



Universiteit Utrecht

MASTER THESIS


Phase and vacancy behaviour of hard slanted cubes: a quantitative analysis on the influence of shape

Author:
Robin van Damme

Email:
Robinvdamme@gmail.com

Supervisor:
Dr. Laura Filion

Utrecht, The Netherlands
July 1, 2016



Abstract

In this thesis, we study the phase and vacancy behaviour of hard slanted cubes (right rhombic prisms) as a function of their shape, defined through the shape parameter θ which is the slant angle of these particles. To study the phase behaviour, we use of Monte Carlo simulations in the isotension-isothermal ensemble. We find that these particles form three different crystal structures depending on the slant angle and the density: a simple cubic crystal, a sheared cubic crystal and a hexagonal random tiling crystal. We summarize these results in a phase diagram. To study the vacancy behaviour, specifically that of the simple cubic crystal, we use free energy calculations; both standard Einstein integration methods as well as a scheme which employs thermodynamic integration over the shape. This scheme will be explained in detail, as it could be useful for applications in other systems as well. We find that the vacancy concentration in this simple cubic crystal of slanted cubes does not depend on the shape of the particles, but only on the packing fraction.

Contents

1	Introduction	3
1.1	Model system	4
1.2	Overview of the contents of this thesis	5
2	Monte Carlo simulation methods	6
2.1	An introduction to statistical thermodynamics	6
2.2	The canonical ensemble	7
2.3	A general note on simulations	8
2.4	Monte Carlo	8
2.4.1	Checking for overlap with the separating axis theorem	9
2.4.2	Monte Carlo simulations in the canonical ensemble	11
2.4.3	Monte Carlo simulations in the isotension-isothermal ensemble	12
2.5	Consistency check: comparison of equations of state for hard cubes	13
3	Standard methods to compute the free energy of solids	15
3.1	Thermodynamic integration	15
3.2	The Frenkel-Ladd (Einstein crystal) method	16
3.3	Einstein molecule method	17
3.4	Method: Free energy of solids	19
3.5	Consistency check: Free energy of a hard sphere FCC crystal	21
4	Crystal phases of slanted cubes	22
4.1	Possible close packed crystals	22
4.2	Isotension Monte Carlo simulation results	22
4.3	Free energies of the sheared cubic, alternating cubic and hexagonal lattices	27
4.4	Phase diagram of slanted cubes	28
5	Using shape switch MC simulations to calculate free energy differences	29
5.1	Shape as a thermodynamic variable	29
5.2	Obtaining a free-energy difference by switching	30
5.3	Implementation of a shape-switch move	30
5.4	Extended sampling / biasing	31
5.5	Methodology of shape-switch simulations	33
5.6	The result: a $\mu_\theta(\theta)$ equation of state	35
6	Phase diagram and vacancy behaviour of a simple cubic crystal of slanted cubes	36
6.1	Shape invariance of vacancy formation cost	36
6.2	Updated phase diagram	37
7	Summary and outlook	39

1 Introduction

Big things are made out of many small things. This seemingly innocuous fact, likely known since the dawn of times, has far-reaching implications. Some of the earliest recorded theories to explain how small things lead to big things were considered by the greek natural philosophers. The most notable and closest to modern science was a theory originating from Democritus and Epicurus a few centuries B.C.E.: that everyday objects were composed of indivisible, discontinuous objects constantly colliding with one another [1]. Plato, in contrast to Democritus and Epicurus, considered reality to be mathematical in nature. To him, geometric shapes made up the fundamental, indivisible constituents of matter; cubes, tetrahedra, octahedra and icosahedra [2]. These theories belong to the group of theories we now call *atomism*. While the underlying small-scale nature of reality in these theories varies, one concept is common among most of them: the behaviour of big things is governed by that of small things.

We now skip ahead more than a thousand years to the year 1662, to the work of the Irish scholar Robert Boyle and the English scholar Robert Hooke on the behaviour of gases [3]. Using the then recently invented air pump, they carefully measured the relation between the pressure and the volume of a closed container of air, and found an inverse relation between the two that is now known as Boyle's law: $P \propto 1/V$. Similar relations were also found for other properties: Charles' law (in 1780), which relates the volume of a gas to its temperature ($V \propto T$) and Avogadro's law (in 1811), which relates the volume of a gas to its number of constituent moles of molecules ($V \propto n$). These relations later led to the formulation of the ideal gas law: $PV = nRT$, a relation that links all the relevant properties of a gas composed of particles that do not interact with one another. They are all *static* laws; they deal only with the behaviour of a gas at rest, or rather one at *equilibrium*. While there may be movement in the gas, it is such that the large-scale quantities (pressure, volume, temperature) do not change. In other words: while the arrangement of small things may change, the big things do not.

That the constituents of matter undergo constant motion was seen in the experiments of the Scottish botanist Robert Brown in 1827 [4], when he observed in a system of pollen what is now known as Brownian motion: the constant, random motion of small particles in a fluid due to collisions with even smaller (and to Brown, invisible) particles that make up the fluid. This experiment had two very important implications. The first was the existence of these tiny (and back then, invisible) particles known nowadays as atoms and molecules. Brown's work, along with that of Einstein in 1905 [5] and Jean Perrin in 1908 [6, 7], was the end to the discussion on the existence of atoms [8]. The second important implication was that the thermally driven motion of these tiny particles causes the motion of larger particles, forming a link between the behaviour of the small (microscopic) and the large (macroscopic). Around 1870, the formal theory of equilibrium statistical thermodynamics was constructed by James Clerk Maxwell [9], Ludwig Boltzmann and Josiah Willard Gibbs [10], providing the theoretical formulation of the link between the microscopic and macroscopic phenomena of matter. This exceptionally versatile framework has survived the coming of both quantum mechanics as well as the theory of relativity, and is the formulation we still use today. *How* we use it, however, is considerably different from its usage in the early 20th century. All but the simplest of systems are difficult to study analytically, and require approximate theories to be used instead of exact ones. The invention of computers and the exponential increase in their computational power, however, allows us to numerically evaluate very complex systems without using significant approximations. As with any newly developed field, the first systems were simple. In 1953, Metropolis et al. [11] performed simulations to determine the equation of state of a system of hard spheres; spheres whose only interaction with each other is an infinite repulsion at contact. These particles are easy in some ways: their interaction is short-ranged and isotropic, making evaluation easy and cheap. But as the years went by and computing power grew, more and more complex systems could be simulated. The emergence and ever-increasing control over the synthesis of colloids, particles of sizes between a few nanometers (10^{-9} m) and a few micrometers (10^{-6} m) that are small enough to undergo Brownian motion (like Brown's pollen) yet large enough to be considered hard, caused an increased interest in the influence of particle shape on their macroscopic behaviour. It is here that we find the context for the present work.

Over the last decade, there has been great deal of interest in the phase behaviour of anisotropic particles. On the experimental side, advances in colloid synthesis allow for the creation of Janus particles [12], patchy particles [13], cubic [14–17], rhombic [17–19], hexagonal [17] or tetrahedral [17, 20] supraparticles, as well as many other shapes. The enormous increase in computing power, on the other hand, now allows for simulation of these colloidal systems involving thousands of particles. One particular type of system of a more fundamental interest is that of the hard core repulsive systems, whose phase behaviour is governed solely by entropy [21]. Here, the anisotropy is given by the shape of the particle. The last few years saw a number of studies investigating the influence of shape. Some investigated the phase behaviour of convex polyhedra [22, 23], platelets [24, 25] or tetrapods [26, 27], others how this behaviour changes when going from spheres to cubes [28, 29] or cubes to octahedra [30]. Of particular interest are the crystal phases formed by these particles. In 2012, Smallegange

et al. published a paper [31] in which they describe a system of hard cube-shaped particles which exhibit an extraordinarily high vacancy concentration (up to 8%) for a crystal at densities near the fluid-solid coexistence region. Compared to the hard sphere system, which has vacancy concentrations $\leq 0.001\%$ [32, 33], this is remarkable indeed. The presence of these vacancies in the cube system actually stabilizes it, allowing a simple cubic crystal with defects to remain stable at a packing fraction $\eta = 0.50$, whereas without these vacancies it would melt. The cause of this unusually high concentration is the low entropic cost of vacancy creation. From the conclusion of [31]:

“Clearly, one vacancy provides additional free volume for multiple nearby particles, decreasing the entropic cost of creating a defect. In most other colloidal crystals, such as hard-sphere face-centered cubic crystals, particles near a vacancy are still confined to their lattice site by their remaining neighbors. As a result, the local entropy gain from a defect is much lower.”

The obvious question is now: does this also happen for other systems? If we, for instance, were to make the particles more anisotropic by stretching them along one axis, would they still form a simple cubic lattice? And how would the vacancy behaviour change? An elegant way to answer these questions would be to start from cubes, gradually increase the anisotropy of our particles, and then study how their phase behaviour changes. This is what we will do. The next section will describe the model system we will be using: one where we can directly tune the anisotropy of the particles by means of a single shape parameter θ .


1.1 Model system

Our model system consists of particles with a well-defined shape interacting via a hard potential. A hard potential is a pairwise potential where there is an infinite energy cost to let two particles overlap, but as long as they do not, there is no energy cost whatsoever. Such a potential is given by:

$$\beta\phi(\mathbf{r}_1, \mathbf{\Omega}_1, \mathbf{r}_2, \mathbf{\Omega}_2) = \begin{cases} \infty & \text{when } (\mathbf{r}_1, \mathbf{\Omega}_1, \mathbf{r}_2, \mathbf{\Omega}_2) \text{ are such that the particles overlap/interpenetrate,} \\ 0 & \text{when } (\mathbf{r}_1, \mathbf{\Omega}_1, \mathbf{r}_2, \mathbf{\Omega}_2) \text{ are such that the particles do not overlap/interpenetrate,} \end{cases} \quad (1)$$

where $\mathbf{r}_1, \mathbf{r}_2$ are the positions of the two particles, $\mathbf{\Omega}_1, \mathbf{\Omega}_2$ are their orientations, $\beta = 1/k_B T$ with k_B Boltzmann's constant and T the temperature, and $\phi(\mathbf{r}_1, \mathbf{\Omega}_1, \mathbf{r}_2, \mathbf{\Omega}_2)$ is the pair potential. Such a model system is a very good approximation of everyday objects like tables or billiard balls, and also a fairly good approximation for large uncharged colloids, but it is a rather poor approximation for any system where electromagnetic interactions play a significant role such as charged colloidal or atomic systems. More importantly, however, hard interactions allow us to study solely the entropic contributions to the phase behaviour, as energy minimization forces nonoverlapping configurations with zero energy. The phase behaviour will then purely be determined by entropy maximization. We call these systems “entropy-driven systems”, and they are surprisingly diverse. For an excellent introduction to this field we direct the reader to the works of Daan Frenkel: Refs. [34, 35]. These type of model systems are of particular interest for the field of *colloids*, particles with sizes between a few nanometers (10^{-9} m) and a few micrometers (10^{-6} m) which are large enough to be considered hard particles in many cases, but are small enough to undergo Brownian motion.

To be able to define such a thing as overlap, we need to define a shape for our particles. The particles we will be considering in this thesis will be a specific subset of polyhedra with the proper name “right rhombic prisms”. For simplicity, however, we will refer to them as “slanted cubes”. After the discovery of defect-rich behaviour of hard cubes, the question arose whether this behaviour is also found in particles with different symmetries. A natural path to choose is then to define a model system which resembles a cube and in a limiting case reproduces one, but is otherwise different. For simplicity, one usually defines such a system with a so-called *shape parameter*, so that the shape is a continuous function of this one parameter. These type of systems have been studied before, studying shapes going from circles to squares [25], spheres to cubes [16, 28, 29], or cubes to octahedra [30]. Following [36], we define our slanted cubes by a set of three vectors:



$$\begin{aligned} v_1 &= (\sigma, 0, 0), \\ v_2 &= (0, \sigma, 0), \\ v_3 &= (\sigma \cos \theta, 0, \sigma \sin \theta), \end{aligned} \quad (2)$$

Figure 1: A slanted cube and the vectors by which it is defined.

such that every edge has length σ . Note that the volume of the particle in this definition is $V_{part} = \sigma^3 \sin \theta$, not $V_{part} = \sigma^3$, as is often the case in model systems. See Fig. 1. Note also that in the limit of $\theta \rightarrow 90^\circ$, we obtain a cube, while in the limit of $\theta \rightarrow 0$ we would have infinitely thin rods ($L/D \rightarrow \infty$). In this thesis we will restrict our studies to $\sim 45^\circ \leq \theta \leq 90^\circ$. Smaller angles would be a system more akin to narrow platelets, but this choice of shape is not a good model system for real physical systems of platelets nor of elongated particles, and hence we will not study it here.

One might wonder whether these slanted cubes are merely interesting from a theoretical perspective. This is not the case. While indeed no system in nature will have truly hard edges, systems of rounded hard cubes can be made reliably for a number of years already [14, 15, 37], and systems of hard slanted cubes in the form of rhombic platelets have also been made [18, 19, 38].

