self-Assembly of Virus Capsids and Other Polyhedral Structures', *Protein Sci.*, **14**, 1518–1525, 2005.

[137] R. Zandi, P. van der Schoot, D. Reguera, W. Kegel and H. Reiss, 'Classical Nucleation Theory of Virus Capsids', *Biophys. J.*, **90**, 1939–1948, 2006.

[138] T. Zhang and R. Schwartz, 'Investigating Scaling Effects on Virus Capsid-Like Self-Assembly Using Discrete Event Simulations', *IEEE T. Nanobio.*, **6**, 235–241, 2007.

[139] A. Zlotnick, 'Theoretical Aspects of Virus Capsid Assembly', *J. Mol. Recog.*, **18**, 479–490, 2005.

[140] P. van der Schoot and R. Bruinsma, 'Electrostatics and the Assembly of an RNA Virus', *Phys. Rev. E*, **71**, 061928, 2005.

[141] T. Hu and B. I. Shklovskii, 'Kinetics of Viral Self-Assembly: Role of the Single-Stranded RNA Antenna', *Phys. Rev. E*, **75**, 051901, 2007.

[142] M. F. Hagan, 'A Theory for Viral Capsid Assembly Around Electrostatic Cores', *J. Chem. Phys.*, **130**, 114902, 2009.

[143] R. Fantoni, A. Giacometti, F. Sciortino and G. Pastore, 'Cluster Theory of Janus Particles', *Soft Matter, Advanced Article; arXiv:1012.1820*, 2011.

[144] D. Frenkel and B. Smit. *Understanding Molecular Simulation, from Algorithms to Applications.* Academic Press (Elsevier), 2002.

[145] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. N. Teller and E. Teller, 'Equation of State Calculations by Fast Computing Machines', *J. Chem. Phys.*, **21**, 1087–1092, 1953.

[146] S. Whitelam and P. L. Geissler, 'Avoiding Unphysical Kinetic Traps in Monte Carlo Simulations of Strongly Attractive Particles', *J. Chem. Phys.*, **127**, 154101, 2007.

[147] S. Chib, F. Nardari and N. Shephard, 'Markov Chain Monte Carlo Methods for Stochastic Volatility Models', *Journal of Econometrics*, **108**, 281–316, 2002.

[148] F. Wang and D. P. Landau, 'Efficient, Multiple-Range Random Walk Algorithm to Calculate the Density of States', *Phys. Rev. Lett.*, **86**, 2050–2053, 2001.

[149] G. Ganzenmüller and P. J. Camp, 'Applications of Wang-Landau Sampling to Determining Phase Equilibria in Complex Fluids', *J. Chem. Phys.*, **127**, 154504, 2007.

[150] A. Louis and M. Sprik. *Computer Simulation Methods in Chemistry and Physics (lecture notes).* University of Cambridge, Cambridge, 2005. URL `http://www-thphys.physics.ox.ac.uk/user/ArdLouis/teaching/M10Mich05.pdf`.

[151] D. A. McQuarrie. *Statistical Mechanics.* University Science Books, Sausalito, California, 2000. ISBN 1891389157.

[152] L. Berthier and W. Kob, 'The Monte Carlo Dynamics of a Binary Lennard-Jones Glass-Forming Mixture', *J. Phys.: Condens. Matter*, **19**, 205130, 2007.

[153] H. E. A. Huitema and J. P. van der Eerden, ‘Can Monte Carlo Simulation Describe Dynamics? A Test on Lennard-Jones Systems’, *J. Chem. Phys.*, **110**, 3267–3274, 1999.

[154] M. P. Allen. *Introduction to Molecular Dynamics Simulation*, volume 23, 1–28. John von Neumann Institute for Computing, Jülich, 2004. ISBN 3000126414.

[155] V. I. Manousiouthakis and M. W. Deem, ‘Strict Detailed Balance is Unnecessary in Monte Carlo Simulation’, *J. Chem. Phys.*, **110**, 2753, 1999.

[156] N. Metropolis, ‘The Beginning of the Monte Carlo Method’, *Los Alamos Sci.*, **15**, 125–130, 1987.

[157] J. P. Sethna. *Statistical Mechanics: Entropy, Order Parameters, and Complexity.* Oxford University Press, USA, 2006. ISBN 0198566778.

[158] A. Maczek. *Statistical Thermodynamics.* Oxford University Press, 1998. ISBN 0198559119.

[159] R. Kubo, N. Saitō and M. Toda. *Statistical Physics: Equilibrium Statistical Mechanics.* Springer, 1998. ISBN 3540536620.

[160] D. Thirumalai, R. D. Mountain and T. Kirkpatrick, ‘Ergodic Behavior in Supercooled Liquids and in Glasses’, *Phys. Rev. A*, **39**, 3563–3574, 1989.

[161] G. B. Arfken and H. J. Weber. *Mathematical Methods for Physicists.* Elsevier Academic, London, 6th edition, 2005. ISBN 9780120885848.

[162] D. M. Smith, ‘Efficient Multiple-Precision Evaluation of Elementary Functions’, *Math. Comp.*, **52**, 131–134, 1989.

[163] F. J. Vesely, ‘Angular Monte Carlo Integration Using Quaternion Parameters: A Spherical Reference Potential for CCl_4’, *J. Comput. Phys.*, **47**, 291–296, 1982.

[164] M. Geradin and A. Cardona, ‘Kinematics and Dynamics of Rigid and Flexible Mechanisms Using Finite Elements and Quaternion Algebra’, *Comput. Mech.*, **4**, 115–135, 1989.

[165] E. W. Weisstein. ‘Quaternion’. Accessed 8-12-2010, URL `http://mathworld.wolfram.com/Quaternion.html`.

[166] S. M. Johnson, J. R. Williams and B. K. Cook, ‘Quaternion-Based Rigid Body Rotation Integration Algorithms for Use in Particle Methods’, *Int. J. Numer. Meth. Eng.*, **74**, 1303–1313, 2008.

[167] G. Marsaglia, A. Zaman and W. W. Tsang, ‘Toward a Universal Random Number Generator’, *Statistics & Probability Letters*, **9**, 35–39, 1990.

[168] R. Vink and J. Horbach, ‘Grand Canonical Monte Carlo Simulation of a Model Colloid–Polymer Mixture: Coexistence Line, Critical Behavior, and Interfacial Tension’, *J. Chem. Phys.*, **121**, 3253, 2004.

[169] K. K. Han and H. S. Son, ‘On the Isothermal-Isobaric Ensemble Partition Function’, *J. Chem. Phys.*, **115**, 7793, 2001.

[170] I. R. McDonald, ‘NpT-Ensemble Monte Carlo Calculations for Binary Liquid Mixtures’, *Mol. Phys.*, **23**, 41–58, 1972.

[171] R. Eppenga and D. Frenkel, ‘Monte Carlo Study of the Isotropic and Nematic Phases of Infinitely Thin Hard Platelets’, *Mol. Phys.*, **52**, 1303–1334, 1984.

[172] S. Whitelam and P. L. Geissler, ‘Erratum: “Avoiding Unphysical Kinetic Traps in Monte Carlo Simulations of Strongly Attractive Particles”’, *J. Chem. Phys.*, **128**, 219901, 2008.

[173] D. Frenkel, ‘Speed-Up of MC Simulations by Sampling of Rejected States’, *Proc. Nat. Acad. Sci. USA*, **101**, 17571–17575, 2004.

[174] G. M. Torrie and J. P. Valleau, ‘Nonphysical Sampling Distributions in Monte Carlo Free-Energy Estimation: Umbrella Sampling’, *J. Comput. Phys.*, **23**, 187–199, 1977.

[175] P. R. Ten Wolde and D. Frenkel, ‘Computer Simulation Study of Gas–Liquid Nucleation in a Lennard-Jones System’, *J. Chem. Phys.*, **109**, 9901, 1998.

[176] D. Chandler. *Introduction to Modern Statistical Mechanics.* Oxford University Press, New York, 1987. ISBN 9870195042771.

[177] R. E. Belardinelli and V. D. Pereyra, ‘Fast Algorithm to Calculate Density of States’, *Phys. Rev. E*, **75**, 2007.

[178] G. Villar. *Positive Design for Self-Assembly: Analysis of Tetrameric Protein Complexes.* Master’s thesis, University of Oxford, UK, 2008.

[179] J. K. Johnson, J. A. Zollweg and K. E. Gubbins, ‘The Lennard-Jones Equation of State Revisited’, *Mol. Phys.*, **78**, 591–618, 1993.

[180] B. Smit. Doctoral thesis, University of Utrecht, Netherlands, 1990.

[181] A. Z. Panagiotopoulos, N. Quirke, M. Stapleton and D. J. Tildesley, ‘Phase Equilibria by Simulation in the Gibbs Ensemble – Alternative Derivation, Generalization and Application to Mixture and Membrane Equilibria’, *Mol. Phys.*, **63**, 1988.

[182] C. Chen, M.-C. Daniel, Z. T. Quinkert, M. De, B. Stein, V. D. Bowman, P. R. Chipman, V. M. Rotello, C. C. Kao and B. Dragnea, ‘Nanoparticle-Templated Assembly of Viral Protein Cages’, *Nano Lett.*, **6**, 611–615, 2006.

[183] J. Sun, C. DuFort, M.-C. Daniel, A. Murali, C. i Chen, K. Gopinath, B. Stein, M. De, V. M. Rotello, A. Holzenburg, C. C. Kao and B. Dragnea, ‘Core-Controlled Polymorphism in Virus-Like Particles’, *Proc. Natl. Acad. Sci. USA*, **104**, 1354–1359, 2007.

[184] B. Dragnea, C. Chen, E.-S. Kwak, B. Stein and C. C. Kao, ‘Gold Nanoparticles as Spectroscopic Enhancers for *in Vitro* Studies on Single Viruses’, *J. Am. Chem. Soc.*, **125**, 6374–6375, 2003.

[185] C. B. Chang, C. M. Knobler, W. M. Gelbart and T. G. Mason, ‘Curvature Dependence of Viral Protein Structures on Encapsidated Nanoemulsion Droplets’, *ACS Nano*, **109**, 281–286, 2008.

[186] Y. Hu, R. Zandi, A. Anavitarte, C. M. Knobler and W. M. Gelbart, ‘Packaging of a Polymer by a Viral Capsid: The Interplay between Polymer Length and Capsid Size’, *Biophys J.*, **94**, 1428–1436, 2008.

[187] F. Caruso, ‘Hollow Capsule Processing Through Colloidal Templating and Self-Assembly’, *Chem.-Eur. J.*, **6**, 413–419, 2000.

[188] R. K. Iler, ‘Multilayers of Colloidal Particles’, *J. Colloid Interface Sci.*, **21**, 569–594, 1966.

[189] Y. A. Vlasov, X.-Z. Bo, J. C. Sturm and D. J. Norris, ‘On-Chip Natural Assembly of Silicon Photonic Bandgap Crystals’, *Nature*, **414**, 289–293, 2001.

[190] K. P. Velikov, C. G. Christova, R. P. A. Dullens and A. van Blaaderen, ‘Layer-by-Layer Growth of Binary Colloidal Crystals’, *Science*, **296**, 106–109, 2002.

[191] Z. Zhang and S. C. Glotzer, ‘Self-Assembly of Patchy Particles’, *Nano Lett.*, **4**, 1407–1413, 2004.

[192] A. W. Wilber, J. P. K. Doye, A. A. Louis and A. C. F. Lewis, ‘Monodisperse Self-Assembly in a Model with Protein-Like Interactions’, *J. Chem. Phys.*, **131**, 175102, 2009.

[193] S. E. Aniagyei, C. DuFort, C. C. Kao and B. Dragnea, 'Self-Assembly Approaches to Nanomaterial Encapsulation in Viral Protein Cages', *J. Mater. Chem.*, **18**, 3763, 2008.

[194] I. Brigger, C. Dubernet and P. Couvreur, 'Nanoparticles in Cancer Therapy and Diagnosis', *Adv. Drug Delivery Rev.*, **54**, 631–651, 2002.

[195] D. Bhatia, S. Mehtab, R. Krishnan, S. Indi, A. Basu and Y. Krishnan, 'Icosahedral DNA Nanocapsules by Modular Assembly', *Angew. Chem. Int. Ed.*, **121**, 4198–4201, 2009.

[196] R. L. Garcea and L. Gissmann, 'Virus-Like Particles as Vaccines and Vessels for the Delivery of Small Molecules', *Curr. Opin. Biotechnol.*, **15**, 513–517, 2004.

[197] D. P. Nayak, E. K. W. Hui and S. Barman, 'Assembly and Budding of Influenza Virus', *Virus Res.*, **106**, 147–165, 2004.

[198] C. Desgranges and J. Delhommelle, 'Phase Equilibria of Molecular Fluids via Hybrid Monte Carlo Wang–Landau Simulations: Applications to Benzene and n-Alkanes', *J. Chem. Phys.*, **130**, 244109, 2009.

[199] M. M. Maye, D. Nykypanchuk, M. Cuisinier, D. van der Lelie and O. Gang, 'Stepwise Surface Encoding for High-Throughput Assembly of Nanoclusters', *Nat. Mater.*, **8**, 388–391, 2009.

[200] J. A. A. W. Elemans, A. E. Rowan and R. J. M. Nolte, 'Mastering Molecular Matter. Supramolecular Architectures by Hierarchical Self-Assembly', *J. Mater. Chem.*, **13**, 2661–2670, 2003.

[201] E. D. Levy, E. B. Erba, C. V. Robinson and S. A. Teichmann, 'Assembly Reflects Evolution of Protein Complexes', *Nature*, **453**, 1262–1265, 2008.

[202] A. Zlotnick and S. J. Stray, 'How Does Your Virus Grow? Understanding and Interfering with Virus Assembly', *Trends Biotechnol.*, **21**, 536–542, 2003.

[203] A. Klug, 'Vom Makromolekül zum biologischen Molekülverband (Nobel-Vortrag)', *Angew. Chem.*, **95**, 579–596, 1983.

[204] C. M. Erben, R. P. Goodman and A. J. Turberfield, 'A Self-Assembled DNA Bipyramid', *J. Am. Chem. Soc.*, **129**, 6992–6993, 2008.

[205] C. M. Erben, R. P. Goodman and A. J. Turberfield, 'Single-Molecule Protein Encapsulation in a Rigid DNA Cage', *Angew. Chem. Int. Edit.*, **129**, 7414–7417, 2006.

[206] M. M. Maye, M. T. Kumara, D. Nykypanchuk, W. B. Sherman and O. Gang, 'Switching Binary States of Nanoparticle Superlattices and Dimer Clusters by DNA Strands', *Nat. Nanotechnol.*, **5**, 116–120, 2010.

[207] E. Bianchi, P. Tartaglia, E. La Nave and F. Sciortino, 'Fully Solvable Equilibrium Self-Assembly Process: Fine-Tuning the Clusters Size and the Connectivity in Patchy Particle Systems', *J. Phys. Chem. B*, **111**, 11765–11769, 2007.

[208] S. C. Glotzer, 'Some Assembly Required', *Science*, **306**, 419–420, 2004.

[209] T. E. Ouldridge, I. G. Johnston, A. A. Louis and J. P. K. Doye, 'The Self-Assembly of DNA Holliday Junctions Studied with a Minimal Model', *J. Chem. Phys.*, **130**, 065101, 2009.

[210] T. E. Ouldridge, A. A. Louis and J. P. K. Doye, 'DNA Nanotweezers Studied with a Coarse-Grained Model of DNA', *Phys. Rev. Lett.*, **104**, 178101, 2010.