1.2 Overview of the contents of this thesis

This thesis is organized as follows: in Chapter 2 we review the basic theory of statistical thermodynamics, how it is applied in Monte Carlo simulations and the “floppy box” method, a relatively modern modification to the standard Monte Carlo simulation of systems in the NPT ensemble that is required to properly model crystals and crystal-crystal transitions. Chapter 3 deals with some of the most common ways to numerically calculate the free energies of complex systems: thermodynamic integration and the Frenkel-Ladd or Einstein crystal method. The Einstein molecule method, a modern variant on the Einstein crystal method and the method we will be using, is also discussed in greater detail. In Chapter 4 we apply the Monte Carlo algorithms described in Chapter 2 to our model system of slanted cubes, discuss the various crystal phases we find and construct a phase diagram of the slant angle-density plane. In Chapter 5 we describe a somewhat new scheme to obtain the free energy of a system by thermodynamic integration over the shape of its constituent particles. This scheme is then applied in Chapter 6 to calculate the free energy of a simple cubic crystal of slanted cubes, which in turn allows us to study the vacancies that are present in this crystal. We also use the knowledge of the free energies of the various phases to construct a more accurate version of the phase diagram found in Chapter 4. Finally, Chapter 7 summarizes the results of this thesis and provides some possible directions for further research.

2 Monte Carlo simulation methods

In this chapter, we will discuss the essential theory and the various components necessary to investigate the phase behaviour of many-particle systems through the use of Monte Carlo simulations. We shall start with a conceptual introduction to statistical thermodynamics, after which we will look at the thermodynamic quantities that are of particular interest for studying with simulations. Then, following a brief general note on the necessity, applications and available methods of simulation, we shall delve into the inner workings of Monte Carlo simulations. First, we review the basics: how to sample phase space with Markov chains, Monte Carlo moves and the Metropolis criterion, the principle of detailed balance and the process of equilibration. We describe how, for hard particles, the Metropolis criterion reduces to checking whether two particles overlap with each other and how we check for this overlap using the separating axis theorem. We then turn to the method by which we sample the available phase space, outlining the various Monte Carlo moves required in simulations of the canonical, isothermal-isobaric and isothermal-isotension ensembles. To conclude this chapter, we check whether the simulation we have constructed yields the correct results by simulating a system where the behaviour is known; we calculate the equation of state for a system of hard cubes, comparing our results with those of Smallenburg et al. [31]. We find an excellent agreement, which indicates our simulation is working as intended and we can move on to the system of interest: the slanted cubes.

2.1 An introduction to statistical thermodynamics

In the field of statistical physics, we are interested in the collective or macroscopic behaviour and properties of system of a large number of particles. Consider a system of identical particles, whose (relevant) individual properties are their position \mathbf{r} , their orientation $\mathbf{\Omega}$, and their momentum \mathbf{p} . These are typically three-component vectors, as we are usually dealing with three-dimensional space. Now, each of the N particles under consideration has its own position and momentum, and these variables each depend explicitly on time. As such, we can define a *phase space* [39] $\Gamma(t) = \{\mathbf{r}(t)^N, \mathbf{p}(t)^N, \mathbf{\Omega}(t)^N\}$ where

$$\begin{aligned}\mathbf{r}(t)^N &= \{\mathbf{r}_1(t), \mathbf{r}_2(t) \dots \mathbf{r}_N(t)\} \\ \mathbf{\Omega}(t)^N &= \{\mathbf{\Omega}_1(t), \mathbf{\Omega}_2(t) \dots \mathbf{\Omega}_N(t)\} \\ \mathbf{p}(t)^N &= \{\mathbf{p}_1(t), \mathbf{p}_2(t) \dots \mathbf{p}_N(t)\}\end{aligned}\tag{3}$$

are the collections of particle positions, orientations and momenta. This is the *microscopic* description of our system of particles. Each possible configuration is described by one unique value of Γ . More specifically, each possible configuration of a d -dimensional system is a point in a $3Nd$ -dimensional hyperspace spanned by the $3 * d$ components of the positions, orientations and momenta of each particle.

Now let us take a step back. We are rarely really interested the specific value of Γ . Rather, we would like to know something about a property of the system that depends on Γ , say, the energy of the system as given by the Hamiltonian $\mathcal{H}(\Gamma(t))^1$, or the pressure $p(\Gamma(t))$. These are *macroscopic* quantities, which depend on the microscopic properties of the system but are not necessarily uniquely determined by them. In most applications of statistical physics, the time scale at which the microscopic properties of the system change are tiny and far smaller than the time between measurements (for reasons both physical and practical). For instance, say we want to measure the pressure of some air in a box at a certain temperature. Since the pressure is related to the movement of the particles, one can imagine that this quantity can be described by some function of the positions and momenta of the particles in the box: $p(\Gamma(t))$. Realistically, our measurement will take a finite amount of time. So what we measure is not really $p(\Gamma(t))$, but rather its time-averaged value \bar{p} over the duration from the start of our measurement ($t = 0$) to the end ($t = \tau$):

$$\bar{p} = \frac{1}{\tau} \int_0^\tau p(\Gamma(t)) dt.\tag{4}$$

During this time, the positions and momenta of the particles will change considerably. But from our experience and intuition, we know that the air is not likely to experience sudden, drastic changes in pressure. Somehow, the way p depends on $\Gamma(t)$ must be such that for all the configurations Γ the system visits, the pressure p must be almost the same. This is the domain of statistical mechanics: the theory that describes the link between microscopic and macroscopic properties of systems of large numbers of particles. This independence leads to a concept known as *ergodicity*. For a system that is ergodic, the following holds:

$$\bar{A} = \frac{1}{\tau} \int_0^\tau A(\Gamma(t)) dt = \int f(\Gamma) A(\Gamma) d\Gamma = \langle A \rangle.\tag{5}$$

Here, $A(\Gamma)$ is some property of the system such as the pressure, volume or temperature, and $f(\Gamma)$ is a normalized probability distribution which assigns a probability density to every value of Γ . Note that we include no time

dependence in the right-hand side of the equation. This is crucial. For an ergodic system, the probability density $f(\Gamma)$ is independent of time, and we can replace the time average \bar{A} with the ensemble average $\langle A \rangle$.

2.2 The canonical ensemble

In statistical thermodynamics, one often speaks of *ensembles*, often with some adjective such as “canonical”, “NPT” or “grand canonical”. This is just an expression of which parameters in a system of particles are fixed. In a canonical ensemble, we assume that the number of particles (N), the volume (V) and the temperature (T) are fixed. A real-world example of a system we could study in a canonical ensemble is a bottle of water in the refrigerator. There is a set amount of water in the bottle (fixed N), the bottle is mostly rigid (fixed V), and the temperature of the system is given by whichever temperature you set your refrigerator to (fixed T). As we all know, a bottle of water is not instantaneously cool just because you put it in the refrigerator. It has to cool first, and it does this by exchanging energy with its surrounding. This, too, is the case for a canonical ensemble. The fact that this system can exchange energy will lead to a certain distribution of states $f(\Gamma)$. If the temperature of the surrounding (and by extension, the system) is very low, particles will be slow and your bottle of water will likely contain ice, a crystal of water where the molecules sit on well-defined positions. If, on the other hand, the temperature is very high, particles will zip around criss-cross in a gas phase. In other words, there is a preference for certain states Γ that depends on the temperature. This preference, the probability distribution $f(\Gamma)$, was already derived by Boltzmann in 1884 (and later refined by Gibbs in 1902) and is:

$$f(\Gamma) = \frac{e^{-\beta\mathcal{H}(\Gamma)}}{\int d\Gamma e^{-\beta\mathcal{H}(\Gamma)}}. \quad (6)$$

Here, again, $\mathcal{H}(\Gamma)$ is the Hamiltonian, which gives the energy of the system, and $\beta = 1/k_B T$, where T is the temperature and k_B is Boltzmann’s constant. For N identical (indistinguishable) particles, we can rewrite this as:

$$f(\Gamma) = \frac{e^{-\beta\mathcal{H}(\Gamma)}}{\int d\Gamma e^{-\beta\mathcal{H}(\Gamma)}} = \frac{1}{N!h^{3N}} \frac{e^{-\beta\mathcal{H}(\Gamma)}}{Z(N, V, T)}, \quad (7)$$

where h is Planck’s constant and $Z(N, V, T)$ is the canonical partition function:

$$Z(N, V, T) = \frac{1}{N!h^{3N}} \int d\Gamma e^{-\beta\mathcal{H}(\Gamma)}. \quad (8)$$

We can often simplify this expression by evaluating the kinetic part of the integral. The kinetic part of the Hamiltonian is $\mathbf{p}^2/2m$. Writing this out and using the Gaussian integral $\int e^{-x^2/a} = \sqrt{\pi/a}$, we get:

$$\begin{aligned} Z(N, V, T) &= \frac{1}{N!h^{3N}} \int d\mathbf{p}^N e^{-\beta \sum_{i=1}^N \mathbf{p}_i^2/2m} \int d\mathbf{r}^N e^{-\beta\Phi(\mathbf{r}^N)} \\ &= \frac{1}{N!h^{3N}} \left(\prod_{i=1}^N \int d\mathbf{p}^N e^{-\beta\mathbf{p}^2/2m} \right) \int d\mathbf{r}^N e^{-\beta\Phi(\mathbf{r}^N)} \\ &= \frac{1}{N!h^{3N}} (2\pi mk_B T)^{3N/2} \int d\mathbf{r}^N e^{-\beta\Phi(\mathbf{r}^N)} \\ &= \frac{Q(N, V, T)}{N!\Lambda^{3N}}. \end{aligned} \quad (9)$$

Here, $\Lambda = \sqrt{h^2/2\pi mk_B T}$ is the thermal de Broglie wavelength and $Q(N, V, T)$ is the configurational integral, given by:

$$Q(N, V, T) = \int d\mathbf{r}^N e^{-\beta\Phi(\mathbf{r}^N)}. \quad (10)$$

This integral is colloquially known as “the difficult part” of the partition function, as the interparticle potential energy $\Phi(\mathbf{r}^N)$ generally consists of pairwise or many-body potentials. We now also define a Helmholtz free energy

$$F(N, V, T) = -k_B T \ln Z(N, V, T), \quad (11)$$

where we make the crucial note that it follows from the laws of thermodynamics that the equilibrium macrostate of a canonical system is the one that minimizes this free energy.

¹There are systems with Hamiltonians that are also explicitly time-dependent: $\mathcal{H} = \mathcal{H}(\Gamma(t), t)$, but such systems are non-ergodic and therefore not of interest to this present work.

2.3 A general note on simulations

There are only a few systems for which one can derive the phase behaviour analytically. Two such systems, and ones we will use in this thesis, are the well-known ideal gas and the somewhat less well-known (noninteracting) Einstein crystal [40]. Perhaps surprisingly, even for a system as seemingly simple as hard spheres no analytical description of the phase behaviour is known, though not for lack of trying [41]. Instead, spurred on by the exponential increase in computing power, numerical simulations have become the norm for determining the phase behaviour for a variety of systems including those with long-range electrostatic interactions [42, 43], hydrodynamics [44] and extreme polydispersity [34, 45]. The two most widely used algorithms for this are the Molecular Dynamics (MD) and Monte Carlo (MC) algorithms. The former constructs a path² through phase space by integrating the equations of motions, while the latter typically utilizes so-called Markov chains to sample from the probability distribution $f(\Gamma)$. In this thesis, we will mainly be using MC algorithms for various applications, but will also make occasional use of MD as it is generally more efficient at simulating larger systems. The inner workings of the former will be explained in the following chapter, but for the latter we will simply use a version of the MD code used in Ref. [31] that was modified by Frank Smallenburg to be able to handle slanted cubes.

2.4 Monte Carlo

The core of Monte Carlo simulations is the evaluation of the statistical average, which is the integral equation

$$\langle A \rangle = \int f(\Gamma) A(\Gamma) d\Gamma. \quad (12)$$

The simplest way of evaluating this integral is by simply looking at what values the integrand takes at different values of Γ . As we have noted earlier with Eq. 5, it is possible to compute equilibrium properties of a system by sampling phase space, as long as one knows both the dependence of the property on the system coordinates $A(\Gamma)$ and the probability distribution $f(\Gamma)$. However, the latter is not known *a priori*. To calculate this integral numerically, one generates random values of Γ and then checks their contribution to the integral by checking $A(\Gamma)$. Unfortunately, uniform random sampling of values from an at least $2dN$ -dimensional phase space is an extremely inefficient way of sampling. Not only that, but it is also a poor reflection of what happens in actual thermodynamic systems. Real thermodynamic systems and especially solids spend the vast majority of their time in relatively small subvolumes of phase space, with separate phases often occupying distinct and largely non-overlapping volumes of phase space. Instead of generating a completely new phase-space coordinate, it would make more sense to start from a point which we know is inside or close to the relevant subvolume and generate new configurations close to this point. The way to do this is through Markov chains. A Markov chain is a series of steps through phase space, where each new step taken depends only on the current and possible next position. A possible step in a Markov chain would be moving a particle from a position \mathbf{r}_i to a new position $\mathbf{r}_j = \mathbf{r}_i + \Delta\mathbf{r}$, where $\Delta\mathbf{r}$ is small compared to, for example, the particle size or the lattice spacing in a crystal. Whether this step is actually taken depends on some probability of acceptance a_{ij} . The choice for this acceptance probability is important, as we will soon see.

To determine the direction of our Markov chain, we need to introduce the concept of importance sampling. Let us take another look at Eqs. 9 and 12. The fundamental assumption of statistical physics tells us that for a system with a given energy and composition (microcanonical ensemble), every microstate Γ is equally probable so that $f(\Gamma) = 1/\int d\Gamma$. For a system with variable energy but a fixed temperature (canonical ensemble), microstates are weighted according to a Boltzmann distribution (Eq. 6). If we then wish to simulate a system in the canonical ensemble, we should use this knowledge somehow in the sampling path we take through phase space. Otherwise, we will be spending a great deal of time in regions where $f(\Gamma)$ is very small, which contribute very little to the integral we wish to calculate. This is an old problem. Back in 1970, Hastings published a paper [47] providing a general way to bias Markov chains so that their sampling converges onto the thermodynamic equilibrium distribution. We briefly describe this method here. Starting from the principle of detailed balance:

$$\pi_i p_{ij} = \pi_j p_{ji}, \quad (13)$$

where π_i is the probability that a system is in a microstate i and p_{ij} is the probability of a Markov chain step taking the system from microstate i to j . We rewrite p_{ij} in the product of two terms:

$$p_{ij} = q_{ij} a_{ij}, \quad (14)$$

²Many-body systems are chaotic, and computer precision is finite. MD simulations therefore do not predict *the* time evolution of a system, but *a* time evolution of it [46].

where q_{ij} describes the probability of attempting a Markov chain step from i to j , and a_{ij} is the probability of accepting this proposed step. Using the work of Metropolis et al. [11] and taking³ $q_{ij} = q_{ji}$, it is possible to rewrite

$$a_{ij} = \begin{cases} 1 & \text{if } \pi_j/\pi_i \geq 1 \\ \pi_j/\pi_i & \text{if } \pi_j/\pi_i < 1. \end{cases} \quad (15)$$

Now, we *know* the probability distributions π_i . For the canonical ensemble, it is

$$\pi_i = f(\Gamma_i) = \frac{e^{-\beta\mathcal{H}(\Gamma_i)}}{\int d\Gamma e^{-\beta\mathcal{H}(\Gamma)}}. \quad (16)$$

Filling this into the acceptance probability and writing $\mathcal{H}(\Gamma_i) = E_i$, we obtain

$$a_{ij} = \begin{cases} 1 & \text{if } e^{-\beta(E_j-E_i)} \geq 1 \\ e^{-\beta(E_j-E_i)} & \text{if } e^{-\beta(E_j-E_i)} < 1 \end{cases} = \min \left[1, e^{-\beta(E_j-E_i)} \right]. \quad (17)$$

This criterion is the core of the Metropolis algorithm. If we simply accept or reject changes in Γ (which we call MC moves or MC steps) according to this criterion, our sampling will converge to the equilibrium probability distribution $f(\Gamma)$. How fast this convergence will be, however, can depend on a great many factors.

In practice, when we want to simulate a certain system in equilibrium, we take precautions to ensure that equilibrium is actually reached before we start sampling. This can be done by picking an initial configuration that is close to whatever configuration we would like to sample. Additionally, we let the system equilibrate by allowing it to complete a number of MC steps before we start sampling whatever property we would like to measure. Doing this properly mostly comes down to common sense. You do not start from a high density crystal if you want to simulate a low density fluid, or from a simple cubic crystal if what you want is a face-centered cubic crystal. You *could*, but you would have to wait a (possibly approaching infinitely) long time for nucleation processes and the crossing of similar free-energy barriers before you reach an equilibrium configuration.

An important point we should make here is how the Metropolis algorithm is simplified for simulations of hard particles. Simply filling in the hard interaction in Eq. 18 yields

$$a_{ij} = \min \left[1, e^{-\beta(E_j-E_i)} \right] = \begin{cases} 1 & \text{if } i \rightarrow j \text{ causes no overlap} \\ 0 & \text{if } i \rightarrow j \text{ causes overlap} \end{cases} \quad (18)$$

The Metropolis criterion is now simplified to finding out whether two geometric shapes overlap. There are a number of ways to detect this, which we will discuss in the next section.

2.4.1 Checking for overlap with the separating axis theorem

In Section 2.4 we saw that simulating hard particles comes down to detecting whether particles overlap with one another. In this section, we will explain the workings of the separating axis theorem (SAT) algorithm, which can detect whether two convex polyhedra (of which our particles are a subset) are overlapping. For the sake of completeness we should also name the most popular alternative, the Gilbert-Johnson-Keerthi (GJK) algorithm [48]. While both essentially do the same thing, there are reasons for choosing SAT over GJK in our case. The first is ease of implementation. The SAT algorithm is relatively simple to implement, while GJK is more complicated. The second reason is that we will mainly be considering simple (few vertices) convex polyhedra, for which SAT is likely to be the more efficient of the two. While GJK offers more flexibility and better performance for more complex polyhedra, we will not need it for studying slanted cubes. Either way, this is a question of numerical efficiency and flexibility, not one of physics, and so we will dwell on it no more.

The core of the SAT algorithm is, unsurprisingly, the Separating Axis Theorem. This theorem states [49, 50]:

Two convex polytopes are nonoverlapping if there exists an axis for which the projections of the polytopes onto this axis are nonoverlapping.

For reference, a polytope is a geometric object with only flat faces. A polyhedron is a three-dimensional polytope, and a convex polytope has the property that a line connecting any two points inside it does not go outside the object. A visual representation is the easiest way to comprehend this, see Fig. 2.

³When constructing an MC simulation, we have to take care that this is actually the case! In practice, this usually comes down to making sure that for every MC step we take, its opposite could theoretically be taken as the next MC step.

So to check for overlap, we need three things: an axis to check, a way of projecting the polytopes onto this axis, and a way to detect whether the projections are separate. Let us start at the beginning. Which axes should we check? Intuitively, one can guess that this depends on the polytope under consideration and indeed, that is the case. It turns out [49, 50] that the only axes that need to be checked are the surface normals of the polytopes in question, and the cross-product of their edges. Furthermore, most polytopes possess symmetries that make some of these axes degenerate, causing a further reduction of axes we need to check. For slanted cubes, we have 3 unique normals and 3 unique edges per slanted cube, so we need at most check $2 \cdot 3 = 6$ surface normals and $3(3 - 1) = 6$ combinations of edges. We implement the overlap check as follows:




Figure 2: (Left) A convex shape. No straight line from any point on the surface to another goes outside the shape. (Right) A concave shape. Some straight lines from the surface must go outside the shape to reach another point on the surface.

Separating axis overlap check:

0. Overlap check is called for particle i . Another particle j is selected to check for overlap with i . In our case, we base this selection on proximity by using a cell list.⁴
1. A vector \mathbf{T} is defined as the vector between the centers of mass of the particles i and j . We perform two early-exit checks by comparing the length of \mathbf{T} to the radii of the outscribed and inscribed spheres of i and j .
 - If $|\mathbf{T}| > r_{out,i} + r_{out,j}$, the particles are nonoverlapping.
 - If $|\mathbf{T}| < r_{in,i} + r_{in,j}$, the particles are overlapping.

If neither applies, we proceed to step 2.

2. From the normals and edge combinations of i and j , we select an axis to check for separation. We apply the following separating axis check to this axis \mathbf{L} :

2.1 Project \mathbf{T} onto \mathbf{L} to obtain the center-of-mass projection p_C :

$$p_C = \mathbf{T} \cdot \mathbf{L}.$$

If $p_C < r_{in,i} + r_{in,j}$, this axis is not a separating axis.

- 2.2 Project the vertices $\mathbf{v}_{i,k}$ and $\mathbf{v}_{j,k}$ of the particles i and j onto the axis \mathbf{L} . We define the maximum and minimum values:

$$p_i^{max} = \max_k (\mathbf{v}_{i,k} \cdot \mathbf{L}) \quad p_i^{min} = \min_k (\mathbf{v}_{i,k} \cdot \mathbf{L})$$

$$p_j^{max} = \max_k (\mathbf{v}_{j,k} \cdot \mathbf{L}) \quad p_j^{min} = \min_k (\mathbf{v}_{j,k} \cdot \mathbf{L})$$

- 2.3 If $p_C > 0$, check whether $p_i^{max} < p_j^{min} + p_C$. If true, \mathbf{L} is a separating axis. If false, \mathbf{L} is not. If $p_C < 0$, check $p_i^{min} > p_j^{max} + p_C$ instead. If true, \mathbf{L} is a separating axis, and if false, \mathbf{L} is not.

If any of these axes is a separating axis, the particles are nonoverlapping. If none of the checked axes was a separating axis, the particles must be overlapping.

It should be noted that this method is not restricted to MC simulations, but is in fact general and could be implemented much the same in an MD or in any other kind of simulation.

⁴A cell list is a computational trick to quickly sort particles based on proximity. The entire simulation volume is divided into smaller subvolumes (cells) into which particles are sorted. Easy proximity detection can then be done by simply only checking the surrounding cells of a particle. For more information, see [42].

2.4.2 Monte Carlo simulations in the canonical ensemble

One of the simplest Monte Carlo simulations, that is also frequently used, is the simulation of a system of particles in an NVT ensemble. Here, the number of particles N , the volume V and the temperature T are fixed. Typically, one applies periodic boundary conditions in order to simulate a small system in contact with a larger one without any interface or surface effects. One common example of an application of such an MC-NVT simulation is the calculation of the equation of state for systems of (non-hard) particles. One can use the virial expression

$$p = \frac{k_B T N}{V} + \frac{1}{Vd} \left\langle \sum_{i=1}^{N-1} \sum_{j=i+1}^N \mathbf{f}_{ij} \mathbf{r}_{ij} \right\rangle, \quad (19)$$

where d is the dimensionality of system, f_{ij} is the force exerted on particle i by particle j and \mathbf{r}_{ij} is the vector connecting their centers of mass, to calculate the pressure. In such a simulation, we sample over the space spanned by all particle positions and orientations. As such, we must define MC moves that allow us to sample this space. To sample the positions, we use particle translation moves and to sample the orientations, we use particle rotations. A translation in an MC-NVT simulation is done as follows:

Monte Carlo translation move:

- 0 A particle i is chosen for translation.
- 1 The position of particle i before the move is saved temporarily.
- 2 A random vector $\Delta \mathbf{r}$ is generated isotropically by generating its three components randomly from a uniform distribution, such that the magnitude of $\Delta \mathbf{r}$ is smaller or equal to some step size Δ_r .
- 3 The position of particle i is changed: $\mathbf{r}_i \rightarrow \mathbf{r}_i + \Delta \mathbf{r}$.
- 4 The Metropolis criterion is evaluated. For hard particles, this comes down to checking for overlaps using the subroutine we described in Section 2.4.1. If the particle overlaps with any other particles in its new location, the move is marked as rejected. Otherwise, it is marked as accepted.
- 5 If the move is marked as rejected, the position is changed back to the old value that was stored in step 1. Otherwise, the move is accepted.

Similarly, a rotation is done as follows:

Monte Carlo rotation move:

0. A particle i is chosen for rotation.
1. The orientation of particle i before the move is saved temporarily.
2. A random rotation step matrix $\Delta \boldsymbol{\omega}$ is generated by generating a random vector isotropically upon the unit sphere and rotating around this vector by a random angle $\Delta_\phi/2 \leq \phi \leq \Delta_\phi/2$.
3. The orientation of particle i is changed: $\boldsymbol{\omega} \rightarrow \boldsymbol{\omega} + \Delta \boldsymbol{\omega}$.
4. The Metropolis criterion is evaluated. For hard particles, this comes down to checking for overlaps using the subroutine we described in Section 2.4.1. If the particle overlaps with any other particles in its new orientation, the move is marked as rejected. Otherwise, it is marked as accepted.
5. If the move is marked as rejected, the orientation is changed back to the old value that was stored in step 1. Otherwise, the move is accepted.

2.4.3 Monte Carlo simulations in the isotension-isothermal ensemble

Unfortunately, the use of a hard potential means that we cannot use the virial expression to evaluate the pressure. This is because the hard potential is discontinuous, and evaluation of the forces needed for the virial expression requires that the potential is continuous. In order to calculate the equation of state, we must choose a different ensemble. In the NPT ensemble, it is not the volume of the system but the pressure that is fixed. Using this ensemble, we *can* measure the equation of state of hard particles, and quite easily at that. In an MC-NPT simulation, we also allow the volume to fluctuate. In the simplest numerical implementation of the NPT ensemble in three dimensions, one defines a simulation volume (box) of L_x by L_y by L_z , and allows the system to change its volume $V = L_x * L_y * L_z$ isotropically in response to an externally applied pressure p . Consider an MC move that isotropically changes the volume by increasing or decreasing the available simulation volume and scaling all coordinates relative to this volume. The Metropolis criterion for this move is:

$$\text{acc}(o \rightarrow n) = \min \left[1, \exp \left[-\beta U(\mathbf{r}_n^N) + \beta U(\mathbf{r}_o^N) - \beta p(V_n - V_o) + N \ln \left(\frac{V_n}{V_o} \right) \right] \right]. \quad (20)$$

Here p is the applied external pressure, V_o and V_n are the volumes before and after the MC move, respectively, and N is the number of particles. For the derivation of this term, we refer to [42]. The volume move is then performed as follows:

Monte Carlo volume scaling move:

1. A copy of the old volume V ; V_o is stored for later use.
2. A new volume is generated: $V_n = V_o \exp[R\Delta_V]$, where R is a uniformly generated random real number in the domain $[0, 1)$.
3. The length scaling factor $f = (V_n/V_o)^{1/3}$ is calculated.
4. The partial Metropolis criterion is checked according to:

$$\text{acc}(o \rightarrow n) = \min \left[1, \exp \left[-\beta p(V_n - V_o) + N \ln \left(\frac{V_n}{V_o} \right) \right] \right],$$

checking whether we should reject this move before even needing to check for overlaps. If the move is rejected here, the function simply quits, having changed nothing. If it is accepted, it proceeds.