[211] T. Foster, 'Immune Evasion by Staphylococci', *Nat. Rev. Microbiol.*, **3**, 948–958, 2005.

[212] M. Parker, 'Cryptic Clues as to how Water-Soluble Protein Toxins Form Pores in Membranes', *Toxicon*, **42**, 1–6, 2003.

[213] H. M. Berman, J. Westbrook, Z. Feng, G. Gilliland, T. N. Bhat, H. Weissig, I. N. Shindyalov and P. E. Bourne, 'The Protein Data Bank', *Nucleic Acids Res.*, **28**, 235–242, 2000.

[214] L. Song, M. R. Hobaugh, C. Shustak, S. Cheley, H. Bayley and J. E. Gouaux, 'Structure of Staphylococcal α-Hemolysin, a Heptameric Transmembrane Pore', *Science*, **274**, 1859–1865, 1996.

[215] S. Das, M. Darshi, S. Cheley, M. Wallace and H. Bayley, 'Membrane Protein Stoichiometry Determined from the Step-Wise Photobleaching of Dye-Labelled Subunits', *ChemBioChem*, **8**, 994–999, 2007.

[216] T. Kawate and E. Gouaux, 'Arresting and Releasing *Staphylococcal* α-Hemolysin at Intermediate Stages of Pore Formation by Engineered Disulfide Bonds', *Protein Sci.*, **12**, 997–1006, 2003.

[217] L. Jayasinghe, G. Miles and H. Bayley, 'Role of the Amino Latch of Staphylococcal α-Hemolysin in Pore Formation: a Co-Operative Interaction between the N Terminus and Position 217', *J. Biol. Chem.*, **281**, 2195–2204, 2006.

[218] B. L. Kagan, A. Finkelstein and M. Colombini, ‘Diphtheria Toxin Fragment Forms Large Pores in Phospholipid Bilayer Membranes’, *Proc. Natl. Acad. Sci. USA*, **78**, 4950, 1981.

[219] F. D. Lowy, ‘Staphylococcus Aureus Infections’, *N. Engl. J. Med.*, **339**, 520, 1998.

[220] G. N. Moll, W. N. Konings and A. J. M. Driessen, ‘Bacteriocins: Mechanism of Membrane Insertion and Pore Formation’, *A. van Leeuw. J. Microb.*, **76**, 185–198, 1999.

[221] J. Delves-Broughton, P. Blackburn, R. Evans and J. Hugenholtz, ‘Applications of the Bacteriocin, Nisin’, *A. van Leeuw. J. Microb.*, **69**, 193–202, 1996.

[222] P. W. Atkins and J. de Paula. *Atkins’ Physical Chemistry.* Oxford University Press, Oxford, 7th edition, 2001. ISBN 9780198792857.

[223] S. Tilley, E. Orlova, R. Gilbert, P. Andrew and H. Saibil, ‘Structural Basis of Pore Formation by the Bacterial Toxin Pneumolysin’, *Cell*, **121**, 247–256, 2005.

[224] L. A. Davies, G. Jackson and L. F. Rull, ‘Simulation Study of the Link between Molecular Association and Reentrant Miscibility for a Mixture of Molecules with Directional Interactions’, *Phys. Rev. Lett.*, **82**, 5285–5288, 1999.

[225] G. Jackson, ‘Theory of Closed-Loop Liquid-Liquid Immiscibility in Mixtures of Molecules with Directional Attractive Forces’, *Mol. Phys.*, **72**, 1365–1385, 1991.

[226] F. F. Nord, M. Bier and S. N. Timasheff, ‘Investigations on Proteins and Polymers. IV. Critical Phenomena in Polyvinyl Alcohol-Acetate Copolymer Solutions’, *J. Am. Chem. Soc.*, **73**, 289–293, 1951.

[227] M. Corti, C. Minero and V. Degiorgio, ‘Cloud Point Transition in Nonionic Micellar Solutions’, *J. Phys. Chem.*, **88**, 309–317, 1984.

[228] P. E. Cladis, ‘A One Hundred Year Perspective of the Reentrant Nematic Phase’, *Mol. Cryst. Liq. Cryst.*, **165**, 85–121, 1988.

[229] Y. Levin and M. E. Fisher, ‘Criticality in the Hard-Sphere Ionic Fluid’, *Physica A*, **225**, 164–220, 1996.

[230] M. E. Fisher and Y. Levin, ‘Criticality in Ionic Fluids: Debye-Hückel Theory, Bjerrum, and Beyond’, *Phys. Rev. Lett.*, **71**, 3826–3829, 1993.

[231] T. L. Beck, J. Jellinek and R. S. Berry, ‘Rare Gas Clusters: Solids, Liquids, Slush, and Magic Numbers’, *J. Chem. Phys.*, **87**, 545, 1987.

[232] T. E. Ouldridge, A. A. Louis and J. P. K. Doye, 'Extracting Bulk Properties of Self-Assembling Systems from Small Simulations', *J. Phys.: Condens. Matter*, **22**, 104102, 2010.

[233] R. Laghaei, E. N. Nasrabad and B. C. Eu, 'Excluded Volume in the Generic van der Waals Equation of State and the Self-Diffusion Coefficient of the Lennard-Jones Fluid', *J. Chem. Phys.*, **124**, 154502, 2006.

[234] P. H. V. Konynenburg and R. L. Scott, 'Critical Lines and Phase Equilibria in Binary van der Waals Mixtures', *Philos. T. R. Soc. A*, **298**, 495–540, 1980.

[235] S. Torquato, T. M. Truskett and P. G. Debenedetti, 'Is Random Close Packing of Spheres Well Defined?', *Phys. Rev. Lett.*, **84**, 2064–2067, 2000.

[236] D. Zwillinger, editor. *CRC Standard Mathematical Tables and Formulae.* CRC Press, Boca Raton, London, 31st edition, 2003. ISBN 1584882913.

[237] R. P. Hermann, F. Grandjean and G. J. Long, 'Einstein Oscillators that Impede Thermal Transport', *Am. J. Phys.*, **73**, 110–118, 2005.

[238] M. Hillert. *Phase Equilibria, Phase Diagrams and Phase Transformations: Their Thermodynamic Basis.* Cambridge University Press, New York, 2nd edition, 2008. ISBN 9780521853514.

[239] A. Giacometti, F. Lado, J. Largo, G. Pastore and F. Sciortino, 'Effects of Patch Size and Number Within a Simple Model of Patchy Colloids', *J. Chem. Phys.*, **132**, 174110, 2010.

[240] E. Bianchi, J. Largo, P. Tartaglia, E. Zaccarelli and F. Sciortino, 'Phase Diagram of Patchy Colloids: Towards Empty Liquids', *Phys. Rev. Lett.*, **97**, 168301, 2006.

[241] S. Torquato and Y. Jiao, 'Dense Packings of the Platonic and Archimedean Solids', *Nature*, **460**, 876–879, 2009.

[242] J. P. K. Doye and W. C. K. Poon, 'Protein Crystallization *in Vivo*', *Curr. Opin. Colloid In.*, **11**, 40–46, 2006.

[243] V. J. Anderson and H. N. W. Lekkerkerker, 'Insights into Phase Transition Kinetics from Colloid Science', *Nature*, **416**, 811–815, 2002.

[244] A. R. West. *Basic Solid State Chemistry.* John Wiley & Sons Ltd, Chichester, 1999. ISBN 0471987557.

[245] Y. He, Y. Tian, Y. Chen, Z. Deng, A. E. Ribbe and C. Mao, 'Sequence Symmetry as a Tool for Designing DNA Nanostructures', *Angew. Chem. Int. Ed.*, **117**, 6852–6854, 2005.

[246] Y. He and C. Mao, 'Balancing flexibility and stress in DNA nanostructures', *Chem. Commun.*, 968–969, 2006.

[247] C. Zhang, M. Su, Y. He, X. Zhao, P.-A. Fang, A. E. Ribbe, W. Jiang and C. Mao, 'Conformational Flexibility Facilitates Self-Assembly of Complex DNA Nanostructures', *Proc. Natl. Acad. Sci. USA*, **105**, 10665–10669, 2008.

[248] N. Goodnight, R. Wang and G. Humphreys, 'Computation on Programmable Graphics Hardware', *IEEE Comput. Graph.*, 12–15, 2005.

[249] D. Goldberg, 'What Every Computer Scientist Should Know About Floating-Point Arithmetic', *ACM Comput. Surv.*, **23**, 5–48, 1991.

[250] nVidia. 'Cuda Website'. Accessed 11-1-2011, URL `http://www.nvidia.com/object/cuda_home_new.html`.

[251] J. E. Stone, D. Gohara and G. Shi, 'OpenCL: A Parallel Programming Standard for Heterogeneous Computing Systems', *Comput. Sci. Eng.*, **12**, 66, 2010.

[252] H. Benson. *University Physics, Revised Edition.* John Wiley & Sons–Interscience, New Jersey, 1996. ISBN 0471006890.

[253] W. Humphrey, A. Dalke and K. Schulten, 'VMD – Visual Molecular Dynamics', *J. Mol. Graph.*, **14**, 33–38, 1996.

[254] POV Raytracer Pty. Ltd. 'POV-Ray website'. Accessed 7-2-2011, URL `http://www.povray.org`.

Appendices

The following appendices contain more technical information relevant to the project that may be useful to the more specialist reader, along with some results that are interesting, but not as directly relevant to our self-assembling systems as those in the main section.

Appendix A

Clusters simulated

A number different shaped clusters were simulated during this project, and this appendix gives details about all of them. In the pictures, the colour scheme is for the first particle type to be shown as red and the locations of its patches are shown in blue (type-A), green (type-B), yellow (type-C), cyan (type-D) and magenta (type-E). The second particle type (if present) is shown in yellow with orange patches (if present). The patches do not have volume, but rather they are schematically represented as such.

DS - Dimer of dimers

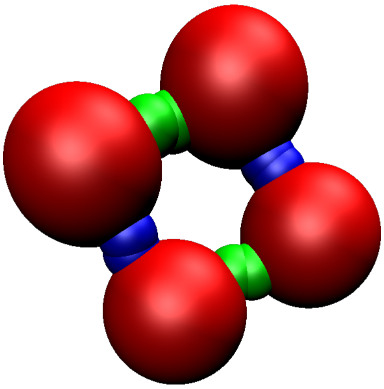

Particles:	4 particles, all of the same type
Patches:	2 patches per monomer, A and B
Bonds:	4 bonds, of two types A-A and B-B
Comments:	Wang-Landau test system

IhO - Uncentred icosahedron

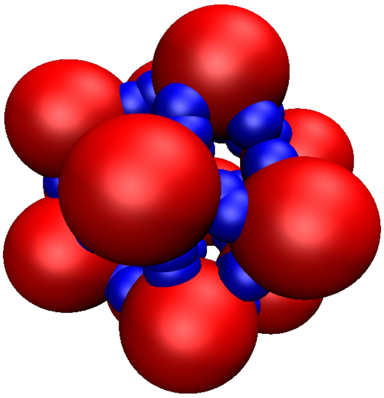

Particles:	12 particles, all of the same type
Patches:	5 patches per monomer, all of the same (type-A)
Bonds:	30 bonds, all type A-A
Comments:	Simple icosahedron

IhC - Centred icosahedron

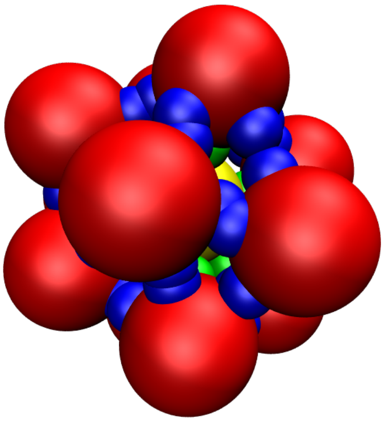

Particles: 13 particles, 12 shell and 1 core

Patches: 6 patches per shell monomer (5 type-A, 1 type-B), 1 patch per core particle (type-C)

Bonds: 42 bonds, 30 A-A, 12 B-C

Comments: Encapsulation test sytem

IhCM - Multiply-patched core, centred icosahedron

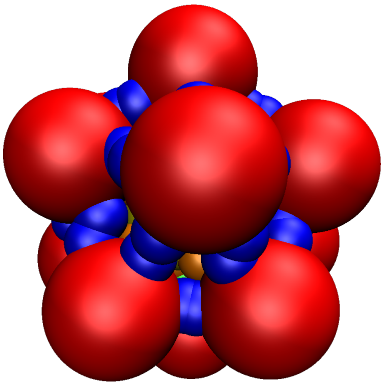

Particles: 13 particles, 12 shell and 1 core

Patches: 6 patches per shell monomer (5 type-A, 1 type-B), 12 patches per core particle (type-C)

Bonds: 42 bonds, 30 A-A, 12 B-C

Comments: Encapsulation test system

DdO - Uncentred dodecahedron

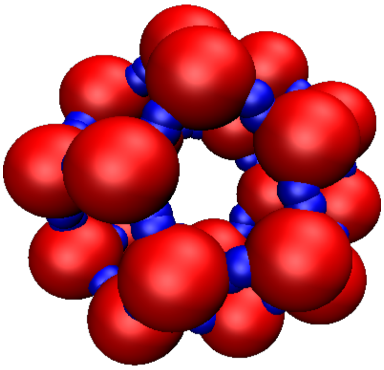

Particles: 20 particles, all of the same time

Patches: 3 patches per monomer, all of the same type (type-A)

Bonds: 30 bonds, all type A-A

Comments: Kinetically inaccessible target

DdC - Centred dodecahedron

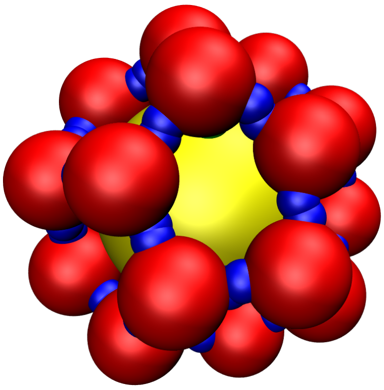

Particles: 21 particles, 20 shell and 1 core

Patches: 4 patches per shell monomer (3 type-A, 1 type-B), 1 patch per core monomer

Bonds: 50 bonds, 30 type A-A, 20 type B-C

Comments: Encapsulation test system

DdCM - Multiply-patched core, centred dodecahedron

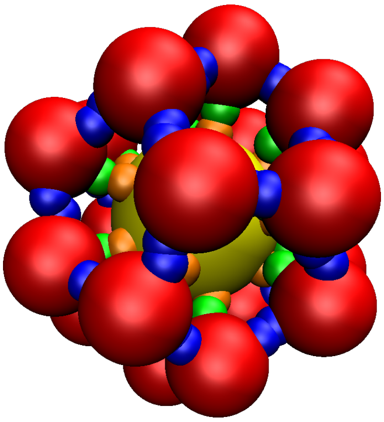

Particles: 21 particles, 20 shell and 1 core

Patches: 4 patches per shell monomer (3 type-A, 1 type-B), 20 patches per core monomer

Bonds: 50 bonds, 30 type A-A, 20 type B-C

Comments: Encapsulation test system

HD - Dimer

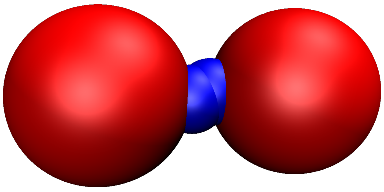

Particles: 2 particles, all of the same type

Patches: 1 patches (type-A)