5. The simulation volume is scaled to the new volume: $L_{i,n} = L_{i,o}f$. All particle positions are scaled as well: $\mathbf{r}_n^N = \mathbf{r}_o^N f$.
6. If the volume has increased, the step is accepted here, as an isotropic increase in volume cannot lead to overlap. If it has decreased, all particles are checked for overlap in the new simulation volume with the new positions. If there is any overlap, the old configuration and volume are restored. Otherwise, it is accepted.

For crystal structures, however, this approach is flawed. If the initial relative dimensions of box are not commensurate with the crystal lattice, it will not fit correctly, and there will be latent stresses acting on the system preventing it from reaching its equilibrium state. In 1981, a paper by Parrinello and Rahman [51] offered a solution: define the box shape by a 3x3 tensor \mathbf{h} , and allow its elements to vary independently. This effectively makes for a simulation of the isothermal-isotension ensemble. While Parrinello and Rahman's original work was designed around an MD simulation, it has since been ported to MC simulations as well. Its implementation is relatively simple. Following the approach of [26, 52, 53], we implement a new MC move in addition to the standard NPT volume changing moves: one which deforms the simulation volume anisotropically. This MC move is done as follows:

Monte Carlo volume deformation move:

1. A copy of \mathbf{h} ; \mathbf{h}_o , is stored for later use.
2. A random element of \mathbf{h} , h_{ij} is chosen. Following [53], we choose \mathbf{h} to be symmetric in order to eliminate rotations of the entire system. We found that this improved stability of the algorithm.
3. A random deformation Δh is generated in the range $-\Delta h_{max} \leq \Delta h \leq \Delta h_{max}$, where Δh_{max} is allowed to vary during equilibration.
4. The volume is deformed by modifying h_{ij} and its symmetric counterpart h_{ji} ;

$$h_{ij} \rightarrow h_{ij} + \Delta h,$$

$$h_{ji} \rightarrow h_{ji} + \Delta h.$$

5. The partial Metropolis criterion is checked according to:

$$\text{acc}(o \rightarrow n) = \min \left[1, \exp \left[-\beta p(V_n - V_o) + N \ln \left(\frac{V_n}{V_o} \right) \right] \right],$$

where $V_n = \det \mathbf{h}$ and $V_o = \det \mathbf{h}_o$.

6. All particles are checked for overlap in the new box. If there is any overlap, the move is rejected and the box is restored to its old configuration; $\mathbf{h} = \mathbf{h}_o$. Otherwise, it is accepted.

2.5 Consistency check: comparison of equations of state for hard cubes

To check whether our MC code is working properly, we calculate the pressure-density equation of state of a system of hard cubes, and compare our results with those of Ref. [31]. There are two things we would like to see in particular: namely that the overall equation of state matches, and that we find evidence of vacancy formation and/or stabilization.

On the first few attempts to perform anisotropic NPT simulations on this system we observed that the system deformed extremely quickly if the deformation moves were allowed from the very start of the simulation. This was due to the fact that our simulations primarily consisted of melting runs for which the initial particle configuration was a perfect crystal lattice. For hard cubes, this crystal is energetically invariant to shearing of the lattice, and since all particles start out perfectly aligned, shearing the crystal lattice will hardly ever lead to overlaps in these first few cycles of the simulation. As such, the system was free to deform. The result of this was that the simulation volume became extremely distorted in the first few moments of simulation. The crystal, also severely deformed, occasionally melted even at very high pressures due to this distortion. Since any real system, especially one composed of hard and otherwise noninteracting particles, will never have such a perfectly orientationally aligned lattice at anything but the close packing limit, this behaviour is clearly unphysical. We fixed this issue by delaying the start of the volume deformations until the crystal has relaxed a little. Volume scaling moves were still allowed during this process. The number of MC cycles used for this partial equilibration was found by trial and error.

The equation of state shown in Fig. 3 was obtained from anisotropic MC-NPT simulations of a system of 1000 hard perfect cubes. We used $1 * 10^6$ Monte-Carlo cycles for equilibration, of which the first $3 * 10^5$ were done keeping the shape of the simulation volume fixed to prevent unphysical shearing of the crystal. A further $2 * 10^6$ Monte-Carlo cycles were used for sampling, where we sampled the density every 100 cycles, in order to obtain independent measurements.

In Fig. 3 we can see that the two equations of state are in near-perfect agreement with each other. A very small deviation can be seen if zoomed in on the coexistence region, where our MC results suggest a slightly lower coexistence pressure. However, Smallemburg et al.'s [31] equation of state is for a system of $N = 8000$ particles, whereas ours is for $N = 1000$ particles. Considering the additional size dependence of the system due to the vacancy delocalization, some deviations are to be expected here. Moreover, Smallemburg's equation of state was obtained from EDMD simulations in an NVT ensemble, while ours was obtained from MC simulations in an NPT ensemble. It is well known that different things will happen in the coexistence region, depending

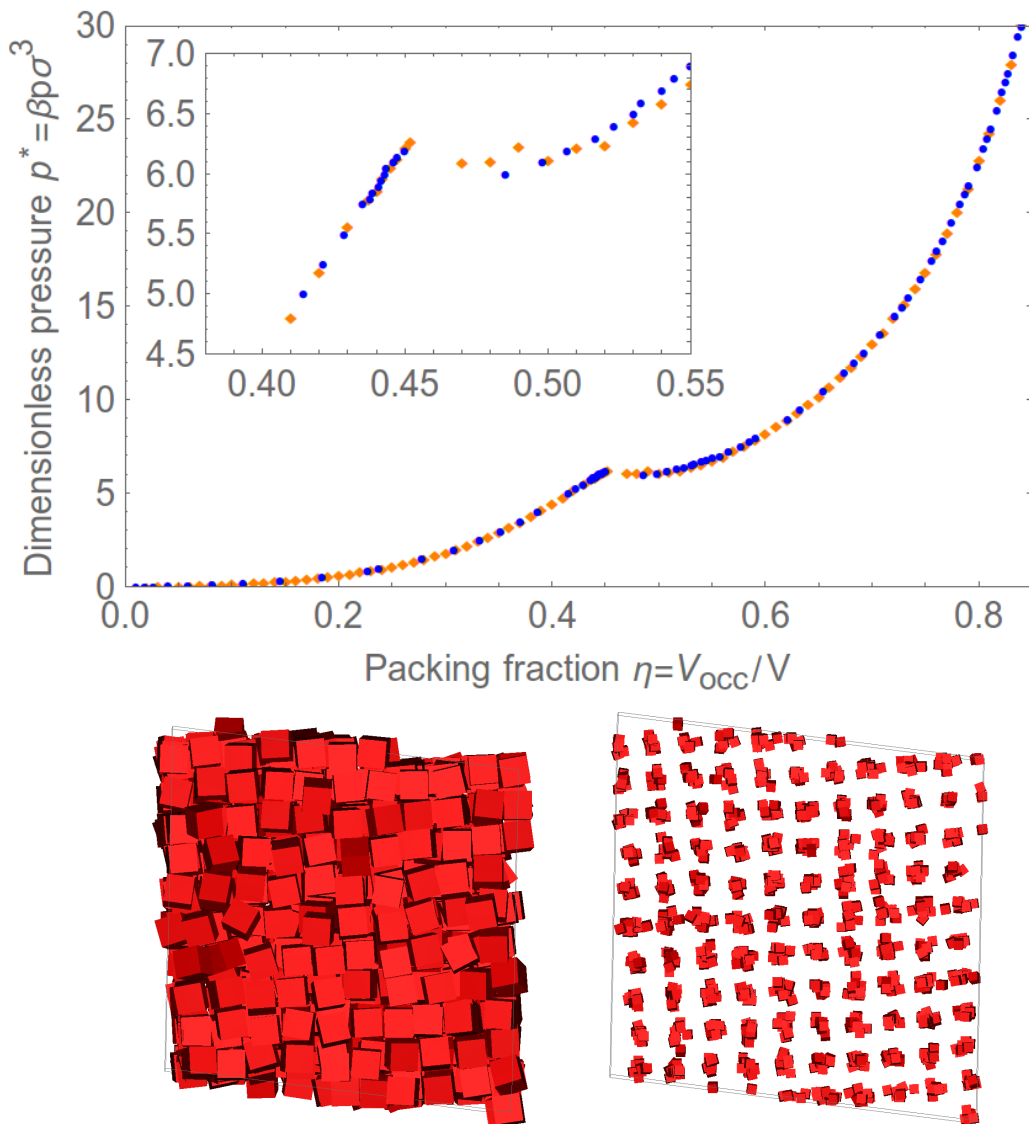


Figure 3: (top) Comparison between the equation of state found in our anisotropic NPT simulation (blue circles) and that found by Smalenburg et al. (orange diamonds) The inset is zoomed in on the fluid-solid coexistence region. (bottom-left) Final configuration of the $p^* = 6.80$ simulation. The crystal has rotated with respect to its initial configuration, something also observed in Refs. [23, 31] and found to be an indication of vacancy formation. (bottom-right) The same as bottom-right, but the particles have been shrunk visually to make the lattice more visible.

on which ensemble you use to simulate it. This behaviour follows from the Gibbs phase law:

$$F = C - P + 2 \Rightarrow P = C - F + 2, \quad (21)$$

where P is the number of phases in coexistence with each other, F the degrees of freedom of the system (the number of independent intensive thermodynamic variables) and C is the number of components. For a one-component system in the NPT ensemble, we have $C = 1$ and $F = 2$, so $P = 1$. Hence, a one-component system in an NPT ensemble cannot have any coexistence. This has both up- and downsides. On one hand, it makes the equations of state of the separate phases clearly distinct. One can see this in the equation of state, as it will “jump” from one phase to another. On the other hand, you will never be able to actually study a coexistence in the NPT ensemble. In contrast, a single-component system in the NVT ensemble is able to phase separate into two phases and you can actually observe a fluid-solid coexistence popping up in simulations in the coexistence region. Estimating the coexistence densities/pressures from the equation of state becomes somewhat trickier, however, as it will display what is known as a van der Waals loop, instead of a clearly visible jump. No matter the ensemble, however, there are better ways to derive coexistence properties, based around knowledge of the free energies.

Visual analysis of the system configuration at the end of the simulation reveal that all simulations of the solid close to the coexistence region have a crystal lattice that is rotated with respect to its initial configuration (Fig. 3, bottom). This behaviour was also first observed by Agarwal et al. [23], and later by Smalenburg et.

al. [31], who showed that this rotation was the result of the crystal trying to achieve its equilibrium vacancy concentration, but being unable to do so by forming additional layers due to the limited system size. To conclude, we find excellent agreement between our results and those of Smallegange et al., and can consider our anisotropic MC-NPT code to be working properly. We are now ready to study the phase behaviour of slanted cubes. But before we do so, we first take a look at another part of our toolkit: free-energy calculations.

3 Standard methods to compute the free energy of solids

One of the most common ways to find the free energy of an arbitrary thermodynamic system is by thermodynamic integration. In this method, we construct a path from a system for which the free energy can be determined exactly (called the reference state) to the system of which we would like to know the free energy. Notable examples of good reference systems (also named in Section 2.3) include the ideal gas and the Einstein crystal. In this chapter we will describe the theory and methods for computing the free energy of simple fluids and crystals. To start off, we review the basic theory of thermodynamic integration. From here, we move on to describe the *de facto* method to compute the free energy: the Frenkel-Ladd (or Einstein crystal) method. We will start by reviewing the original method proposed by Frenkel and Ladd in Ref. [40], after which we will discuss a modern variant on it. This latter method was introduced by Vega and Noya in Ref. [54], and its main attractions are the simpler implementation and the elegant way in which it deals with center-of-mass drift when the particles are only weakly bound to their lattice sites. Next, we address some more practical concerns on how to get the desired accuracy out of the aforementioned methods. To end this chapter and to check whether our implementation of the thermodynamic integration scheme is yielding the correct results, we calculate the free energy of a face-centered cubic crystal of hard spheres, and compare our values to literature values.