Bonds: 1 bond, of type A-A

Comments: Hierarchical intermediate (also possible with different bond types)

HT - Triangle

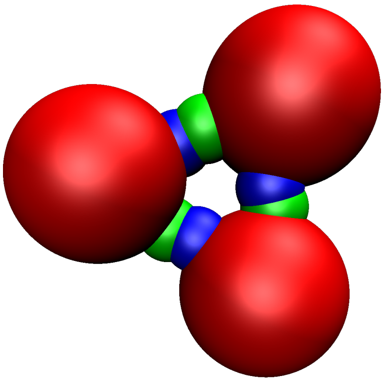

Particles: 3 particles, all of the same type

Patches: 2 patches per monomer, A and B

Bonds: 3 bonds, all of type A-B

Comments: Assembly-competent hierarchical intermediate

HS - Square

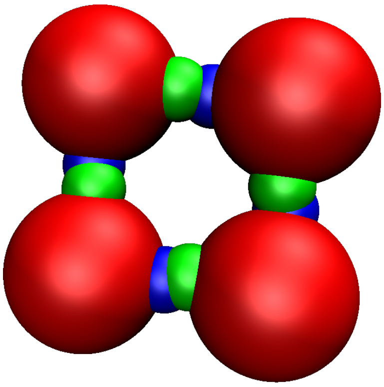

Particles: 4 particles, all of the same type

Patches: 2 patches per monomer, A and B

Bonds: 4 bonds, all of type A-B

Comments: Assembly-competent hierarchical intermediate

IWC - Assembly-competent hierarchically forming icosahedron

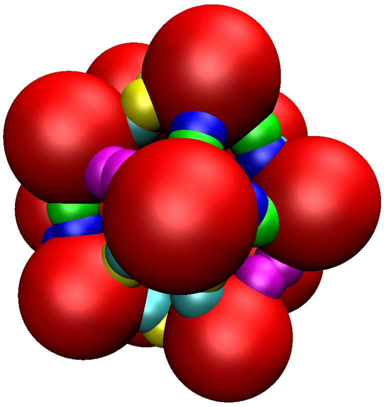

Particles:	12 particles, all of the same type
Patches:	5 patches per monomer, 1 each of types A to E
Bonds:	30 bonds, 12 of type A-B, 12 of type C-D, 6 of type E-E
Comments:	Hierarchical assembly target

IW - Non-assembly competent hierarchically forming icosahedron

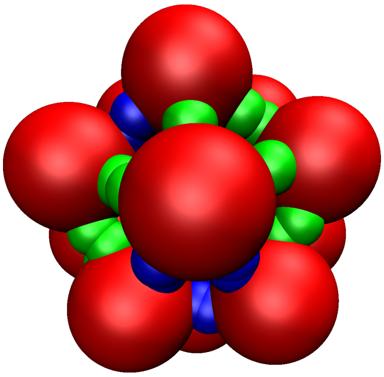

Particles:	12 particles, all of the same type
Patches:	5 patches per monomer, 2 type-A and 3 type-B
Bonds:	30 bonds, 12 of type A-A, 18 of type B-B
Comments:	Hierarchical assembly target

cO - Cuboctahedron

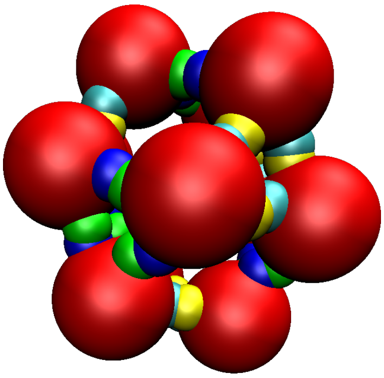

Particles:	12 particles, all of the same type
Patches:	4 patches per monomer, of types A to D
Bonds:	24 bonds, 12 A-B and 12 C-D
Comments:	Hierarchical assembly target

tT - Truncated tetrahedron

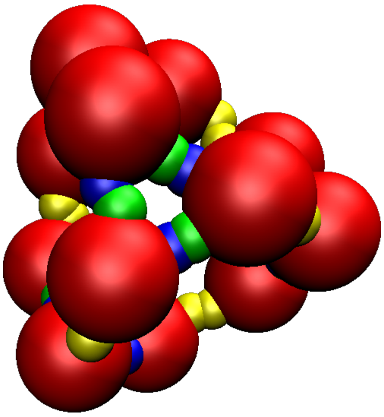

Particles: 12 particles, all of the same type

Patches: 3 patches per monomer, of types A to C

Bonds: 18 bonds, 12 A-B and 6 C-C

Comments: Hierarchical assembly target

tO - Truncated octahedron

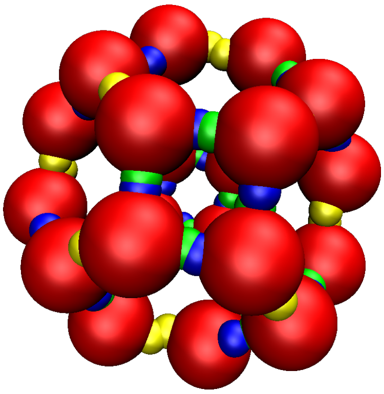

Particles: 24 particles, all of the same type

Patches: 3 patches per monomer, of types A to C

Bonds: 36 bonds, 24 A-B and 12 C-C

Comments: Hierarchical assembly target

rC - Rhombicuboctahedron

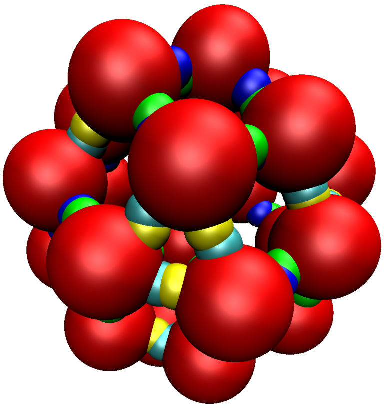

Particles: 24 particles, all of the same type

Patches: 4 patches per monomer, of types A to D

Bonds: 48 bonds, 24 A-B and 24 C-D

Comments: Hierarchical assembly target

sC - Snub cube

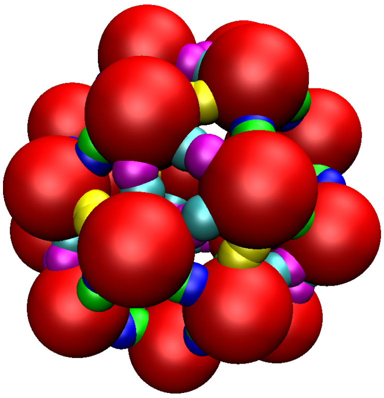

Particles: 24 particles, all of the same type

Patches: 5 patches per monomer, of types A to E

Bonds: 60 bonds, 24 A-B, 12 C-C and 24 D-E

Comments: Hierarchical assembly target

P7 - Heptameric Pore

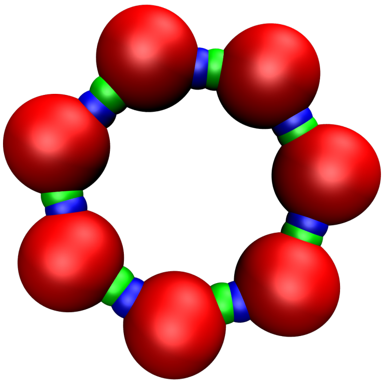

Particles: 7 particles, all of the same type

Patches: 2 patches per monomer, A and B

Bonds: 7 bonds, all type A-B

Comments: Model of a protein complex

Oh - Octahedron

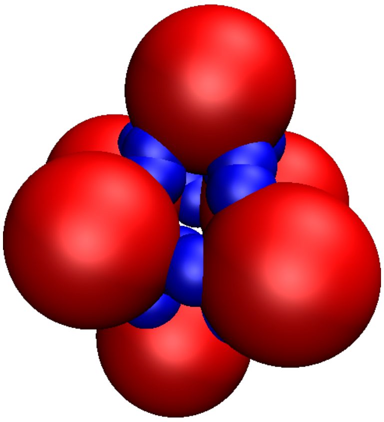

Particles: 6 particles, all of the same type

Patches: 4 patches per monomer, all type A

Bonds: 12 bonds all of type A-A

Comments: Forms complicated liquid-crystal-like structures

Appendix B

Other areas of interest

This appendix contains the results of several small sub-projects, each of which is rather self-contained, but nevertheless interesting. The projects include the effect of changing several parameters in the simulation setup, and show that our results are general and thus our conclusions are applicable to self-assembling patchy systems in general, and not just our model. We also look at the use of using graphics processors (using CUDA) to run our simulations rather than central processing units (using Fortran).

B.1 Changing the move type

Past research in the group [46] has looked into the effect of cluster moves [146] (see section 2.2.3) rather than single-particle moves. However, the moves were rotations or random displacements in a cube: this biases the system to make larger moves (to make a small move requires three random numbers to be small). Our move types are explained in section 2.2.1. A small test into the effect of using moves randomly distributed in a cube was carried out (all other simulations in this thesis use moves randomly distributed in a sphere). All thermodynamic properties of the system would not be expected to change when using a different move type (as detailed balance is followed by the program). However, the dynamics may be affected. We consider this effect here. Both methods of generating a translational move used the same maximum move size of $0.3\,\sigma^{\mathrm{LJ}}$ (in each dimension for the cubic moves, thus the maximum move size is $\approx 0.5\,\sigma^{\mathrm{LJ}}$, in total for the spherically symmetric moves), thus the mean size of move was greater for the cubic moves.

B.1.1 Centred icosahedra system

We tested the formation of the centred icosahedral cluster-forming model, described in chapter 3, using the two different methods of selecting translational moves with the virtual

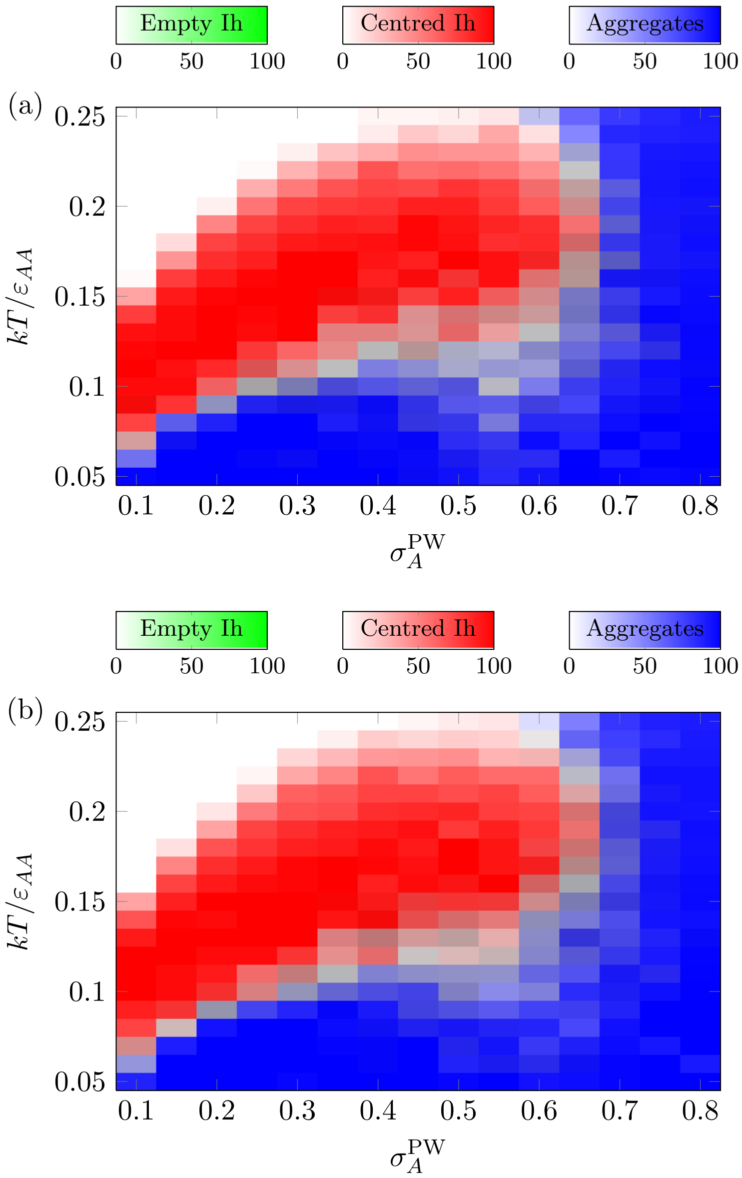

Figure B.1: The yield of simulations in the centred icosahedron forming system as a function of patch width of the shell-shell interactions and temperature, using moves distributed evenly in (a) a sphere and (b) a cube. Other settings are the same as in table 3.1.

move Monte Carlo algorithm of Geissler and Whitelam [146]. We considered a variety of different patch widths examining both the time resolved and final yields. The yield plots are shown in figure B.1. They clearly show that there is no major difference in dynamic yield between the two different move types. The slight differences in the aggregates that can be seen at low temperatures and patch widths of around 0.55 radians can be explained in terms of the size of the maximum move. It is easier for smaller clusters to aggregate when longer cluster moves are possible, this is especially important when they are initially forming.

The effect on the actual dynamics can be seen in our time resolved plots in figure B.2.

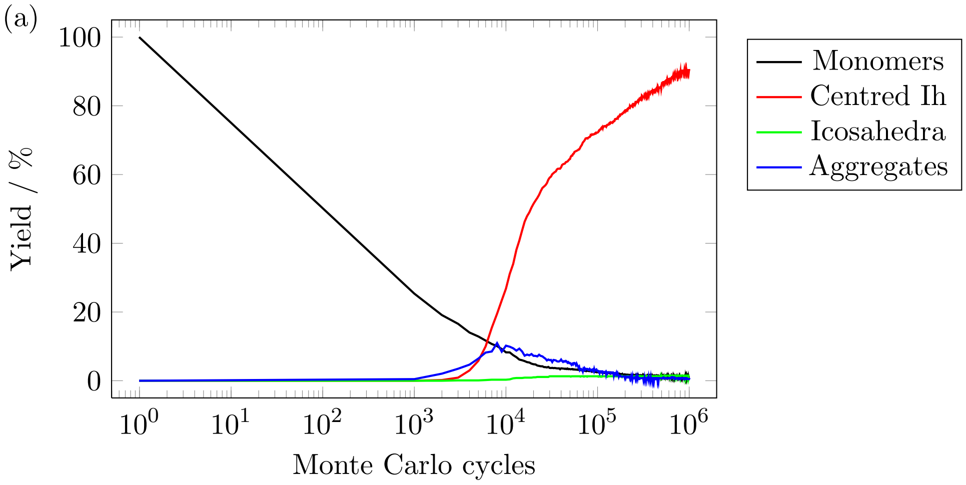

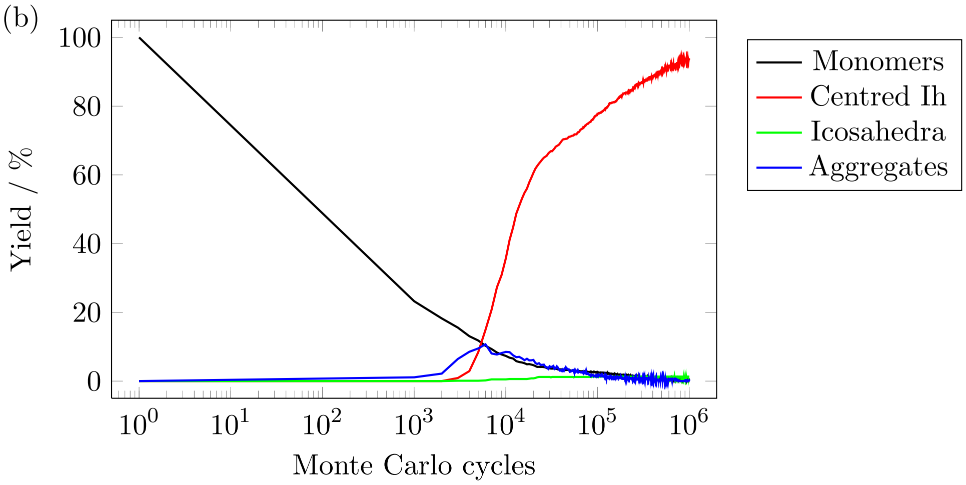

Figure B.2: The yield as a function of time for $kT = 0.14\,\varepsilon$ for the centred icosahedron forming system using moves distributed evenly in (a) a sphere and (b) a cube. Other settings are the same as in table 3.1.