3.1 Thermodynamic integration

Let us first look at the simplest example of thermodynamic integration; measuring the free energy of a simple liquid. This can be done by integrating the pressure-density equation of state $p(\rho)$ over the density. As a reference system, we can use the ideal gas, for which the free energy is known.

$$\begin{aligned} F_{liq} &= F_{idgas} + \Delta F \\ \Rightarrow F(\rho) &= F_{idgas}(\rho = 0) + \int_0^\rho \frac{\partial F(\rho')}{\partial \rho'} d\rho'. \end{aligned} \quad (22)$$

For an ensemble with fixed temperature ($dT = 0$) and number of particles ($dN = 0$):

$$dF = -pdV \Rightarrow \frac{\partial(F/N)}{\partial \rho} = -p(\rho) \frac{\partial(V/N)}{\partial \rho} = -p(\rho) \frac{\partial(1/\rho)}{\partial \rho} = \frac{p(\rho)}{\rho^2} \Rightarrow \frac{\partial(\beta F/N)}{\partial \rho} = \frac{\beta p(\rho)}{\rho^2} \quad (23)$$

and so the free energy of a liquid can be calculated from the following expression:

$$\frac{F(\rho)}{Nk_B T} = \frac{F_{idgas}(\rho = 0)}{Nk_B T} + \frac{1}{Nk_B T} \int_0^\rho \frac{\beta p(\rho')}{\rho'^2} d\rho'. \quad (24)$$

In practice, we cannot actually integrate to $\rho = 0$ as the integrand will diverge. Instead, we use the fact that for the ideal gas we have $\beta p(\rho) = \rho$, and take the integration path:

$$\frac{F(\rho_2)}{Nk_B T} = \frac{F_{idgas}(\rho_1)}{Nk_B T} + \frac{1}{Nk_B T} \int_{\rho_1}^{\rho_2} \left(\frac{\beta p(\rho')}{\rho'^2} - \frac{1}{\rho'} \right) d\rho'. \quad (25)$$

In this example, we have illustrated thermodynamic integration over the density. However, the method is far more general than this. One can do the same for *any* thermodynamic variable, or even for systems with vastly different potentials. This latter application is known as Hamiltonian integration, and we will be making use of this method frequently in this thesis. The principle is simple [55]: one constructs a Hamiltonian that couples two systems with different potentials by means of a parameter λ , for which one limit of λ corresponds to the composite system interacting with only the potential of the reference system, and another limit of λ corresponding to the system interacting only with the potential of the system for which the free energy is unknown. For example, say we have a reference system with an energy $U_{ref}(\mathbf{r}^N)$ and another system with an energy $U_1(\mathbf{r}^N)$. A coupled Hamiltonian could then be:

$$U(\lambda, \mathbf{r}^N) = \lambda U_{ref}(\mathbf{r}^N) + (1 - \lambda) U_1(\mathbf{r}^N) \quad (26)$$

such that $U(1, \mathbf{r}^N) = U_{ref}(\mathbf{r}^N)$ and $U(0, \mathbf{r}^N) = U_1(\mathbf{r}^N)$. The free energy of this coupled system is then given by:

$$F(N, V, T, \lambda) = -k_B T \ln Z(N, V, T, \lambda) = -k_B T \ln \left[\frac{1}{N! \Lambda^{3N}} \int \exp(-\beta U(\lambda, \mathbf{r}^N)) d\mathbf{r}^N \right]. \quad (27)$$

Taking the derivative with respect to λ yields

$$\begin{aligned}
\frac{\partial F(N, V, T, \lambda)}{\partial \lambda} &= -k_B T \frac{\partial}{\partial \lambda} \ln \left[\frac{1}{N! \Lambda^{3N}} \int \exp(-\beta U(\lambda, \mathbf{r}^N)) d\mathbf{r}^N \right] \\
&= \frac{-k_B T}{Q(N, V, T, \lambda)} \left(\int \frac{\partial}{\partial \lambda} \exp(-\beta U(\lambda, \mathbf{r}^N)) d\mathbf{r}^N \right) \\
&= \frac{-k_B T}{Q(N, V, T, \lambda)} \left(\int (-\beta U(\lambda, \mathbf{r}^N)) \exp(-\beta U(\lambda, \mathbf{r}^N)) d\mathbf{r}^N \right) \\
&= \left\langle \frac{\partial U(\lambda, \mathbf{r}^N)}{\partial \lambda} \right\rangle_{N, V, T, \lambda}.
\end{aligned} \tag{28}$$

which means that

$$\begin{aligned}
F(\lambda = 0) &= F(\lambda = 1) + \int_1^0 \frac{\partial F(\lambda)}{\partial \lambda} d\lambda \\
&= F(\lambda = 1) - \int_0^1 \left\langle \frac{\partial U(\lambda, \mathbf{r}^N)}{\partial \lambda} \right\rangle_{N, V, T, \lambda} d\lambda \\
\Rightarrow F_1 &= F_{ref} - \int_0^1 \left\langle \frac{\partial U(\lambda, \mathbf{r}^N)}{\partial \lambda} \right\rangle_{N, V, T, \lambda} d\lambda.
\end{aligned} \tag{29}$$

This last step requires an important assumption: it assumes that the derivative $\partial F/\partial \lambda$ is continuous along the entire λ -integration path. Many systems exhibit first-order phase transitions in which there are discontinuities in the first derivative of the free energy. We cannot integrate over this transition. For fluids, we can almost always use the ideal gas as a reference system. But we usually cannot reach a crystal configuration from the ideal gas without crossing a first-order phase transition. Hence, for crystals we need a different reference system. This reference system has to be a system for which the free energy is known, which is already quite rare, and it has to be able to reach solid configurations without crossing any first-order phase transitions. Luckily, such a system exists; it is the *Einstein crystal*.

3.2 The Frenkel-Ladd (Einstein crystal) method

In 1984, Frenkel and Ladd [40] proposed a way to calculate the free energy of a solid. Their method was based around Hamiltonian integration using an Einstein crystal as a reference system. In an Einstein crystal, particles are bound to their lattice sites (\mathbf{r}_0^N) via a harmonic potential:

$$U_{Ein}(\lambda, \mathbf{r}^N) = \sum_{i=1}^N \lambda (\mathbf{r}_i - \mathbf{r}_0)^2. \tag{30}$$

The free energy of a non-interacting Einstein crystal is given by the analytical expression [40]

$$F_{Ein, id}(\lambda) = \Phi_0 - k_B T \ln \left[\frac{\pi k_B T}{\lambda} \right]^{3N/2} + C(T) + \mathcal{O}(1/\lambda), \tag{31}$$

with Φ_0 the potential energy of the static lattice, $C(T)$ the kinetic contribution to the free energy at a temperature T , and $\mathcal{O}(1/\lambda)$ all higher-order terms. To get from the free energy of an Einstein crystal to that of a crystal of hard spheres, they constructed a coupled potential that includes both the potential of the solid they want to get to as well as the potential of the Einstein crystal:

$$\mathcal{U}(\lambda, \mathbf{r}^N) = \mathcal{U}_{HS}(\mathbf{r}^N) + U_{Ein}(\lambda, \mathbf{r}^N), \tag{32}$$

This is a slightly different coupling than in Eq. 29. The reason for this is the hard sphere potential. One cannot gradually lower the contribution of a hard potential, like what was done in Eq. 29. Instead, the noninteracting limit was reached by binding the particles so tightly to their lattice sites (high enough λ) that they hardly ever got near enough for the hard potential to play a role. This means that the integration path was also slightly different:

$$F_{solid} = F(\lambda = 0) = F_{Ein, id}(\lambda) - \int_0^\lambda \left\langle \frac{\partial U(\lambda', \mathbf{r}^N)}{\partial \lambda'} \right\rangle_{N, V, T, \lambda'} d\lambda'. \tag{33}$$

The derivative $\partial U/\partial \lambda$ is simply the mean squared displacement of the particles from their lattice sites, and can easily be calculated in an MC-NVT simulation. Frenkel and Ladd used this method to calculate the free-energy difference between the hcp and fcc crystal phases of a system of hard spheres, two phases which are

very similar in both structure and free energy. However, the Frenkel-Ladd method has a drawback: for low coupling constants, the crystal can start to drift as a whole. This manifests itself as a quasi-divergence in the translational integrand as $\lambda \rightarrow 0$. Frenkel and Ladd solved this issue by fixing the center of mass of the system in their simulations. Fixing the center of mass of a system and correctly accounting for the free energy cost this incurs is, however, a somewhat arduous process. Recently, Vega and Noya [54] published a paper which describes a variation of the Frenkel-Ladd method where this quasi-divergence is cleverly avoided. It is this method that we will be using for our own free-energy calculations, and the next section will explain it in detail.

3.3 Einstein molecule method

Developed by Vega and Noya and published in 2007 [54], the Einstein molecule method is a variation on the Einstein crystal method with the added advantage of inherently accounting for the drift of the center of mass of the system. Vega and Noya consider a system that very closely resembles an Einstein crystal, but where the positions of all particles are described relative to one particular particle which they call the *lattice carrier*. Most importantly, the positions of the lattice sites are also defined relative to this lattice carrier. Of course, this is completely equivalent to the regular Einstein crystal. But whereas fixing the center of mass was a complicated process in the former, in the latter it is as easy as fixing the position of the lattice carrier (both theoretically and computationally). To see how this works, let us look at the free energies. First, the free energy of our reference system, the ideal Einstein molecule, is given by:

$$\frac{F_{Ein-mol,id}}{Nk_B T} = \frac{1}{N} \ln [\rho \Lambda^3] + \frac{3}{2} \left(1 - \frac{1}{N}\right) \ln \left[\frac{\beta \lambda}{\pi \Lambda^2}\right] \quad (34)$$

Fixing the center of mass, as done in Ref. [40] by Frenkel and Ladd, requires a free-energy cost of

$$F_{\text{fix CM}} = k_B T \ln [V/\Lambda^3], \quad (35)$$

while the free-energy cost of fixing the position of one particle, the lattice carrier, is given by [54, 55]:

$$F_{\text{fix lat.}} = k_B T \ln [V/\Lambda^3]. \quad (36)$$

As the astute reader will no doubt see, these are the same. Therefore, fixing the center of mass is numerically equivalent to fixing the position of one of the particles. Implementation-wise, however, the latter is significantly simpler; one simply fixes the position of one particular particle in the simulation by never allowing it to move. The translational energy can then be simplified to:

$$U_{Ein-mol,tr}(\lambda_{tr}, \mathbf{r}^{N-1}) = \sum_{i=1}^{N-1} \lambda_{tr} (\mathbf{r}_i - \mathbf{r}_0)^2. \quad (37)$$

The second reason this method is so convenient can be seen from Fig. 4. Modified from Ref. [55], this figure schematically represents the steps taken to obtain the free energy of a crystal for the Frenkel-Ladd and the Einstein molecule method. One can see that in the Einstein molecule method, the cost of fixing the position of the lattice carrier needs to be added in one step, then subtracted in another. The complicated method of accounting for CM drift in the Frenkel-Ladd method is now reduced to merely fixing the position of one of the particles in our simulation, and we lose no accuracy whatsoever.

Up until now, we have only looked at translational contributions to the free energy. But particles that have an anisotropy in their shape or in their potential also possess rotational degrees of freedom that contribute to the free energy. Despite the Einstein-crystal system not having any rotational degrees of freedom, we can still use it as a reference system as long as we ensure that the contribution of the rotational degrees of freedom vanishes in the Einstein crystal limit. In another paper by Noya and Vega [56], it is shown that if a system possess rotational degrees of freedom, it is necessary to allow the lattice-carrier particle to rotate in order to obtain the structural equivalence with a system with a fixed center of mass. The fact that all particles are allowed to rotate also means that the way one deals with orientational degrees of freedom is the same in both the Einstein crystal and the Einstein molecule approaches. The full integration path, including the orientational degrees of freedom, now becomes:

$$F_{solid} = F_{Ein-mol,id,tr}(\lambda_{tr}) + F_{Ein,id,or}(\lambda_{or}) - \int_0^{\lambda_{tr}} \left\langle \frac{\partial U_{tr}(\lambda'_{tr}, \mathbf{r}^{N-1})}{\partial \lambda'_{tr}} \right\rangle_{N,V,T,\lambda'_{tr}} d\lambda'_{tr} - \int_0^{\lambda_{or}} \left\langle \frac{\partial U_{or}(\lambda'_{or}, \mathbf{\Omega}^N)}{\partial \lambda'_{or}} \right\rangle_{N,V,T,\lambda'_{or}} d\lambda'_{or}. \quad (38)$$

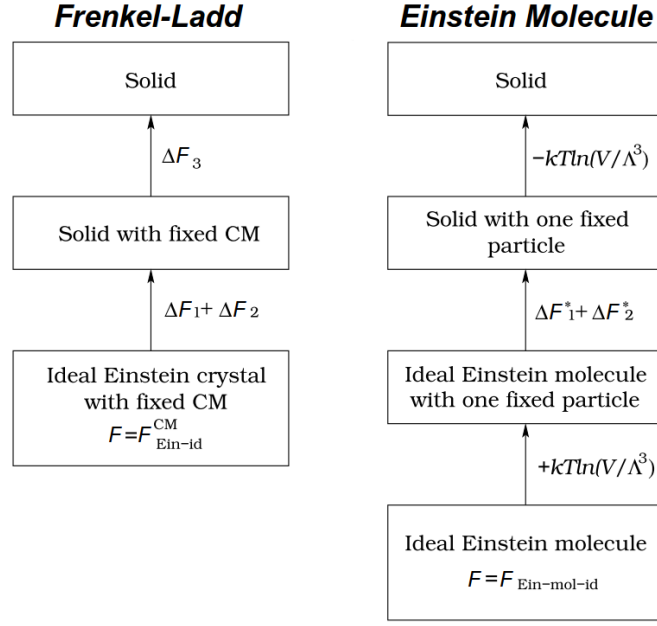


Figure 4: Image modified from Ref. [55]. Integration paths for the Frenkel-Ladd and Einstein molecule methods. ΔF_1 and ΔF_1^* indicate the free-energy difference between the high λ limit system and the true noninteracting Einstein crystal/molecule system, ΔF_2 and ΔF_2^* the free-energy difference resulting from the Hamiltonian integration (the integrals over λ_{tr} and λ_{or} from Eq. 38), and ΔF_3 indicates the free-energy difference resulting from fixing the center of mass in the Einstein crystal method.

The orientational free energy of the ideal Einstein crystal is then given by:

$$\begin{aligned}
 F_{\text{Ein,id,or}}(\lambda_{or}) &= -k_B T \ln \left[\int \exp \left[-\beta U_{\text{Ein,or}}(\lambda_{or}, \hat{\Omega}^N) \right] d\hat{\Omega}^N \right] \\
 &= -k_B T \ln \left[\left(\int \exp \left[-\beta u_{\text{Ein,or}}(\lambda_{or}, \hat{\Omega}) \right] d\hat{\Omega} \right)^N \right] \\
 &= -k_B T N \ln \left[\left(\int \exp \left[-\beta u_{\text{Ein,or}}(\lambda_{or}, \hat{\Omega}) \right] d\hat{\Omega} \right) \right],
 \end{aligned} \tag{39}$$

where $U_{\text{Ein,or}}$ is the orientational energy of the system gained by coupling to some orientational field. This is a one-body effect, and so we can split the integral over Ω^N into N integrals over the one-body orientational energy $u_{\text{Ein,or}}$. This one-particle integral can be evaluated separately from the rest of the calculations. As there are no pairwise or many-body interactions that need to be calculated, it is very cheap to evaluate numerically. The actual form of $U_{\text{Ein,or}}$ depends on the symmetry of the particles in question. In principle, any potential can be used as long as its thermodynamic integration path does not cross any first-order phase transitions. An absence of divergences is also desirable. Vega et al. [55, 56] have developed a number of potentials suitable for certain symmetry groups. So now would be a good time to check which symmetry groups our (slanted) cubes belong to.

The perfect cubes for which the vacancy-rich behaviour was first shown belong to point symmetry group O_h . For them, a suitable potential is:

$$U_{\text{Ein,or}}(\lambda_{or}) = \sum_{i=1}^N \lambda_{or} \sin(\psi_{a,\min})^2 + \lambda_{or} \sin(\psi_{b,\min})^2 \tag{40}$$

with $\psi_{a,\min}$ and $\psi_{b,\min}$ the minimum pair of angles between two separate faces ($\hat{\mathbf{u}}_j, \hat{\mathbf{u}}_k$) and a predefined $\hat{\mathbf{x}}$ or $\hat{\mathbf{y}}$ axis. We can simplify this expression to:

$$U_{\text{Ein,or}}(\lambda_{or}) = \sum_{i=1}^N \lambda_{or} \min_{j \neq k} \left(2 - (\hat{\mathbf{u}}_{i,j} \cdot \hat{\mathbf{x}})^2 + (\hat{\mathbf{u}}_{i,k} \cdot \hat{\mathbf{y}})^2 \right) \tag{41}$$

The slanted cubes belong to a different symmetry group, D_{2h} . For them, we use the potential:

$$U_{\text{Ein,or}}(\lambda_{or}) = \sum_{i=1}^N \lambda_{or} \sin(\psi_a)^2 + \lambda_{or} \sin(\psi_b)^2 \tag{42}$$

where ψ_a and ψ_b are now the angles between respectively the surface normals $\hat{\mathbf{u}}_y$ and $\hat{\mathbf{u}}_{xz}$ of the particles and their equilibrium values $\hat{\mathbf{u}}_{y,0}$ and $\hat{\mathbf{u}}_{xz,0}$, which correspond to the orientation the particles would have in the perfect crystal lattice. For a visualization, see Fig. 5. The only requirement for the two chosen unit vectors is that they are orthogonal, so that the particle orientation is determined uniquely by fixing them.




Figure 5: A slanted cube where we have labeled its surface normals $\hat{\mathbf{u}}_{xz}$ and $\hat{\mathbf{u}}_y$, the equilibrium orientation of the surface normals $\hat{\mathbf{u}}_{xz,0}$ and $\hat{\mathbf{u}}_{y,0}$ and the angles ψ_a and ψ_b between them. For a strong orientational coupling (high λ_{or}), the slanted cubes will be aligned according to their orientations in the perfect crystal lattice.

3.4 Method: Free energy of solids

To evaluate the integrals from Eq. 38, we use a Gauss-Legendre quadrature⁵, where we have transformed them in the same manner as Frenkel et al. in [40, 42];

$$\Delta F = - \int_0^\lambda \langle \partial U / \partial \lambda \rangle d\lambda = - \int_{\ln[c]}^{\ln[\lambda+c]} \langle \partial U / \partial \lambda \rangle \cdot (\lambda + c) d \ln [\lambda + c], \quad (43)$$

where $c \approx k_B T \sigma^2 / \langle (\mathbf{r} - \mathbf{r}_0)^2 \rangle_{\lambda=0}$. The integrand at each value of λ will be evaluated using the Einstein Molecule (EM) method. We implement this method by constructing an MC simulation in an NVT λ ensemble using the potential

$$U(\lambda_{tr}, \lambda_{or}, \mathbf{r}^N) = U_{HC}(\mathbf{r}^N) + U_{Ein,tr}(\lambda_{tr}, \mathbf{r}^N) + U_{Ein,or}(\lambda_{or}, \mathbf{\Omega}^N). \quad (44)$$

Here, U_{HC} is the hard-core repulsive potential, $U_{Ein,tr}(\lambda_{tr})$ is the translation Einstein molecule potential as defined in Eq. 37, and $U_{Ein,or}(\lambda_{or})$ is the orientational Einstein molecule potential as defined in Eq. 41 or 42, depending on the lattice symmetry. Note that the hard-core potential is independent of coupling constants. This is necessary for hard particles, as the hard-core potential is discontinuous and we cannot vary it smoothly. The consequences of this modification will be discussed later in this section. We choose the equilibrium positions and orientations of the particles such that they correspond with those of the ideal lattice. We modify the Metropolis criterion for every MC move to also include the Einstein crystal potentials, and evaluate $U(\lambda_{tr}, \lambda_{or}, \mathbf{r}^N)$ to obtain the averages $\langle \partial U_{tr} / \partial \lambda_{tr} \rangle$ and $\langle \partial U_{or} / \partial \lambda_{or} \rangle$.

But which values of λ_{tr} and λ_{or} should we choose? And how many integration points do we need for a reasonable accuracy? We have a couple of criteria that the integration needs to meet: one end of the integration should be the no coupling limit, the other the ideal Einstein crystal limit, and there need to be enough points in between to capture the curvature of $\langle \partial U / \partial \lambda \rangle$. Let us go through this list in order. The first point seems easy. At $\lambda = 0$, the Einstein crystal part of the potential is zero and there is no coupling. But we do not use $\lambda = 0$ as an integration point using a Gauss-Legendre quadrature. Instead, we should check that whichever lowest value of λ we *do* use actually approximates zero coupling well enough. We can do this by checking whether the integrands have become constant in λ . See Fig. 6. At some point, the coupling is so weak that the system has effectively only hard interactions.

Next, let us see whether the ideal Einstein crystal limit is reached. Going back to Eq. 44, we see that the hard-core potential is independent of the coupling constant. So even in the $\lambda \rightarrow \infty$ limit, the hard-core potential is still there. But while it seems odd to call this a noninteracting limit, in reality the particles will be so tightly bound to their ideal positions and orientations that the hard-core potential hardly matters anymore. Of course, we can only approach this limit asymptotically. We will need to truncate this integral at some point. But where? Well, at whichever point brings us as close to the noninteracting limit as we desire. The Einstein

⁵A Gauss-Legendre quadrature is a choice of integration points which provides very good results if the function you want to integrate can be described (or well approximated) by a polynomial.




Figure 6: Values for $(1/N) \langle \partial U_{or} / \partial \lambda_{or}^* \rangle$ for various λ_{or}^* . Dots indicate simulated data points from the coupled system. The solid blue line indicates Eq. 51; the value for the ideal Einstein molecule. We can see that for high λ_{or}^* we approach the noninteracting limit and for low λ_{or}^* the function goes towards a constant.

crystal is an analytically solvable system, and so we can just derive $(\partial U / \partial \lambda)_{N,V,T}$. Fixing the position of one of the particles, we start from the partition function:

$$\begin{aligned} Z(N, V, T, \lambda_{tr}, \lambda_{or}) &= \frac{1}{N! \Lambda^{3N}} \int d\mathbf{r}_1 \int \exp[-\beta U_{tr}(\lambda_{tr}, \mathbf{r}^{N-1})] d\mathbf{r}^{N-1} \int \exp[-\beta U_{or}(\lambda_{or}, \mathbf{\Omega}^N)] d\mathbf{\Omega}^N \\ &= \frac{V}{N! \Lambda^{3N}} Q_{tr}(\lambda_{tr}, \mathbf{r}^{N-1}) Q_{or}(\lambda_{or}, \mathbf{\Omega}^N). \end{aligned} \quad (45)$$

We will start at the translational integrand, which is just the mean squared displacement:

$$\left\langle \frac{\partial U_{tr}}{\partial (\beta \lambda_{tr})} \right\rangle = \left\langle (\mathbf{r} - \mathbf{r}_0)^2 \right\rangle = \frac{1}{Q_{tr}} \int (\mathbf{r} - \mathbf{r}_0)^2 \exp[-\beta \lambda_{tr} (\mathbf{r} - \mathbf{r}_0)^2] d\mathbf{r} = -\frac{1}{Q_{tr}} \frac{\partial Q_{tr}}{\partial (\beta \lambda_{tr})} \quad (46)$$

To find Q_{tr} and its derivative, we use that $U_{tr} = \sum_{i=2}^N \lambda_{tr} (\mathbf{r}_i - \mathbf{r}_{i,0})^2$ is a quadratic term, and so we can use the Gaussian integral $\int_{-\infty}^{\infty} \exp[-ax^2] = \sqrt{\pi/a}$:

$$\begin{aligned} Q_{tr} &= \int \exp \left[-\beta \lambda_{tr} \sum_{i=2}^N (\mathbf{r}_i - \mathbf{r}_{i,0})^2 \right] d\mathbf{r}^{N-1} \\ &= \left(\int \exp[-\beta \lambda_{tr} (x_i - x_{i,0})^2] dx \right)^{3(N-1)} \\ &= \left(\sqrt{\frac{\pi}{\beta \lambda_{tr}}} \right)^{3(N-1)} \\ &= \left(\frac{\pi}{\beta \lambda_{tr}} \right)^{3(N-1)/2}, \end{aligned} \quad (47)$$

$$\begin{aligned} \frac{\partial Q_{tr}}{\partial (\beta \lambda_{tr})} &= \frac{\partial}{\partial (\beta \lambda_{tr})} \left(\frac{\pi}{\beta \lambda_{tr}} \right)^{3(N-1)/2} \\ &= \left(\frac{3(N-1)}{2} \left(\frac{\pi}{\beta \lambda_{tr}} \right)^{3(N-1)/2-1} \right) \left(-\frac{\pi}{(\beta \lambda_{tr})^2} \right) \\ &= -\frac{3(N-1)}{2} \pi^{3(N-1)/2} \left(\frac{1}{\beta \lambda_{tr}} \right)^{3(N-1)/2+1}. \end{aligned} \quad (48)$$

The mean squared displacement then becomes:

$$\begin{aligned} \left\langle (\mathbf{r} - \mathbf{r}_0)^2 \right\rangle &= -\frac{1}{Q_{tr}} \frac{\partial Q_{tr}}{\partial (\beta \lambda_{tr})} \\ &= -\left(\left(\frac{\pi}{\beta \lambda_{tr}} \right)^{-3(N-1)/2} \right) \left(-\frac{3(N-1)}{2} \pi^{3(N-1)/2} \left(\frac{1}{\beta \lambda_{tr}} \right)^{3(N-1)/2+1} \right) \\ &= \frac{3(N-1)}{2} \frac{1}{\beta \lambda_{tr}}. \end{aligned} \quad (49)$$

Converting this to a dimensionless per-particle number:

$$\begin{aligned}\sigma^2 \langle (\mathbf{r}' - \mathbf{r}'_0)^2 \rangle &= \frac{1}{N} \frac{3(N-1)}{2} \frac{1}{\beta \lambda_{tr}} \\ \Rightarrow \langle (\mathbf{r}' - \mathbf{r}'_0)^2 \rangle &= \frac{3(N-1)}{2N} \frac{1}{\lambda_{tr}^*}\end{aligned}\tag{50}$$

Where $\lambda^* = \sigma^2 \beta \lambda$ is the reduced translational coupling constant. A nearly analogous derivation yields the orientational expression:

$$\frac{\partial U_{or}}{\partial (\beta \lambda_{or})} = \frac{3}{2} \frac{1}{\lambda_{or}^*}\tag{51}$$

We can now decide which values of λ_{tr} and λ_{or} we need by looking at the difference between the integrands we calculate numerically and the ideal expression. The number of integration points and the value of the constant c can then be chosen by looking at the low coupling limit.

3.5 Consistency check: Free energy of a hard sphere FCC crystal

Before we start calculating the free energies of the phases of our slanted cube system, it is useful to first check whether our code yields free energies that are consistent with literature values. Our first check will be to calculate the free energy of an FCC crystal of hard spheres. The question of the relative stability of the FCC and HCP crystal phases of hard spheres was a longstanding point of discussion, and literature has a number of exceptionally accurate free energies for us to use as reference. We are only evaluating Eq. 43 for the Einstein molecule system, though, and this specific term is only found in Ref. [54].