The plots show a slight change in the dynamics, with the cube moves taking longer to form the target cluster. The time scale of aggregate formation is, however, similar for both cube and spherically-distributed moves.

B.2 Changing the range of interaction

Some initial previous work [46] suggested that changing the range of the interactions in a simple patchy model can have a profound effect on the behaviour of a simple system. Furthermore, the range of the interaction potential can generally be tuned in colloidal systems. In this section we look more into this effect on the uncentred and centred icosahedron-forming systems. The range of interactions is a delicate balancing act: although

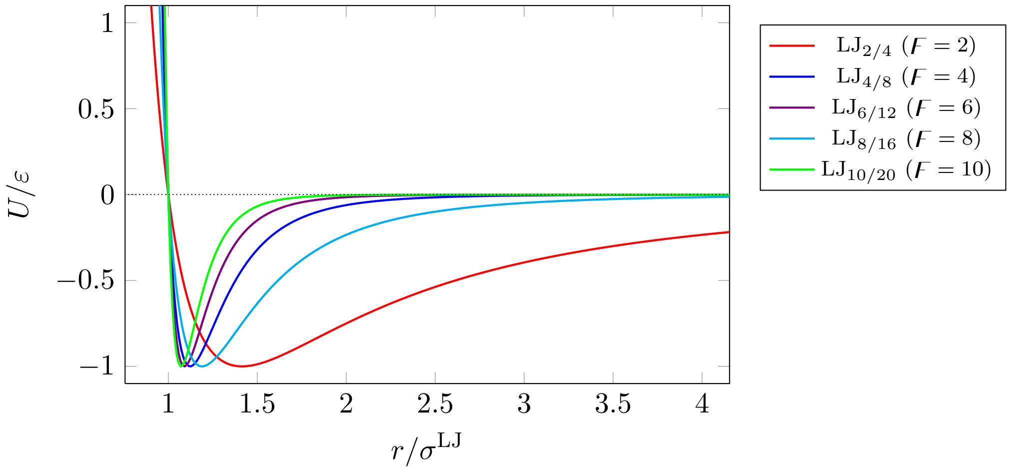

Figure B.3: Example interaction energy curves for different powers of the Lennard-Jones-type potential, note that $\digamma = 6$ corresponds to Lennard-Jones.

long range interactions stabilise all structures, they do so for aggregates, glasses and other high density structures more strongly than they do for a gas of clusters.

Figure B.3 shows the different interaction energy curves for different ranges of interaction. In the case of the longer ranged interactions a cutoff of $4.5\,\sigma$ was used during the simulations. The potentials are named according to the value of the exponents of the attractive and repulsive parts (equation B.1). In this section we just consider the effect of changing one parameter $\digamma$ such that the interactions $\mathrm{LJ}_{\digamma/2\digamma}$ are defined as:

$$\mathrm{LJ}_{\digamma/2\digamma}(\boldsymbol{r}_{ij}) = 4\varepsilon\left[\left(\frac{\sigma^{\mathrm{LJ}}}{|\boldsymbol{r}_{ij}|}\right)^{2\digamma} - \left(\frac{\sigma^{\mathrm{LJ}}}{|\boldsymbol{r}_{ij}|}\right)^{\digamma}\right] \tag{B.1}$$

Figures B.4 and B.6 show the yield for a range of range of interactions. The interaction has only a small effect for values of $\digamma$ greater than four in both the case of the centred and the uncentred system. However, at lower values of $\digamma$ (longer range interactions), the behaviour of the two systems differs slightly.

Uncentred system

The uncentred system shows aggregation occurs for all temperatures tested at low values of $\digamma$. While at values of 4 and above formation of uncentred icosahedron seems to be the stable structure. The aggregate that is formed at low $\digamma$ values can be seen in figure B.5(a) and (b), and it is fundamentally different in structure from the open kinetic aggregates that we normally see in our simulations (figure B.5(c)). It is also interesting that, in the regime of correct particle formation, the kinetic aggregation temperature does not significantly change with increasing $\digamma$. The formation temperature of the target clusters, on the other

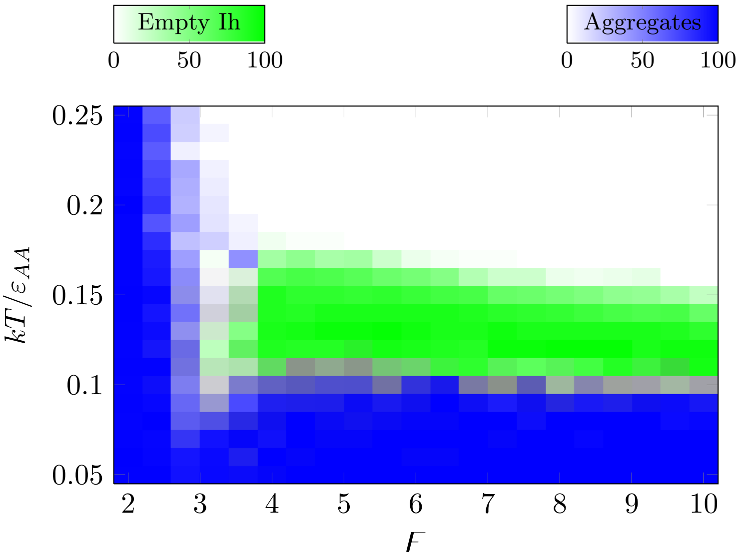

Figure B.4: Yield plot as a function of temperature and the exponent $\digamma$ in the uncentred system. The settings used are the same as those for the IhO system in table 3.1.

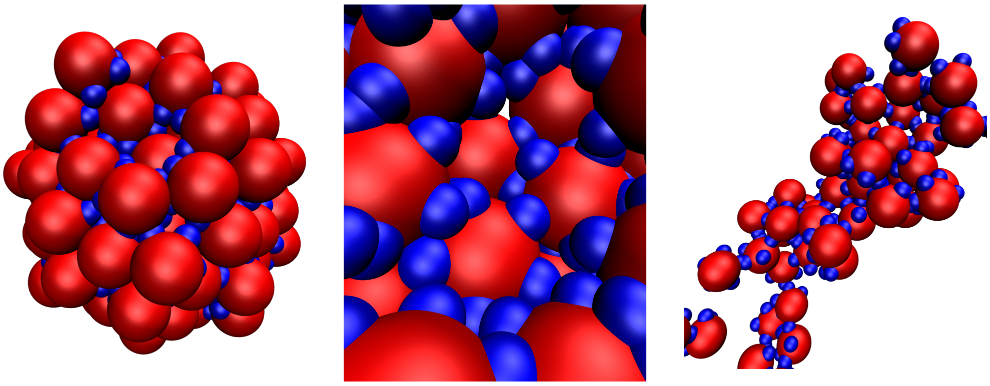

(a) (b) (c)

Figure B.5: The aggregate found at low values of $\digamma$, as viewed from the (a) outside and (b) the centre. The configuration was taken from a run at a temperature of $0.15\varepsilon/k$ with $\digamma = 2$. (c) Shows the aggregate found at higher values of $\digamma$. Note the more open structure compared with those formed for longer range interactions. The configuration is taken from a run at a temperature of $0.05\varepsilon/k$ with $\digamma = 6$.

hand, does decrease slightly as we shorten the range of the interactions. This is most likely due to the decrease in entropy of the formed clusters.

All particles in the aggregate formed at long range interactions (figure B.5 (a) and (b)) are positioned with the patches facing towards the centre of the aggregate. They build in layers with patches pointing towards other patches, but the distance between particles is larger than is normally seen in the Lennard-Jones-range ($\digamma = 6$) potential, as the distance at which the minimum value of the potential occurs is longer.

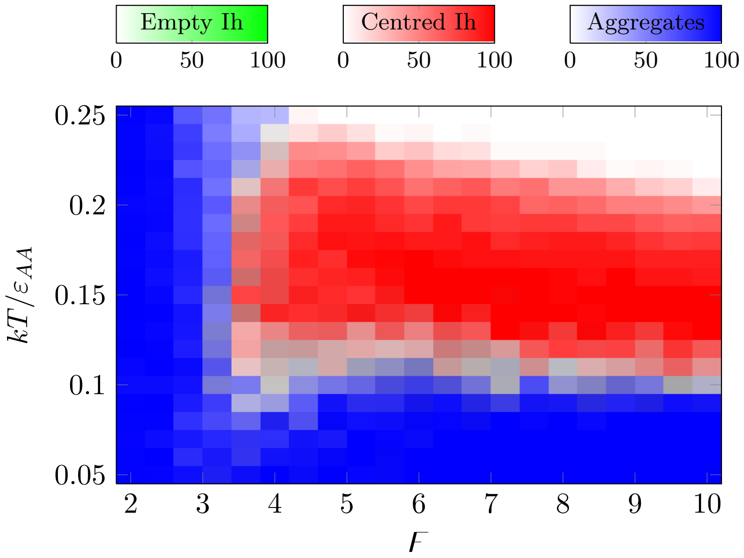

Figure B.6: Yield plot as a function of temperature and the exponent F in the uncentred system. The settings used are the same as those for the IhC system in table 3.1.

Centred system

Similarly, in the centred system, long range interactions (low F regime) encourage aggregate formation. The effect of the longer range interactions outside of this aggregate forming region is to increase the transition temperature. Again the aggregates formed at low temperature are different for the short range and long range interactions (figure B.7(b) and (c)). With longer range interactions we get an increase in the number of large clusters with multiple core particles (figures B.7(a) and (b)), as opposed to the more open aggregate structure formed with shorter range interactions that involve chains and partially complete icosahedra (figure B.7(c)). There are a number of pentagonal faces present on the aggregates formed for longer range-interactions (figure B.7(a) and (b)), while only triangular faces are seen in the aggregates formed for shorter range interactions (figure B.7(c)) It is interesting that our low temperature kinetic aggregates (figure B.7(b)) for longer range interactions are generally smaller than the aggregates formed at higher temperature (where they are more liquid like and form around all available core particles, figure B.7(a)). Unlike the uncentred system the shell particles mainly form a monolayer around the core particles.

B.3 Effects of system size on simulations

Other research within our group has looked at the possible effects of attempting to extract bulk properties from small systems [232]. For our self-assembling clustering system the most relevant of these is the effect of the size of the box on our density of states. If we consider the differences in states for a choice of boxes, one which contains a maximum of one cluster

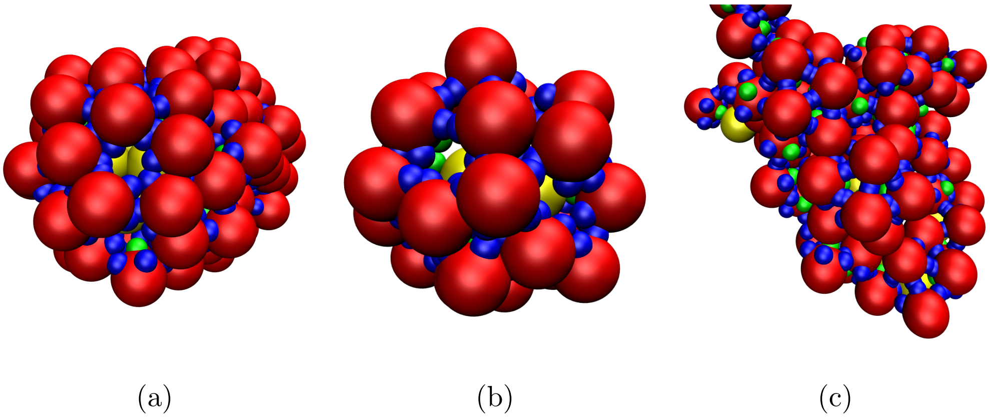

(a) (b) (c)

Figure B.7: The aggregate found at F of (a) and (b) 2, (c) 6. Taken from a runs at temperatures of (a) $0.14\varepsilon/k$, (b) and (c) $0.05\varepsilon/k$.

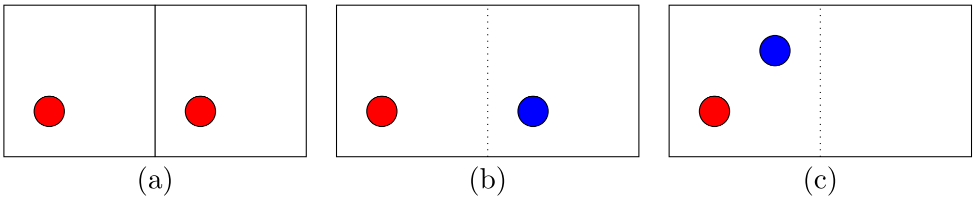

(a) (b) (c)

Figure B.8: Schematic of the issues with a small box size. (a) Small box with periodic boundary shown for one repeat, (b) Larger box with two particles in an example setup that would be included in our periodic smaller box, (c) Larger box in a configuration that is not included when we use a smaller box.

(box A), and another with the same density of particles that can contain a maximum of two clusters (box B). We generally assume that the states which we sample with box A follow the behaviour of the larger system, especially thanks to the periodic boundary condition (figure B.8(a) and (b)). However, this assumes that there is no asymmetry in the overall configurations of our larger box. For example, if states where both particles in our box B are within the same half of the box (figure B.8(c)) are important to our thermodynamics and kinetics, then we do not get good sampling by using a small box with a periodic boundary condition. This section deals with the impact of this system size effect on dynamics and thermodynamics in our simulations.

B.3.1 The Wang-Landau algorithm and the effect of system size

We explored the thermodynamic consequences of this effect, first using the canonical ensemble Wang-Landau algorithm. We considered the centred icosahedral system with a number of different numbers of particles, but with the same number density ($N/V = 0.15\,(\sigma^{\mathrm{LJ}})^{-3}$); all other simulations settings are given in table 3.1. We used 200 bins for every $29\,\varepsilon$ covered, and set the minimum energy to $-29\,\varepsilon$ per possible completed cluster.

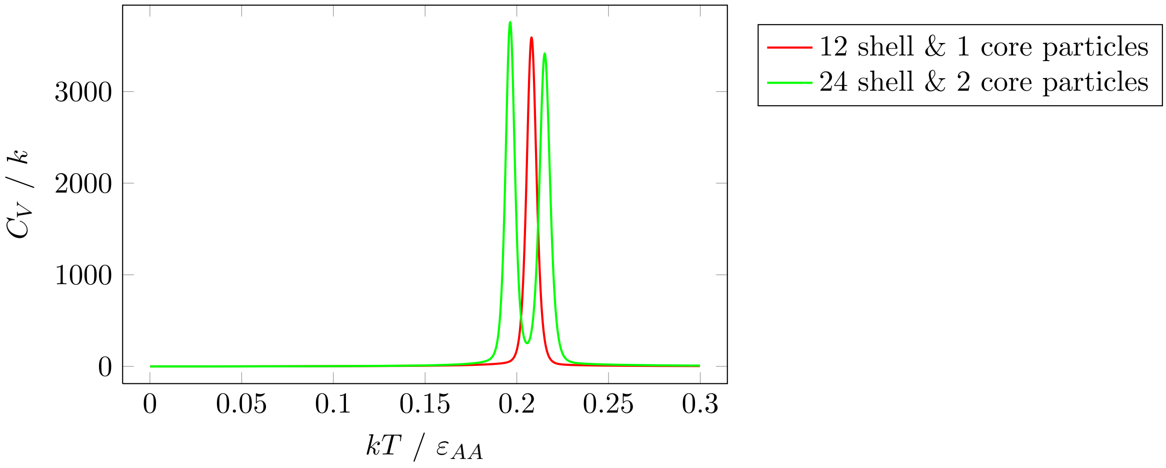

Figure B.9: Changes in the heat capacity for a maximum of one or two completed clusters in the centred icosahedral system showing a clear splitting of the peak into two when two clusters are present.

This ensured that the resolution of the bins was the same for all runs. The results for the algorithm are rather interesting. When we increase the system so that it may have two completed clusters, rather than the usual one, we see a splitting of the peak. These two peaks correspond to transitions between boxes with zero and one completed cluster, and one and two completed clusters (figure B.9). These transitions are still centred around the transition temperature from the single cluster, and have a similar integral under each peak (as the transitions have a similar change in energy). When we further increase the system size, we find that larger boxes, which have a higher maximum number of clusters, have more peaks, but these tend to overlap more and more as the system size increases before giving a broad peak (made up of many small peaks overlapping), as is shown in figure B.10. The centre of the peaks remain in approximately the same position throughout. Therefore, our use of a single cluster gives approximately the same transition temperature, but without the need to simulate hundreds of particles required to see a single broad peak, and we are able to converge results in the shortest possible time (the time for convergence scales approximately as the square of the number of clusters, owing to the increase in the number of bins and configurations which need to be sampled). Multiple peaks corresponding to a single transition are seen in the flipping system we examine in section B.6.1.

Our Wang-Landau runs do not correctly take the behaviour of aggregation into account (this is normally a situation where most particles end up in a single large cluster on one side of the box, as in figure B.8(c)). However, Wang-Landau runs of larger sizes of system which attempted to simulate regions where aggregation occurs kinetically were unable to converge within six months.

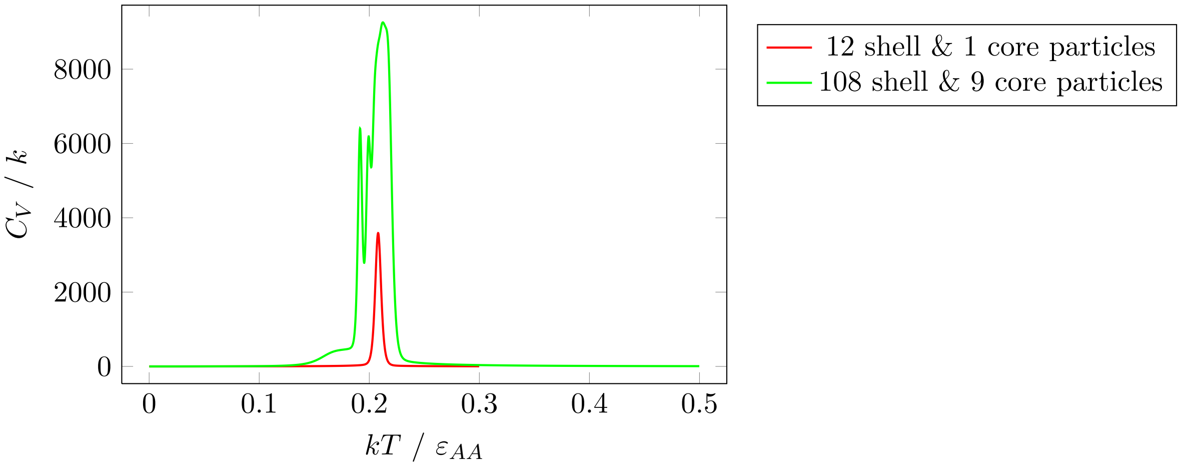

Figure B.10: The heat capacity for a maximum of one or nine completed clusters in the centred icosahedral system showing the broadening of the peak and the overlap of the smaller peaks as more clusters are present.

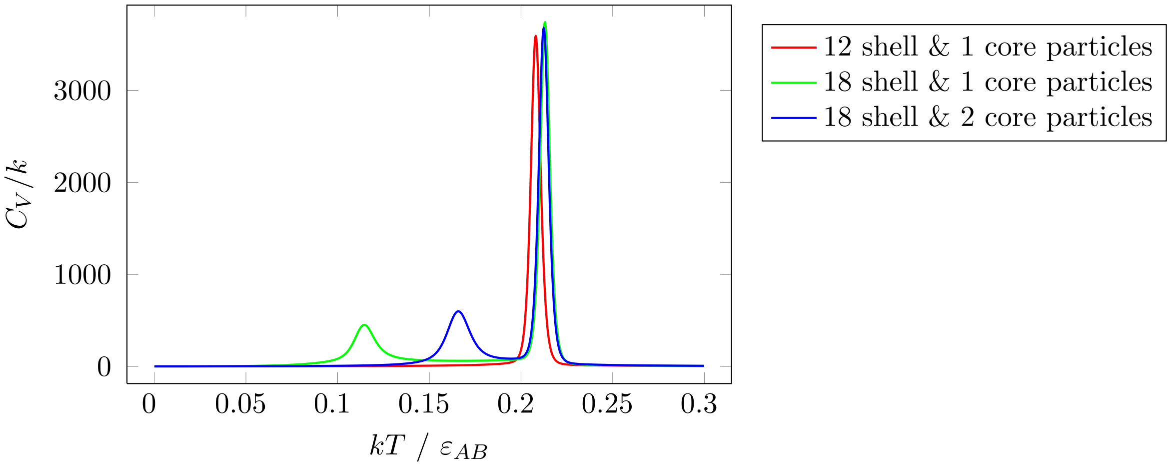

Figure B.11: Heat capacity for non-integer numbers of clusters.

Incomplete clusters

We also investigated the heat capacity predicted by the Wang-Landau algorithm when the number of particles is not an integer multiple of the cluster size. We examined the centred icosahedral system, this time with 18 shell particles, and either one or two central particles. The results (figure B.11) show that we slightly shift the transition temperature of our completed cluster, but that we also add a second, smaller, peak. This peak corresponds to the part formation of an icosahedral cluster. Small peaks like this are seen in the high resolution heat capacity of the centred dodecahedral system in section B.5.1.

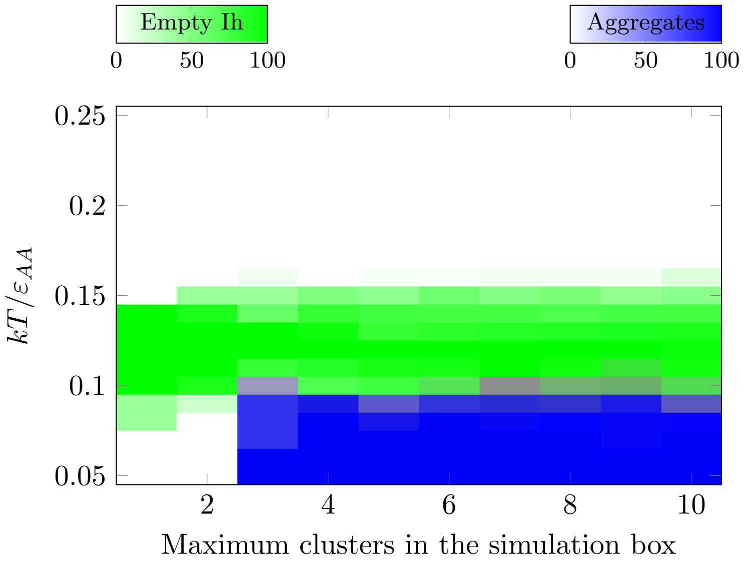

Figure B.12: Yields of icosahedra and aggregates in the uncentred system as a function of the maximum number of clusters possible.

B.3.2 Size effects on the grand canonical Wang-Landau algorithm

As mentioned in section 6.3.3, in a Grand Canonical simulation with clustering there is an effect of the system size. If we have a continuous transition between monomers and clusters, then, for a non-infinite system below the clustering temperature, the densities which correspond to an integer number of clusters will appear more stable than those which do not. Preliminary results show a number of islands of stability which change density if the box size is changed. It is possible to thus uncover which phase changes are continuous and which are discontinuous, by rerunning the box at a different density. This doubles the computation time required compared to a single run, and is very sensitive to any noise.

B.3.3 Effects of system size on dynamic yields

We examined the effect of the system size on the dynamics of formation for the uncentred icosahedral system. While keeping the number density constant at $0.15\,(\sigma^{\mathrm{LJ}})^{-3}$, we ran a number of repeats for simulation boxes which had a different number of particles in. We always kept the total number of particles divisible exactly by 12, so that a yield of 100% was possible in all cases. The results can be seen in figure B.12. This clearly shows that there is not a strong effect of the system size on the dynamic yields, though for systems with only one or two possible clusters we see no “aggregates” (as our definition of aggregate is for a cluster to have more than 25 particles, this is unsurprising).

B.4 The effect of changing the cutoff distance on the kinetics and thermodynamics of the system

In our simulations, we generally cut off the potential at $3.5\,\sigma^{\mathrm{LJ}}$ and shifted the potential up so that there was no interaction at this distance (see section 2.2.1). However, it was necessary with certain systems to use a smaller cutoff distance (the cutoff distance must be less than or equal to half the box length, so if we use a small number of particles, but keep the density constant, we need to decrease our cutoff distance). The cutoff distance being small was especially an issue when we were interested in the thermodynamics of a single cluster (for example, in umbrella sampling or the Wang-Landau algorithm), or were using volume moves (in, for example, the Gibbs ensemble).

Some work into the effect of changing the cutoff of a potential on Lennard Jones particles has been carried out [179], and the effect of the cutoff has been shown to be noticeable in mean field treatment of this system below $3\,\sigma^{\mathrm{LJ}}$. We can make a few approximations to estimate the effect of the cutoff with and without patches, and then run some simulations to test how this affects the results.

Table B.1 shows the value of the energy at the cutoff distance at a range of cutoff distances, which is equal to the change compared to the non-cutoff case of the minimum value of the interaction that occurs in the cutoff and shift potential (see section 2.2.1). The changes in energies are small in all cases. However, changing the cutoff distance also means that the attractions between all particles outside of this range are no longer included in our energy calculations. We can estimate this effect by assuming the number density at distances beyond the cutoff distance is constant and equal to that of the bulk ($\rho_{\mathrm{bulk}} = 0.15/(\sigma^{\mathrm{LJ}})^3$ in our simulations). We then integrate from the cutoff to infinite distance (the third column in table B.1). We find the approximate energies lost due to our cutoff choice. Provided the densities of the particles outside the cutoff are approximately constant, then we would just expect the changing cutoff to scale the temperature slightly.

In clustering systems, the density outside of the cutoff is not necessarily simply that of the bulk, because we may form a cluster. For the icosahedron, at its minimum energy conformation, the distance between nearest neighbours $1.12\,\sigma^{\mathrm{LJ}}$, between next-nearest neighbours is $1.82\,\sigma^{\mathrm{LJ}}$, and between opposite vertices is $2.14\,\sigma^{\mathrm{LJ}}$. For a system that has formed clusters, the radial distribution function clearly shows that there is a density considerably different from the bulk at these distances (figure B.13). This leads us to expect that a system with a cutoff distance of $2\,\sigma^{\mathrm{LJ}}$ may exhibit different behaviour. However, we have assumed a Lennard-Jones like behaviour for all of our particles, without any angular modulation. When we include the fact that patches from particles on opposite

$\frac{\lvert \boldsymbol{r}^{\mathrm{cutoff}} \rvert}{\sigma^{\mathrm{LJ}}}$	$\frac{U^{\mathrm{LJ}}(\lvert \boldsymbol{r}^{\mathrm{cutoff}} \rvert)}{\varepsilon}$	$\rho_{\mathrm{bulk}} \int_{\lvert \boldsymbol{r}^{\mathrm{cutoff}} \rvert}^{\infty} \frac{U^{\mathrm{LJ}}(r)}{\varepsilon} 4\pi r^2 \mathrm{d}r$
2	−0.0615	−0.3125
2.5	−0.0163	−0.1606
3	−0.0055	−0.0930
3.5	−0.0022	−0.0586
4	−0.0010	−0.0393
4.5	−0.0005	−0.0276
5	−0.0003	−0.0201

Table B.1: Cutoff energies.

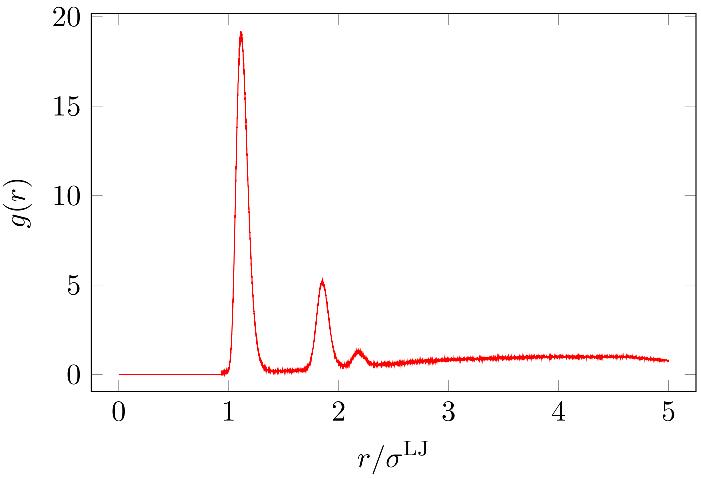

Figure B.13: Icosahedral forming region radial distribution function ($\sigma_A^{\mathrm{PW}} = 0.35$, $kT/\varepsilon_{AA} = 0.15$, other settings are given in table 3.1).

vertices are not pointing towards each other, then this effect is significantly lower. For example, for our uncentred icosahedron-forming model (IhO) using a patch width of 0.35 radians (which we used for our standard model), then the interaction energy between pairs of particles on opposite vertices is of the order of $10^{-5}\,\varepsilon$. Even for a patch width of 0.8 radians (the widest used), this effect only accounts for approximately $10^{-2}\,\varepsilon$.

We can explore what effect the cutoff distance has on our model by looking at the yields and thermodynamics of a simple system (we have chosen the uncentred icosahedron in this case).