	N	λ_{max}^*	ΔF_2
Ref. [54]	256	632.026	-3.0116
PW	256	632.026	-3.0120
Ref. [54]	1372	1000	-3.6862
PW	1372	1000	-3.6847

Table 1: Comparison of the free energies of an FCC crystal of hard spheres at a density of $\rho^* = 1.04086$. The free energies in the present work (PW) were obtained using the Einstein molecule method with a 20 point Gauss-Legendre quadrature.

As can be seen from Table 1, we find the same free energy up to a difference of $\sim 0.002 k_B T/N$. For reference, the free-energy difference between the FCC and HCP phases of hard spheres is on the order of $0.001 k_B T/N$ [57], and Vega et al. note an uncertainty of $0.002 k_B T/N$ in their final estimates for the free energies [54]. So we can consider this check passed. Thus, our toolkit is completed. The next chapter will delve into the phase behaviour of slanted cubes, studying the various crystal phases that this system exhibits using both the isotension-isothermal Monte Carlo simulations outlined in Chapter 2 as well as the free-energy calculations outlined in this chapter.

4 Crystal phases of slanted cubes

In this chapter we use the isotension-isothermal Monte Carlo simulations described in Section 2.4 to explore the phase behavior of hard slanted cubes, which we have described in Section 1.1. The main focus here will be on the crystal phases we find. We calculate pressure-density equations of state for multiple slant angles ranging from 45° to 90° , and characterize the different phases found by looking at their positional and orientational ordering. Using the free-energy calculations described in the previous chapter, we calculate the free energies of various high-density crystal candidates and find that a sheared cubic crystal lattice is the stable high-density crystal structure for both $\theta = 72^\circ$ as well as $\theta = 60^\circ$ slanted particles. From the equations of state we calculate and knowledge of the equilibrium high-density crystal phase we can draw a rough phase diagram in the $\theta - \eta$ plane, marking the regions of slant angles and densities for which certain phases are stable. Accurate determination of the phase boundaries from a MC-NPT simulation alone is hindered by hysteresis, slow nucleation and finite size effects, so we will only use the apparant phase boundaries from our simulations here as rough estimates. In Section 6.2 we will use further free-energy calculations to refine these estimates.

4.1 Possible close packed crystals

Before we can begin, we have to select an initial configuration to start our simulations with. Typically, one explores the pressure regime by simulating separate systems for a range of pressures starting from different initial configurations, usually a low-density fluid and/or a high-density crystal. A low-density fluid can be generated independent of whichever shape the particles have and will equilibrate quickly, but the crystal configuration has to be picked carefully as the equilibration is slow. For the high-pressure regime, there are two arguments that will help us in our choice. First is the close packing argument. At high pressures, the equilibrium crystal configuration must be the one with the highest packing fraction or density. Since slanted cubes can form space-filling lattices, the high-density crystal lattice must be a space-filling one. The second argument is similarity to cubes. At least for slant angles θ close to 90° , the crystal lattice will likely be simple cubic or similar to it. Based on these two arguments, we identify two possible space-filling lattice candidates (Fig. 7): Sheared Cubic (ShC) and Alternating Cubic (AC), which both reduce to a space-filling simple cubic lattice for $\theta = 90^\circ$. For $\theta \neq 90^\circ$, the cubic symmetry of the particles is broken and the particles must align according to a different symmetry to form a space-filling lattice. The ShC lattice consists of layers of one of the two possible 2D rhombus square tilings. Alternatively, it can be thought of as a simple cubic lattice that is sheared along one direction in the same way the particles are. The AC lattice is a simple cubic lattice with respect to the positions, but where each layer of particles is rotated 180° around their x-axis. It corresponds to the other possible 2D rhombus square tiling. We expect the ShC lattice to be the most stable one. In the ShC lattice, particles could theoretically move collectively along rows in all three dimensions, leading to an increase in entropy and also allowing for the possibility of delocalized vacancies. In the AC lattice, particles are trapped along one axis. Later, in Section 4.3, we will do free-energy calculations to test this hypothesis. We expect the the ShC lattice to be the more stable one. In Ref. [58], Paik et al. report that they observe rhombic platelets (with dimensions corresponding to a slant angle of 71°) forming a structure that very closely resembles the ShC lattice. One additional note has to be made: the ShC lattice is only commensurate with a square simulation volume for a certain set of slant angles. These angles are

$$n_z \cos \theta = n \rightarrow \theta = \arccos \left(\frac{n}{n_z} \right),$$

where n is a positive integer, θ is the slant angle and n_z is the number of lattice sites in the z -direction under our definition. To simplify our simulations, we will restrict ourselves to slant angles that meet this criterion. For reference, these angles are listed in degrees in Table 2. In the text, we will often refer to slant angles without listing the decimals, but they are always included in the calculations.

θ	90.000°	84.261°	78.463°	72.542°	66.422°	60.000°	53.130°	45.573°	36.870°	25.842°
----------	----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------

Table 2: Slant angles for which a $10 \times 10 \times 10$ periodic ShC lattice is commensurate with a square simulation volume.

4.2 Isotension Monte Carlo simulation results

To get a general idea of the phase behaviour of this system, we will look at the $p(\eta)$ equations of state and visualizations of the particle configurations. We obtain these by performing anisotropic NPT simulations on a system of $N = 1000$ particles at a range of reduced pressures between $\beta p \sigma^3 = 0$ and $\beta p \sigma^3 = 30$. As the initial configuration we use the ShC lattice. To avoid immediate extreme shearing of the system as discussed in Section 2.5, we only allow volume deformation move after a equilibration period of $> 10^3$ MC cycles. A longer




Figure 7: Possible close packed crystal lattices. (Left) The ShC lattice, created by shearing a cubic lattice similar to the particle shape. This lattice should in theory allow for delocalized vacancies much like the SC lattice. (Right) The AC lattice, formed by alternating the orientation of the particles on a cubic lattice.

period of $3 \cdot 10^5$ MC cycles was given for the system to equilibrate before sampling. Afterwards, for $2 \cdot 10^6$ MC cycles the density was sampled every 100 cycles. For slant angles $66^\circ \leq \theta \leq 90^\circ$, the equations of state are shown in Fig. 8. As expected and consistent with Ref. [31] we can see only two branches for $\theta = 90^\circ$: a fluid branch (F) and a simple cubic branch (SC). However, for the slanted cubes we observe something interesting: at intermediate densities, the slanted cubes form an SC lattice instead of the predicted ShC lattice. This suggests that for slanted cubes, the low density crystal phase is also SC, and not ShC. The SC/ShC phase transition also appears to be first order, showing a gap in the equation of state.

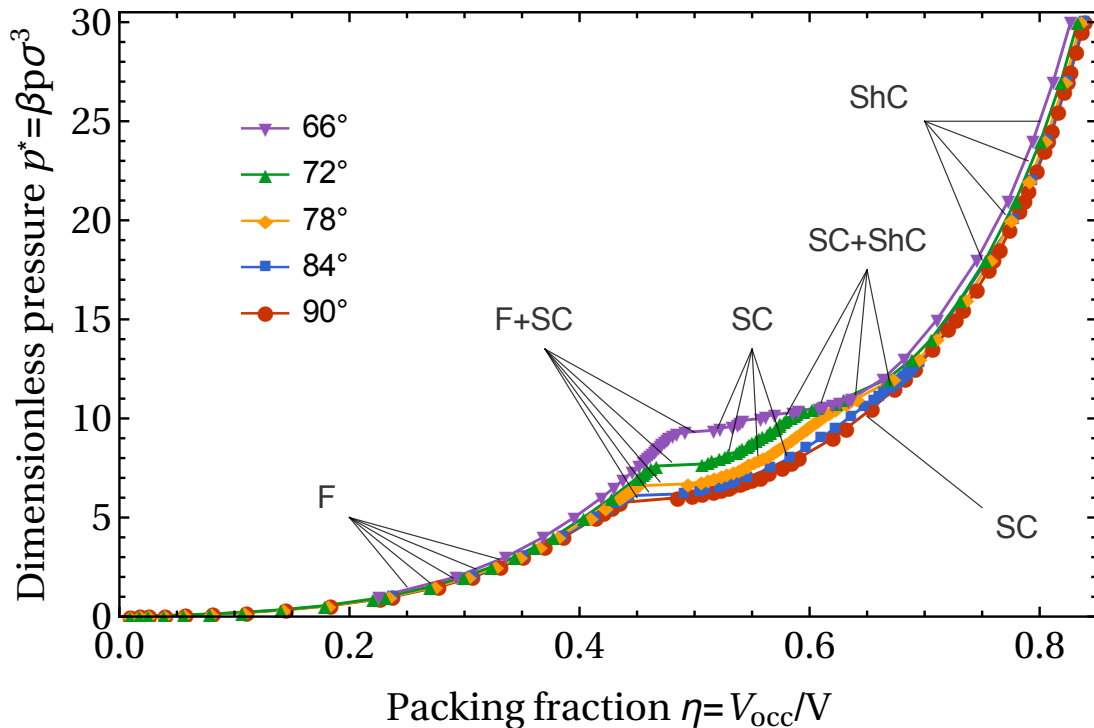


Figure 8: Equations of state for high slant angle particles obtained from MC-NPT simulations starting from a ShC crystal. Indicated are the sheared cubic (ShC), simple cubic (SC) and fluid (F) phases, along with their coexistence regions.

To gain a better understanding of what is happening, let us take a closer look at these phases. In Fig. 3, we show snapshots and orientation scatterplots for the fluid, simple cubic and sheared cubic phases. The orientation scatterplots show the \hat{x} (red), \hat{y} (green), and \hat{z} (blue) projections of the orientation matrix on the unit sphere, which are defined in the reference system of the particle (see Fig. 1).




Table 3: Phases of slanted cubes. Top row: snapshots with particles depicted their true size. Middle row: particles shrunk to a fraction of their true size to show the underlying crystal lattice. Bottom: scatterplots of the particle orientations, where the red, green and blue colors represent the inner products of the orientation matrix of the particles with respectively the \hat{x} , \hat{y} and \hat{z} unit vectors in the reference frame of the particle. The left column shows the fluid phase, the middle column the cubic phase, and the right column the slanted phase.

On the left in Fig. 3, we see the fluid phase. Clearly, there is very little ordering in either the positions or the orientations. The middle column shows the cubic phase. For crystal structures with densities close to coexistence ($\eta \sim 0.55$), we observe something peculiar: while the crystal structure appears to be simple cubic, the simulation volume is not cubic. If we repeat these simulations, choosing the initial condition to be an SC lattice that is aligned with the cubic simulation volume, we observe the same thing (see Fig. 9). The system appears to prefer a non-cubic simulation volume and a rotated simple cubic lattice over the simple cubic lattice that is commensurate with the simulation volume. This is the same behaviour as found for cubes, and it hints that for slanted cubes, too, the SC lattice has a high concentration of vacancies at equilibrium. Later work by Berend van der Meer (unpublished as of yet) shows that the vacancies in this system can delocalize over as many as 20 lattice sites, suggesting that additional finite-size effects may play a role in this “rotated crystal” behaviour for system sizes on the order of 10^3 particles.




Figure 9: (Left) Initial configuration: an SC lattice of slanted cubes with the appropriate orientation distribution. (Middle) Final configuration after $2 \cdot 10^6$ MC cycles. While the lattice is still SC, it has rotated with respect to its initial configuration. The final simulation volume shape is also non-cubic. (Right) Same as middle, but with the particles shrunk to a fraction of their real size, showing the simple cubic ordering.

The third phase this system exhibits is the ShC phase, shown in the right column of Fig. 3. Whether this is actually the equilibrium phase is not something we can conclude from just NPT simulations. The only thing we can really conclude is that it appears to be very stable up to the timescales we have simulated. However, crystal dynamics are notoriously slow, so this can hardly be considered strong evidence for actual stability. Additional NPT simulations of the AC phase offer no resolution. As shown in Fig. 10, the equation of state is identical to that of the ShC lattice to within our accuracy. Like ShC, it too is stable over the timescale of simulation, and it appears to exhibit a first order phase transition to/from the SC phase. The stability of both phases over the simulated timescale leads us to believe that they are nearly degenerate with respect to the free energy. In a sense, this is reminiscent of the difference between the FCC and HCP phases of hard spheres. In Section 4.3 we will do free-energy calculations on both crystal lattices and see that this is indeed the case, with some slight evidence pointing towards the ShC crystal being the preferred phase.

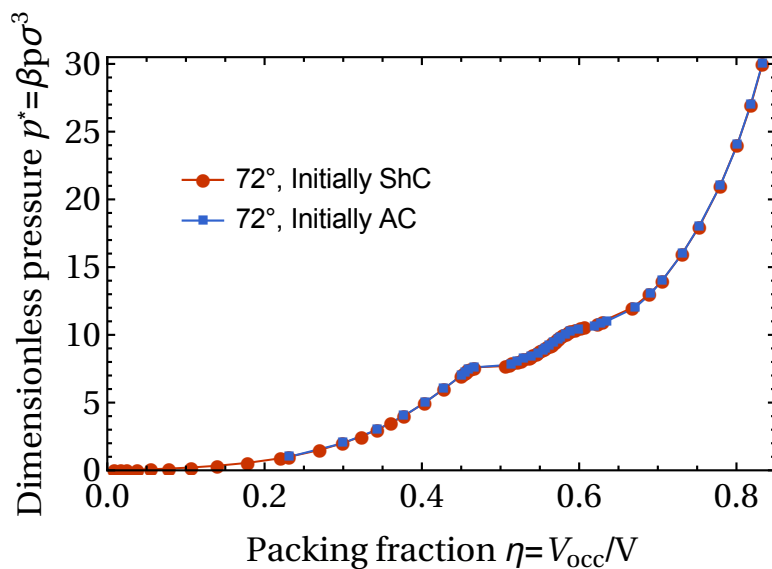


Figure 10: Equations of state for 72° slanted particles obtained from MC-NPT simulations starting from either a sheared cubic (ShC) crystal or an alternating cubic (AC) crystal. Data points are joined by lines to guide the eye. The two equations of state are virtually indistinguishable, even the locations of the phase transitions.