B.4.1 Dynamic effects

Figure B.14 is a plot of the yield of uncentred icosahedra as a function of temperature and cutoff distance (other settings as in table 3.1). At values of cutoff ($\lvert \boldsymbol{r}^{\mathrm{cutoff}} \rvert$) greater than $3\,\sigma^{\mathrm{LJ}}$ there is no difference in yield. Below this level there is a slight decrease in the temperature of transition between monomers and clusters. This is the pattern we would expect from the underlying mathematics. As mentioned previously, the change in

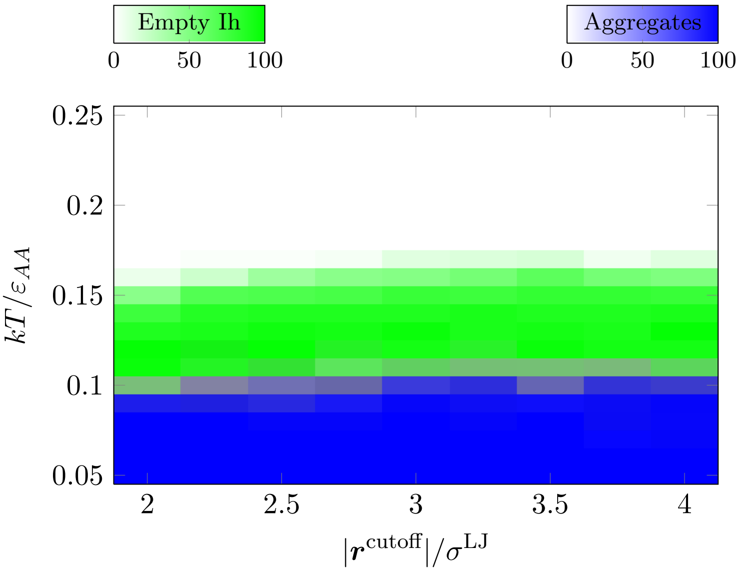

Figure B.14: Yield plot of temperature and cutoff distance. All other settings are given in table 3.1.

transition temperature, even for small cutoffs, are primarily due to the change in well depth caused by the cut and shift nature of our potential, rather than the cutoff causing a major change in behaviour. We do not expect the cutoff to change the dynamics fundamentally at this density.

B.4.2 Thermodynamic effects

In order to find the effect of the cutoff distance on the heat capacity, we used a lower number density of $0.025/(\sigma^{\text{LJ}})^3$ on an uncentred icosahedral forming system. This density allowed the simulation of a single cluster and thus our investigation into the effect of the cutoff distance was not complicated by effects due to the formation of multiple clusters (section B.3.1). For cutoff lengths above $2.75\,\sigma^{\text{LJ}}$ there is no change in the predicted transition temperature. For shorter cutoff lengths we see a slight suppression of transition temperature, similar to the results of our dynamic runs. There is no marked change in behaviour. This supports our use of a shorter cutoff for our Wang-Landau runs, as the effect is just to shift the temperature slightly.

B.5 Wang-Landau specific simulation settings

There are many variables available when running a simulation in the Wang-Landau algorithm. In this section we investigate some of these settings and compare the predictions of the algorithm and the time taken for convergence.

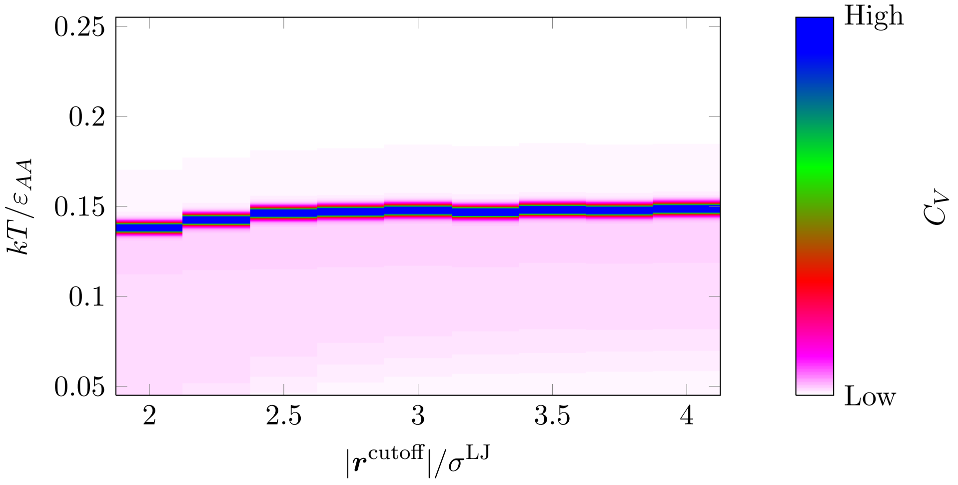

Figure B.15: Plot of heat capacity, temperature and cutoff. All other settings are given in table 3.1.

B.5.1 The effect of changing the minimum energy measured

When using the Wang-Landau algorithm it is necessary to decide on a minimum energy that will be included in the histogram estimates. This energy needs to be accessible on the time scale of the runs (so cannot be too low as the density of states decreases rapidly at low energies), and must also not exclude regions with a high density of states, otherwise we would be cutting off information relevant to our heat capacity calculations. We first examine the simple behaviour that is demonstrated by the uncentred icosahedron, and then consider the more complicated effects of the minimum energy on the centred dodecahedra system.

Uncentred icosahedra-forming system

We systematically looked at the uncentred icosahedral system, modifying the lower energy bound for the Wang-Landau algorithm run. The theoretical minimum in energy is $-30\,\varepsilon$, and in a normal dynamic run the energy gets to a value of approximately $-27\,\varepsilon$. All runs that were tested with a minimum energy below $-28\,\varepsilon$ did not converge within six months. This slow convergence is due to the low density of states associated with this region meaning that there is little probability of attempting moves into the region. The results are shown in figure B.16. They show some minor divergence, but no major changes for minimum Wang-Landau energies between $-28\,\varepsilon$ and $-22\,\varepsilon$, which is a significant range. The predicted heat capacities when higher minimum energies were used were less accurate and showed wider peaks. The time taken for convergence of the algorithm was approximately constant for all runs with minimum energies above $-26\,\varepsilon$ (the run time was limited by the number of bins, rather than the energy landscape). However, compared to the run with a minimum energy

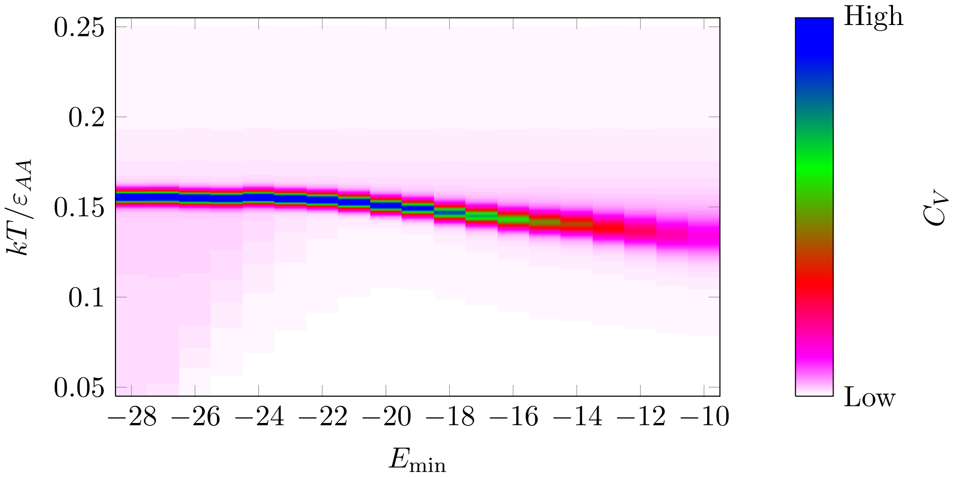

Figure B.16: The heat capacity of a system of 12 icosahedral shell particles as a function of temperature and the minimum energy used in the Wang-Landau run. All other settings are given in table 3.1.

of $-26\,\varepsilon$, it took approximately twice as long for a run with a minimum Wang-Landau energy of $-27\,\varepsilon$ to converge, and four times as long for a run with a minimum energy of $-28\,\varepsilon$ to converge. Since the prediction of the clustering temperature is unchanged by the minimum Wang-Landau energy in this region, it is possible to have significantly shorter simulations by using a higher minimum energy.

Heat capacity at low temperature Generally, we have been looking at the heat capacity to find the transition temperatures between different cluster types. The heat capacity of the formed cluster can also be found using the Wang-Landau algorithm. In figure B.16, and our other heat capacity diagrams, this shows as a pink tinge that continues to low temperatures (the heat capacity of the empty icosahedron is approximately 34 k, compared to a peak of approximately 4000 k). Figure B.16 shows that this information is not available for all minimum energies, and in order for the algorithm to predict the heat capacity at low temperatures, we need to have a much lower minimum energy than we need to accurately place the transition temperature. However, this greatly increases the time required to run the simulations, so we do not normally attempt to find this, except when we are interested in trying to find the nature of the species.

The heat capacity at even lower temperature also isn't well modelled in our runs. At very low temperature, even when we use a low minimum energy, we see a strange peak and then the heat capacity decreases to zero. This occurs due to the finite bin size. At low temperature the approximation that summation and integration are the same (equation 2.80) breaks down and as a results the predicted heat capacity become ill-behaved (the algorithm predicts similar heat capacity to that of a two level quantised system). This is moved to below a temperature of $0.05\,\varepsilon/k$ by the use of a very low minimum energy (it needs to be -28ε for this system), but, as long as the system clusters, the peak is small and has a distinctive shape that allows easy identification.

Centred dodecahedra-forming system

Unfortunately the predictions of the Wang-Landau algorithm do not always vary so simply as a function of the minimum histogram energy. For the centred dodecahedra-forming system, this complicated behaviour is exemplified when we consider the predictions of the algorithm at different relative sizes of core and shell particles. We initially ran the system at a variety of internal particle sizes using a minimum Wang-Landau energy of -25 ε_{AA} (figure B.17(a)), which gave results in under a few hours' computing time (though this is quite a high minimum energy to use). We then used a varying minimum Wang-Landau energy based on the value achieved during the dynamic runs (figure B.17(b)), these took significantly longer to run (it took approximately 3 months of processor time per simulation), but should have given more accurate results.

The overall shapes of the two figures are clearly similar, but, surprisingly, there appears to be more noise in the system with a lower minimum energy. This apparent noise has two main components. Firstly, we resolve the transitions between clusters with between 15 and 19 particles that simulations have shown exist in this system; our previous minimum energy was above the minimum energy of some of the more strained clusters, and thus did not resolve the transitions. Secondly, these smaller centred clusters do not use all particles in the box, and as we have shown in section B.3.1, this leads to a second peak in the heat capacity that is just an artefact corresponding to clustering of the remaining particles. In this case, this is simply an unstructured cluster similar to that found at very small central particle sizes. A few runs with large numbers of particles were also attempted, but they had many more peaks, most of which could not be so easily separated. Figure B.18 shows the transition lines and relevant regions (estimated using kinetic data). Note that our diagram excludes the formation of larger centred clusters seen in the simulations owing to the number of particles simulated.

These results clearly show that the minimum energy can have an effect on our heat capacity prediction, but that the general behaviour is identifiable even when we use a high minimum energy.

B.5.2 The effect of changing the number of bins

The number of bins used for the energies would be expected to have a large impact on the resolution of the algorithm. When we started using the Wang-Landau algorithm (on

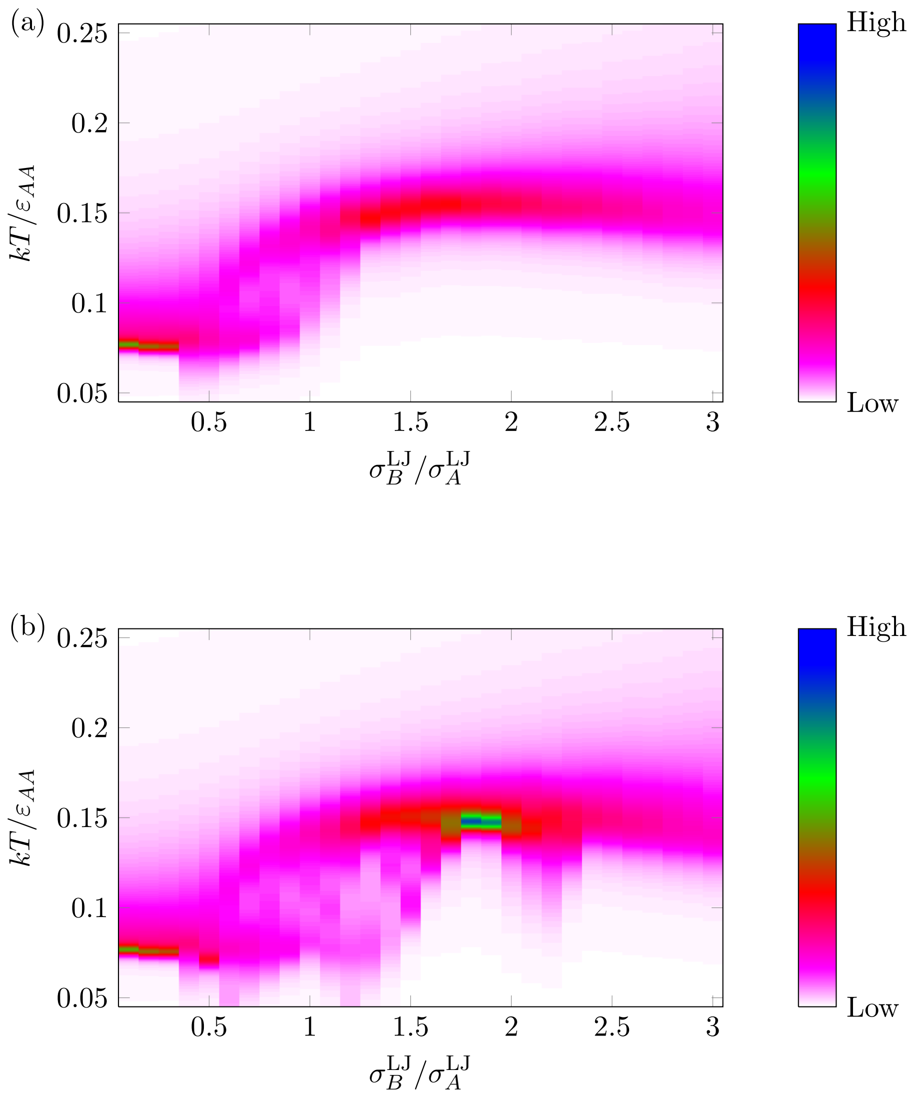

Figure B.17: The heat capacity of a system of 20 shell and one core particle as a function of temperature and interaction strength between the shell and core particles with (a) a minimum Wang-Landau energy of -25 ε_{AA} and (b) a varying minimum Wang-Landau energy based on the value achieved during the dynamic runs. Other settings are the same as table 3.1.

the centred icosahedra-forming system) 20,000 bins were used to attempt to minimise the limiting effects of this resolution, and the width of the bins was permitted to change so that, in areas where the density of states was higher, better resolution was obtained. This meant that, in order to get a reasonable quality of data, it took many months of computing time for each point to be calculated. We compared the effect of using fixed width bins and changing the number of bins on the quality of results obtained, and the time taken for the algorithm to converge. Figure B.19 shows the difference in results for the heat capacity when 20000 or 200 bins are used for the centred icosahedral system. The results do not differ in any significant way. We also tested the uncentred system for a variety of numbers

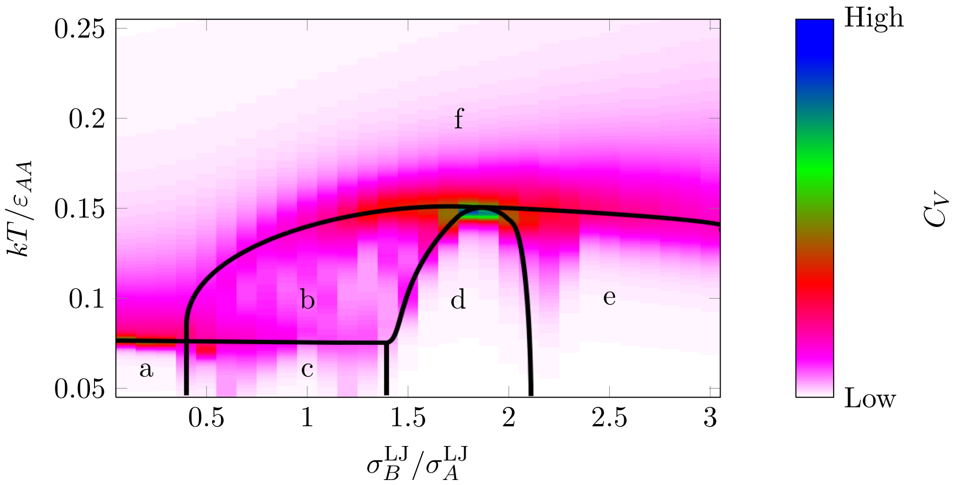

Figure B.18: Approximate relevant transition lines overlaid onto the heat capacity plot for a system of 20 shell and one core particle as a function of temperature and interaction strength between the shell and core particles with a variable minimum energy. The labels denote: (a) the aggregate region, (b) the region of small centred clusters (and monomers), (c) the region of small centred clusters and unstructured uncentred clusters, (d) the region of successful centred dodecahedron assembly, (e) the region where we see unstructured uncentred clusters (we expect more structure in the heat capacity of this region, but other runs have not been able to resolve this), and (f) the monomer gas region.

of bins, and this did not show an effect as long as the number of bins was greater than 100.

More surprising is the time taken for the convergence of the algorithm. At low numbers of bins the time taken for convergence is actually longer than for higher numbers of bins! For the uncentred icosahedra-forming system the fastest convergence was seen with 200 bins, with more bins having a steadily increasing time for convergence in proportion to the number of bins.

Low temperature effects

As mentioned in section B.5.1, the heat capacity at very low temperature can become ill behaved. It is not just the minimum energy that is relevant but also the number of bins used. The more bins that are present, and the lower in energy they sample, the lower the temperature at which the algorithm predicts strange heat capacity behaviour. Since we are normally just looking for the location of the transition, and using many bins and low energy means the algorithm takes considerably longer, we generally used 200 bins for most of our runs.

B.5.3 Flat histograms and the Wang-Landau algorithm

The time taken for the Wang-Landau algorithm to complete is also strongly dependent on the choice of definition of a flat histogram. If we wait for the visit histogram to become truly

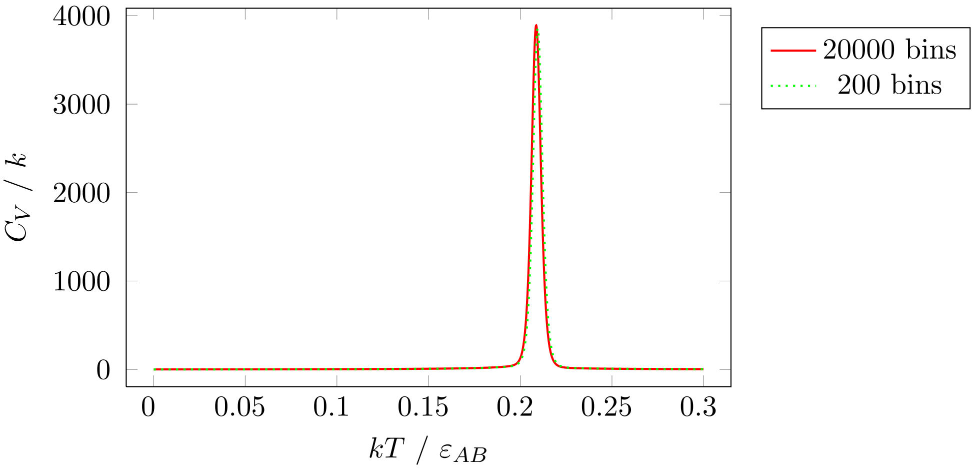

Figure B.19: Comparison of predicted heat capacity for differing numbers of bins for the centred icosahedra forming system.

flat then convergence takes an infinite time, even though the results would be expected to give the clearest picture. In fact, owing to the nature of the algorithm, it is possible to have the simple requirement that all bins have within 90% of the mean visits as the condition for "flatness" without affecting the final outcome. Simpler rules, for example, that each bin must just be visited a thousand or so times give similar final results, but may take longer to converge as the F-value decreases quickly at the start so more visits are needed later, if random noise causes uneven initial visits (thus an uneven estimate of the density of states).

B.6 Using an internal degree of freedom in a simulation

As mentioned in section 1.3.1, proteins may effectively possess an internal degree of freedom, which allows them to change the nature of their inter-protein interactions. Although the first generation of patchy particles is not expected to have this internal degree of freedom, we look at the utility of the principle of internal degrees of freedom for some self-assembling systems.

B.6.1 Different final target structures

We first looked at creating a system which could form either an icosahedron, or a tetrahedron, depending on the state of the internal degree of freedom. The monomer conformation which formed a tetrahedron was set to be lower in energy than the icosahedra-forming conformation, and the difference in energy is $U_{\mathrm{Ih}} - U_{\mathrm{Td}} = \Delta_{\mathrm{trans}} U$. Thermodynamically, we expect the icosahedron to be more stable at lower temperatures, but the tetrahedron should be able to form more quickly when both have the same energies. We aimed to create a system able to change between states with different stable monodisperse clusters present. We varied the relative internal energies of the two forms. Particles were only attracted to other particles which had the same configuration, and the interactions for the attractions were equal in both configurations ($\varepsilon_{AA} = \varepsilon_{BB}$). We looked at the dynamic yields when a system was started in a box full of tetrahedra or icosahedra forming particles, and these can be seen in figure B.20.

Figure B.20(a) shows that the formation of tetrahedra occurs at transition energies above 0.2 ε_{AA}. Below this energy, we are able to form both tetrahedra and icosahedra. It should be noted that the temperature below which we normally see aggregates for the uncentred icosahedral system is at approximately 0.1 ε_{AA}/k (see figure 3.5), which is the same temperature that we see aggregates in this system where we have no internal energy change between the two states. The aggregates formed have similar structure, and indeed most particles are in the higher energy state (which would be expected to form an icosahedron). We investigated whether the formation of tetrahedra was kinetic or thermodynamic in origin by simulating the equivalent system, but starting each particle as an icosahedron-forming particle (figure B.20(b)). This system shows a different behaviour in the low temperature region when the transition energy is low. We then looked at the thermodynamics using the Wang-Landau algorithm.

The thermodynamics of the system can be seen in figure B.21. The thermodynamics suggest a slightly different transition pattern to our kinetic runs. When the transition energies are 0.3 ε_{AA}, we see only a transition to tetrahedra (note, that this transition is wide as we are, in effect, simulating the formation of 3 tetrahedra, thus have some size effects (see section B.3 for more details)). When the transition energy is below 0.25 ε_{AA}, we see a single transition between icosahedra and monomers only. This suggests that the tetrahedra we see formed in this region are a kinetic trap, caused by the rapid assembly; once part of a complete cluster has formed, it is too energetically costly for particles to change their internal structure.

B.6.2 Encapsulation and flipping

Chapter 3 deals with some of the uses of encapsulation to improve the yield of self-assembling structures. Using the centred icosahedral system described in section 3.5, we designed a system where the monomer could have one of two states, being either a core or a shell particle (with the shell particle being lower in energy, and being the initial state for all particles).

Figure B.22 clearly shows that the centred icosahedral system is stable where the energy gap is below a flipping energy of 2.75 ε_{AA}. This is energy much lower than one would expect

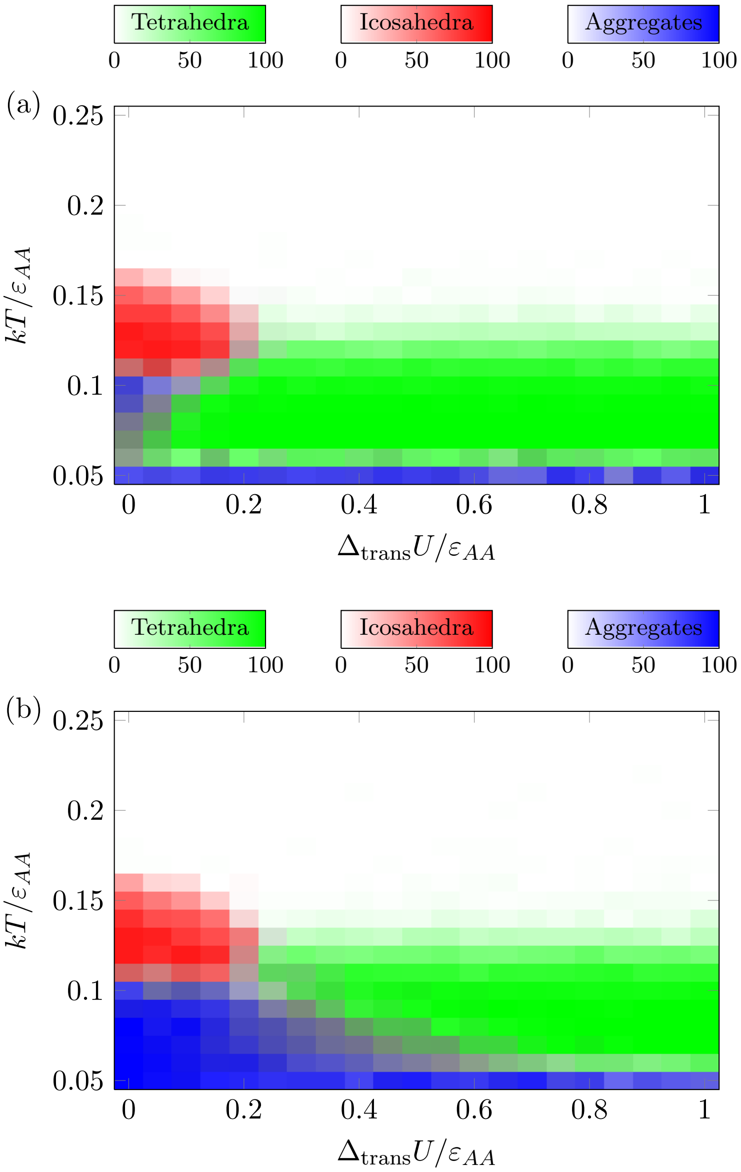

Figure B.20: Yield of icosahedra and tetrahedra starting from (a) a box of tetrahedra forming particles and (b) a box of icosahedra forming particles as a function of energy difference of the two states. $\sigma^{\text{PW}} = 0.35$, $\dfrac{N}{V} = 0.15\,(\sigma^{\text{LJ}})^{-3}$, $U_{\text{Ih}} - U_{\text{Td}} = \Delta_{\text{trans}}U$.

from simple energetic arguments (there is a possible gain of $7\,\varepsilon$ when a bonded shell particle turns into a core particle and becomes fully bonded). This may be a useful property for particles to have, though we appreciate that this type of behaviour in a man-made system is far in advance of current experimental science.

The thermodynamics of this system were not run, owing to the possible effect of the number of particles in a small simulation box. If 13 particles were run, then the system would be unfairly biased towards forming a centred icosahedron, while if 12 particles were run, the system would be unable to form a centred icosahedron. If we were to use

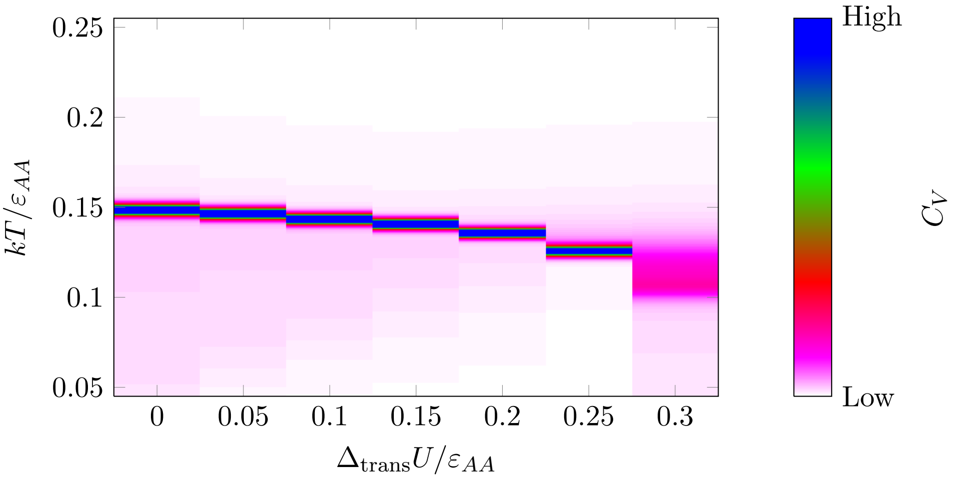

Figure B.21: Plot of the heat capacity of the icosahedra / tetrahedra forming flipping system as a function of the energy difference of the two states.

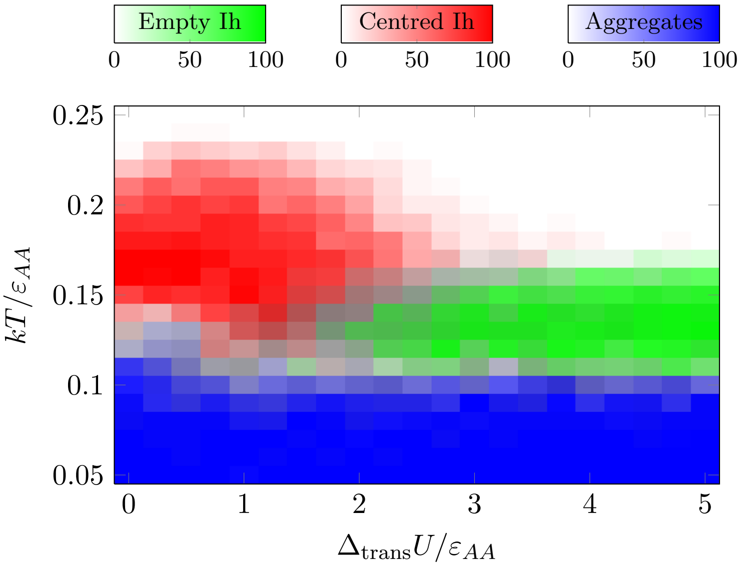

Figure B.22: Yields of centred and uncentred icosahedra when particles may flip between being shell and core particles as a function of the energy of transition. This simulation uses 130 particles and all other settings are given in table 3.1.

larger boxes then we would have significantly longer times for convergence, as well as wide transitions that may overlap with the region of interest (see section B.3 for further details).

B.7 The use of graphics processors in patchy particle Monte-Carlo simulations

Recent technological and software developments have led to a sudden increase in the ease of programming for the multi-parallel units within graphics cards [248]. General-purpose computing on graphics processing units (GPGPU) programming, as the paradigm has come to be known, involves making use of the hugely parallel nature of graphics cards to run calculations.

We used a CUDA port of a simple test version of our code and ran it on nVidia graphics chips. This is an area for future development, and we just present initial results that prove the method is viable. Thanks is given to the Oxford supercomputing centre for the use of their `skynet` cluster for use in some of the initial testing. All of our test runs use single precision floating point numbers, as the implementation of faster double precision floating point libraries and processing methods for GPGPU programming were still in development. A single precision floating point number has approximately 8 significant figures of accuracy for most numbers and calculations [249], which is sufficient for our needs.

We looked at the difference in the speeds of calculating the energy of particles in a model Lennard-Jones system using a single core of an AMD Opteron$^{\mathrm{TM}}$ 2350 processor (with a clock speed of approximately 2 GHz) with either a Fortran or C version of the calculation code, to using a single nVidia Quadro FX 370 graphics card (driver version 185.18.08, core clock speed 360 MHz) and a CUDA ported version. Both were run on the same desktop computer, which was temporarily removed from the network and run with few other programs. We chose random positions and rotations of 1024 particles, then found the difference in time when calculating the energy of a different number of particles in turn (a different number of "moves" being attempted with the energy of just the "moved" particle being calculated) between the two different architectures. The particle which was "moved" was simply the step number modulo the number of particles, since this ensured that there were no effects due to random number generation. In order to ensure that the calculations were performed in every step, in was necessary to disable certain optimisations in the compiler. The time taken in the CUDA case includes the time required to copy all of the data between the graphics and main memory. GPU computation is fast for highly parallelisable runs, however, it is not for highly branched programs. The energy calculations of Lennard-Jones spheres is parallelisable, while calculations of cluster sizes and other analyses are not expected to be.

The results are shown in figure B.23, and clearly show that the CUDA implementation is faster than the Fortran implementation for large numbers of steps. However, at low step numbers the time taken to copy data to/from the graphics card becomes significant. GPU programming may not therefore be suitable for runs which required a high level of branching at each step (e.g. umbrella sampling of cluster size). The plot in figure B.24 shows the ratio of times taken, with a logarithmic x-axis. The plot shows that at least 100 GPU-steps need to occur before the contents of the memory of the graphics card and motherboard are copied to each other, in order for us to get a speed advantage over

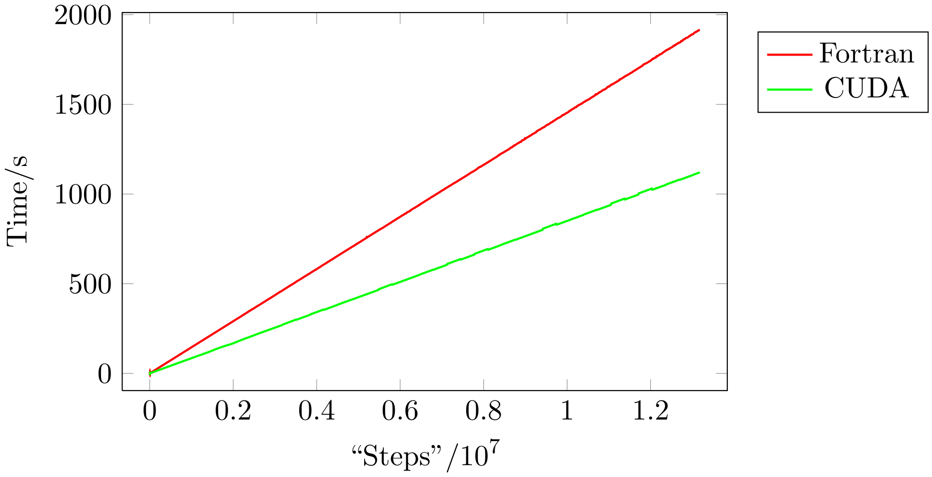

Figure B.23: Time taken to calculate the energy of one particle interacting with all other particles in a system of 1024 particles a certain number of times (equivalent to a certain number of moves) for the Fortran and CUDA programs

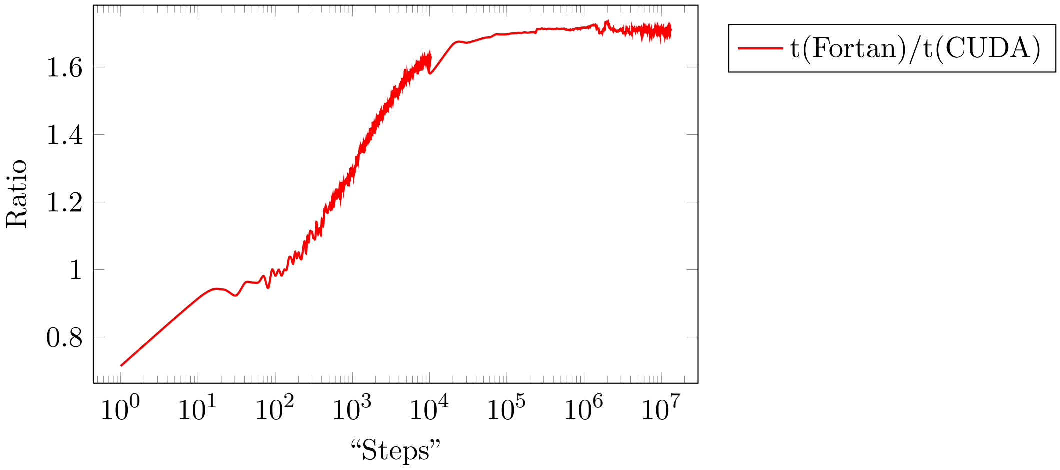

Figure B.24: Ratio of times taken to calculate the energy one particle with all other particles in a system of 1024 particles a certain number of times (equivalent to a certain number of moves) using the Fortran and CUDA versions of the programs.

the CPU-only implementation. Combining CPU and GPGPU programs is easy, though requires some memory coordination. We can see that an approximate upper limit in increasing the speed of the program is a factor of 1.7. This is an impressive increase, but the time required to fully re-implement the code means that this is not as significant as it appears at first.

B.7.1 Conclusions on GPGPU programming

The results clearly show that there is utility in GPU programming, and that it would be worthwhile implementing future code to use this more effectively. There are, however, several barriers to use including the need to learn a new programming language fully (especially one that is still changing), and new issues involving memory usage and thread occupancies. Anyone who is interested in learning more about this should read the plethora of information available on the nVidia website dedicated to CUDA [250], and we recommend that they also consider using the newly standardised openCL programming language too [251].

B.8 Conclusions

We have shown that changing the parameters we have used for our simulations does not strongly modify general behaviour. While the temperature of clustering may change slightly, the overall behaviour generally remains the same (clustering happens, and encapsulation gives correct formation over a wider range of conditions than non-encapsulated assembly). At certain extremes we have found the formation of aggregates occurs more easily, and that we are able to destabilise thermodynamic liquid phases by decreasing the interaction range.

We have also shown that our choices for our settings for our Wang-Landau thermodynamic runs agree with other settings, and that our settings give faster convergence than many others, though sometimes at the expense of predicting the low temperature heat capacity somewhat inaccurately.

We have shown that internal degrees of freedom may be used to give some interesting properties, and to create meta-stable states that persist for a long while. This behaviour may have many useful applications, for example, it could be used as an indicator as to whether a blood product has been kept in reasonable conditions. If the icosahedral-tetrahedral flipping system were created with mainly tetrahedra formed, then it would remain so until the product were heated and then cooled, when icosahedra would form. If these states had different measurable properties (colour would be ideal, though the chemistry required may be complex), then this would make an excellent indicator.

The use of general purpose graphics processor units programming has been shown to give a speed up compared to our current CPU runs. This is worth further investigating, and will hopefully be taken on as a project by a new member of the group.

Appendix C

Mathematical derivations

This appendix deals with some details of mathematical derivations that are skipped over in the text.

C.1 Heat capacity as the variance of energy

The heat capacity at constant volume can be described as the differentital of the expectation value of in internal energy with respect to the temperature, *i.e.*:

$$C_V = \left.\frac{\partial \langle U(T) \rangle}{\partial T}\right|_V . \tag{C.1}$$

We know that the expectation of the energy can be written in terms of an integral over the energy, the density of states ($g(U)$) and the Boltzmann factor,

$$\langle U(T) \rangle = \frac{\int U \times g(U) \times \mathrm{e}^{(-U/kT)} \mathrm{d}U}{Q(T)}, \tag{C.2}$$

and that the partition function, Q, also depends on the temperature,

$$Q(T) = \int g(U) \times \mathrm{e}^{(-U/kT)} \mathrm{d}U. \tag{C.3}$$

We rewrite equation C.2 in a simpler form, as:

$$\langle U(T) \rangle = \frac{\upsilon(T)}{Q(T)}, \tag{C.4}$$

We can then differentiate this using the product rule,

$$\left.\frac{\partial \langle U(T) \rangle}{\partial T}\right|_V = \upsilon \left.\frac{\partial \frac{1}{Q}}{\partial T}\right|_V + \frac{1}{Q} \left.\frac{\partial \upsilon}{\partial T}\right|_V . \tag{C.5}$$

We find the value of each part, using the chain rule multiple times:

$$\begin{aligned}
\left.\frac{\partial \frac{1}{Q}}{\partial T}\right|_V &= \left.\frac{\partial \frac{1}{Q}}{\partial Q}\right|_V \times \left.\frac{\partial Q}{\partial T}\right|_V \\
&= -\frac{1}{Q^2} \times \left.\frac{\partial \int g(U) \times \mathrm{e}^{(-U/kT)\mathrm{d}U}}{\partial T}\right|_V \\
&= -\frac{1}{Q^2} \times \int g(U) \times \left.\frac{\partial \mathrm{e}^{(-U/kT)}}{\partial T}\right|_V \mathrm{d}U \\
&= -\frac{1}{Q^2} \times \int g(U) \times \left.\frac{\partial \mathrm{e}^{(-U/kT)}}{\partial (1/kT)}\right|_V \times \left.\frac{\partial (1/kT)}{\partial T}\right|_V \\
&= -\frac{1}{Q^2} \times \int \frac{U \times g(U) \times \mathrm{e}^{(-U/kT)}}{kT^2} \mathrm{d}U \\
&= -\frac{1}{Q^2} \times \frac{\int U \times g(U) \times \mathrm{e}^{(-U/kT)} \mathrm{d}U}{kT^2}
\end{aligned} \tag{C.6}$$

Similarly, we can see that:

$$\left.\frac{\partial v}{\partial T}\right|_V = \frac{\int U^2 \times g(U) \times \mathrm{e}^{(-U/kT)} \mathrm{d}U}{kT^2} \tag{C.7}$$

Thus, our overall equation is:

$$\begin{aligned}
C_V &= v \left.\frac{\partial \frac{1}{Q}}{\partial T}\right|_V + \frac{1}{Q} \left.\frac{\partial v}{\partial T}\right|_V \\
&= \int U \times g(U) \times \mathrm{e}^{(-U/kT)} \mathrm{d}U \times -\frac{1}{Q^2} \times \int \frac{U \times g(U) \times \mathrm{e}^{(-U/kT)}}{kT^2} \mathrm{d}U \\
&\quad + \frac{\int U^2 \times g(U) \times \mathrm{e}^{(-U/kT)} \mathrm{d}U}{Q \times kT^2} \\
&= \frac{\langle U(T) \rangle^2 + \langle U^2(T) \rangle}{kT^2}.
\end{aligned} \tag{C.8}$$

C.2 Angular frequencies in partition functions

We can use the standard identity of the angular frequency (where μ is the effective mass, *i.e.* the reduced mass in the case of vibrational motion, and the moment of inertia in the case of torsional motion) [252]:

$$\omega = \sqrt{\frac{k}{\mu}} \tag{C.9}$$

We can similarly use the standard equation for the spring constant (k), relating to degree of freedom, X [252]:

$$k = \frac{\partial^2 U}{\partial X^2} \tag{C.10}$$

Thus the angular frequencies are given by:

$$\omega^2 = \frac{1}{\mu}\frac{\partial^2 U}{\partial X^2} \tag{C.11}$$

In our system we therefore have:

$$\begin{aligned}\omega^2_{\text{RAD}} &= \frac{2}{m}\left(\frac{\partial^2 U^{\text{LJ}}}{\partial r^2}\right)_{r=2^{1/6}\sigma} \\ &= 72 \times 2^{2/3} \times \frac{\varepsilon}{m\sigma^2}\end{aligned} \tag{C.12}$$

$$\omega^2_{\text{ANG}} = -\frac{\varepsilon}{I_{\text{mon}}}\left(\frac{\partial^2 U^{\text{ANG}}}{\partial \theta^2_{kij}}\right)_{\left.\begin{smallmatrix}\theta_{kij}\\ \theta_{lji}\end{smallmatrix}\right\}=0} = \frac{\varepsilon}{I_{\text{mon}}(\sigma^{\text{PW}})^2}. \tag{C.13}$$

For bending motion, the relevant derivative is that with respect to angle, and so we must replace the reduced mass with the moment of inertia in equation C.11. Thus,

$$\omega^2_{\text{ANG}} = -\frac{\varepsilon}{I_{\text{mon}}}\left(\frac{\partial^2 U^{\text{ANG}}}{\partial \theta^2_{kij}}\right)_{\left.\begin{smallmatrix}\theta_{kij}\\ \theta_{lji}\end{smallmatrix}\right\}=0} = \frac{\varepsilon}{I_{\text{mon}}(\sigma^{\text{PW}})^2}. \tag{C.14}$$

Appendix D

Program information

D.1 MCP

The program central to this thesis (MCP) is based on MCPatches by A W Wilber [46]. Approximately 25 % of the code derives directly from MCPatches. It is maintained on an svn repository with the whole group contributing and maintaining the code base. The majority of the current version of the code has been written by myself, Alex Wilber and Aleks Reinhardt. Other contributors include Jessica Andréani, Eva Noya, Parvinder Thiara, Gabriel Villar, Matthew Llewelyn-Jones, Trinidad Méndez-Morales, and Jonathan Doye.

D.1.1 Settings for approximate diffusional behaviour

In order for the Monte Carlo moves to give results which mimic diffusion the following settings were used [46]. Maximum move size was set to $0.3\,\sigma^{\mathrm{LJ}}$, the maximum rotation move size 0.3 (this is the maximum length of the quaternion which is added and then the whole new quaternion normalised), probability of attempting a rotation move was 0.4.

D.1.2 A note on compilers

We used the Intel Fortran compiler, version 9.1 (20060818). This has a few bugs that people should be aware of if attempting to use it. The most interesting bug was found when attempting Gibbs Ensemble volume moves. If the change in move size (`deltaV`) is randomly chosen, and the new edge length of box A (`newSideA`) and B (`newSideB`) was calculated using

$$\begin{aligned} &\texttt{newSideA} = \texttt{(oldVolA-deltaV)**(1.0/3.0)} \qquad \text{and} \\ &\texttt{newSideB} = \texttt{(oldVolB+deltaV)**(1.0/3.0)}, \end{aligned} \tag{D.1}$$

then there is a systematic error of the order of 10^{-4} in the total volume at each step. In a standard length run this could cause an error in the volume of up to 50 % if left unchecked. Forcing all intermediate values calculated to be double precision changes the error to a random error with maximum size of approximately 10^{-6}.

D.2 Other programs used

Visualisations of our system and other proteins were created using VMD [253]. The output of the program was in protein databank format [213]. The images were rendered using POV-Ray [254], and edited using the GIMP.

Perl was used as a scripting language to deal with the vast amounts of data produced by the runs. The scripts were maintained by the group on its Subversion repository, and are available on request.

This thesis was produced using pdfLaTeX. Many illustrations were produced using PGF/TikZ, Gnuplot and ImageMagick.

Appendix E

Publications

References [4–7] include work from this thesis.

References [1–3] are papers to which I contributed during my DPhil by assisting other groups which have been accepted for publication.