At smaller slant angles, the region where the SC crystal is stable shrinks. At a slant angle of 60° , we no longer find a simple cubic crystal. Instead, we find a phase that resembles layers of hexagonal tiling (Fig. 11). We first checked whether this phase was not simply a stage of melting to a fluid by looking at the orientation distribution. This is shown in Fig. 12. The orientation distribution quickly changes to a 6-fold symmetric one in-plane, which becomes more pronounced as a function of time. This is evidence against the notion that the crystal is simply melting, as that would require the orientations to become more isotropic, not less.




Figure 11: The random hexagonal tiling phase found for $\theta = 60^\circ$ at a packing fraction of $\eta = 0.60$. On the left are snapshots obtained from MC-NPT simulations, on the right from EDMD simulations. Both a ShC and a perfect rhombille lattice melt into a hexagonal random tiling structure.




Figure 12: Time evolution of the orientation distribution during a simulation on a $\theta = 60^\circ$ system. On the left Fig. the initial configuration, in the middle a snapshot of the system $\sim 5 \cdot 10^4$ MD timesteps later, and on the right the distribution at the end of the simulation, after $\sim 5 \cdot 10^5$ MD timesteps. The distribution quickly goes to a layered 6-fold symmetric one, and narrows from there, suggesting stability of the HRT phase.

To check whether it was not an intermediate between two crystal phases, we changed the initial configuration to a crystal composed of layers of rhombille tiling, a crystal configuration that is only space-filling for $\theta = 60^\circ$, and repeated our anisotropic MC-NPT simulations. To our surprise, these simulations also ended up with an aperiodic lattice of layers of 2D hexagonal tiling. It appears that for a certain density range, the stable phase of the system is a crystal composed of layers of 2D hexagonal random tilings. The closest analogue we can find is the 2D ($\sim 60^\circ$) rhombic tiling as studied by Whitelam et al. [59]. But while they study a two-dimensional system that is enthalpically stabilized by means of attractive patches, ours is three-dimensional and purely entropic, resulting from solely hard-core repulsions. Regarding the region of stability for this random tiling phase; following Ref. [59], we expect this phase to be stable for a small region surrounding $\theta = 60^\circ$. While Whitelam et al. did not report the density dependence of the stability of their random tiling phase, they suggest that it is stable for $56^\circ \leq \theta \leq 64^\circ$. This is consistent with our results, as we have seen no evidence of the random tiling phase being stable anywhere for $\theta = 53^\circ$ or $\theta = 66^\circ$. Theoretically, it could extend upwards to $\eta = 1$ as it is a space-filling configuration. However, this will only be the case for a slant angle of *exactly* 60° . The inability of random tiling to fill space for $\theta \neq 60^\circ$ prohibits its stability at packing fractions nearing $\eta = 1$. Based on the jump in symmetry, we expect the ShC/HRT phase transition to be first-order. Simulations of the AC, ShC and HRT phase for $\theta = 60^\circ$ yield equations of state that are indistinguishable, at least relative to the accuracy of our simulation. No jump or cusp in the equation of state is observed when an ShC lattice melts into an HRT lattice. While this phase is certainly interesting enough to warrant further investigation, it does not really fall within the scope of this research, and so we will leave it for others to explore. Instead, we will focus our attention on the vacancy-rich simple cubic phase.

For slant angles $\theta \approx 45^\circ$ and $\theta \approx 53^\circ$ we observed no intermediate phase between the ShC crystal and the fluid. Snapshots from simulations in the coexistence region (Fig. 13) show regions of ShC-like structure coexisting with a fluid phase, lending evidence to the statement that there is indeed no intermediate phase.




Figure 13: Fluid-solid coexistence for $\theta = 53^\circ$. The solid phase is quite clearly not simple cubic, but whether it is a very defective ShC, AC or hexagonal random tiling is hard to discern. These snapshots were obtained from long EDMD simulations, using a ShC crystal as initial configuration. The coexistence forms quickly, and remains for long timescales.

4.3 Free energies of the sheared cubic, alternating cubic and hexagonal lattices

In order to conclusively answer the question of the relative stabilities of the various crystal phases we found, we use the Einstein molecule method as described in Section 3.3 to calculate the free energies of these crystals for particles with slant angles $\theta = 72^\circ$ and $\theta = 60^\circ$. For the particles with slant angle $\theta = 72^\circ$, we calculate the free energies of the sheared cubic and of the alternating cubic lattices for two packing fractions; $\eta = 0.75$ and $\eta = 0.90$. For the particles with $\theta = 60^\circ$ we calculate the free energy of the sheared cubic and of the layered perfect rhombille crystal shown on the far left in Fig. 11 in Section 4.2. The results are summarized in Table 4. For the sake of completeness we list all contributions separately. The crystal with the lowest free energy is the equilibrium structure. What we can see from Table 4 is that the sheared cubic lattice appears to be the equilibrium structure for $\theta = 72^\circ$, though the difference in free energy with the alternating cubic crystal is small ($\sim 0.01 k_B T/N$). For the $\theta = 60^\circ$ particles, we can see that the difference between the two perfect lattices is even smaller ($0.004 k_B T/N$). The presence of the random tiling phase at lower densities then prompts us to ask the question: would a “mix” of these two phases (i.e. the random tiling phase) not be a more stable configuration due to the gain in combinatorial entropy? The fact that there are many more possible random tilings leads to an increase in the number of microstates that constitute the random tiling macrostate, leading to an increase in entropy. If the free volume gain due to ordering perfectly into either crystal structure is low, then this combinatorial entropy could cause the equilibrium phase to be one that is randomly tiling, even at high densities. We suspect this is the case.

	η	λ^*	$f_{Ein-mol,id,tr}$	$f_{Ein,id,or}$	$\Delta f_{2,tr}$	$\Delta f_{2,or}$	f_{solid}
Sheared cubic (72°)	0.90	8000	11.7519	15.094(2)	-2.270(10)	-2.564(5)	22.011(11)
Alternating cubic (72°)	0.90	8000	11.7519	15.094(2)	-2.234(10)	-2.577(5)	22.035(11)
$f_{ShC} - f_{AC}$							-0.024(16)
Sheared cubic (72°)	0.75	4000	10.7130	14.058(5)	-5.043(5)	-5.501(2)	14.228(7)
Alternating cubic (72°)	0.75	4000	10.7130	14.058(5)	-5.012(5)	-5.512(2)	14.237(7)
$f_{ShC} - f_{AC}$							-0.011(10)
Sheared cubic (60°)	0.75	4000	10.7131	-	-6.634(1)	-6.991(1)	$-2.912(1) + f_{Ein,id,or}$
Rhombille (60°)	0.75	4000	10.7119	-	-6.585(1)	-7.036(1)	$-2.908(1) + f_{Ein,id,or}$
$f_{ShC} - f_{rhomb}$							-0.004(2)

Table 4: Reduced free energies ($f = F/Nk_B T$) of the ShC, AC and rhombille crystal structures for various slant angles and packing fractions. For all crystals listed here $N = 1000$. Note that $f_{Ein,id,or}$, the orientational free energy per particle of the ideal Einstein crystal, is independent of the crystal structure or the number of particles. Therefore, its uncertainty does not contribute to the differences in free energy that are listed.

4.4 Phase diagram of slanted cubes

Based on the equations of state we have calculated and the crystal structures we have seen, we can construct a $\theta - \eta$ phase diagram, in which we mark the regions of stability and coexistence of the various phases. We can look at the gaps in the equations of state to determine roughly where the phase boundaries are for various slant angles, and in this way obtain an estimate for the coexistence densities. The resulting phase diagram is shown in Fig. 14. We can see three crystal domains. For all slant angles except $\theta = 60^\circ$, the sheared cubic lattice is the equilibrium high density crystal lattice. At $\theta = 60^\circ$, a random tiling phase is found at intermediate densities, but from both literature (Ref. [59]) and our own simulations we expect it to be stable in a small region near $\theta = 60^\circ$ up to close packing as well. The highest density for which this phase formed spontaneously from a sheared cubic crystal is indicated by the star, and the dashed lines indicate our guess for the regions of stability near $\theta = 60^\circ$. The simple cubic phase forms the third crystal phase of our phase diagram. From being the only crystal phase encountered at $\theta = 90^\circ$, its region of stability shrinks with a decreasing slant angle, though it remains the equilibrium phase at intermediate densities even for particles that are fairly acubic. The sheared cubic and simple cubic phases are separated by a first-order phase transition, as are the simple cubic and the fluid phases. Note that the accuracy of our phase diagram here depends on the accuracy with which we can estimate the coexistence densities from the equations of state, a process which is hindered by hysteresis and similar free-energy barriers. More accurate values for the coexistence densities can be obtained by common tangent constructions between the free energies of the coexisting phases, and we will do so in Section 6.2 to check how good our estimates from this chapter are.

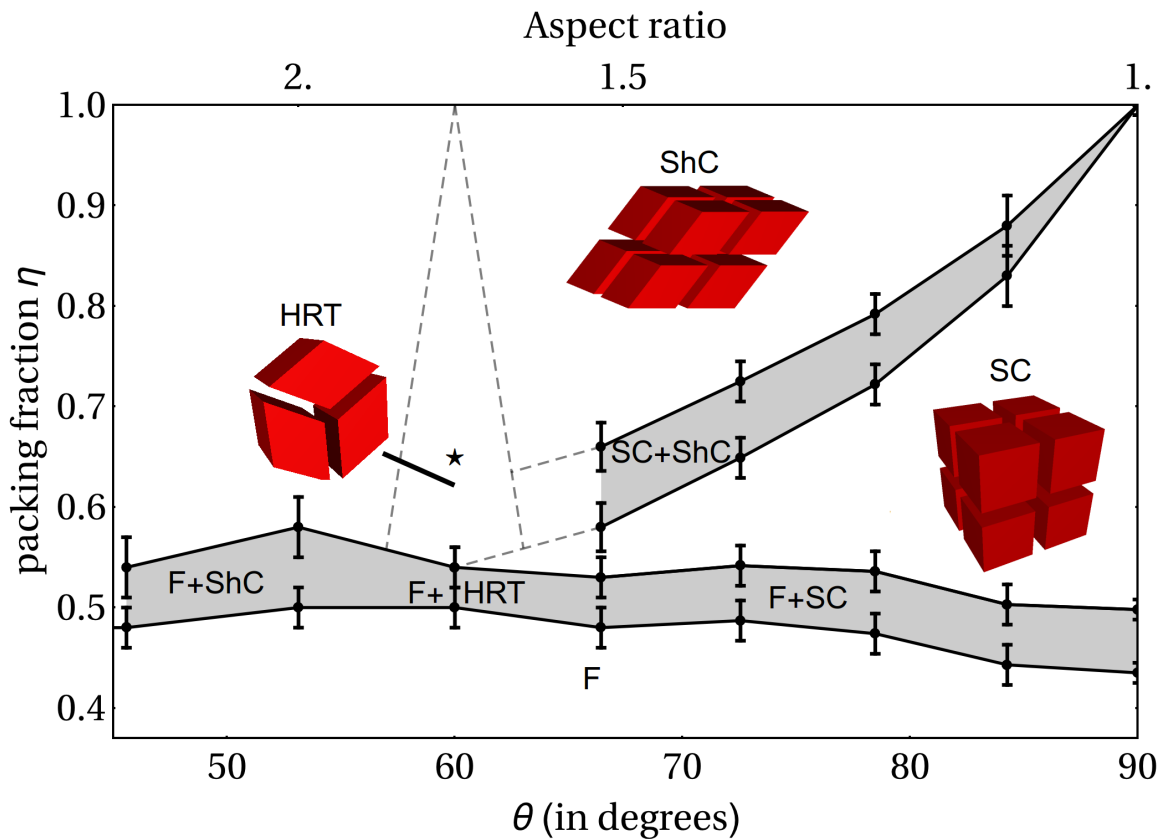


Figure 14: Phase diagram cutout of the $\theta - \eta$ plane for the slanted cube system. Indicated are the fluid (F), simple cubic (SC), sheared cubic (ShC), and hexagonal random tiling (HRT) phases. The shaded areas denote coexistence regions, the star indicates the highest density for which the HRT phase was observed, and the dashed lines indicate our guess for the phase boundaries we have not been able to observe directly.

5 Using shape switch MC simulations to calculate free energy differences

In this chapter, we will describe a new method to obtain free energies from Monte Carlo simulations; one which is based around thermodynamic integration over a shape parameter and was inspired by the lattice switch method of Bruce and Wilding [57, 60] and recent work by van Anders and Glotzer [61]. Integration over a shape parameter has been done before in Ref. [62] for a system of hard dumbbells, where one can integrate over the center-to-center distance L . However, this was only possible because the derivative of the free energy with respect to the center-to-center distance $\partial F/\partial L$ could be expressed in terms of the pressure-density equation of state $p(\rho)$. In this work, we construct a *general* scheme for integration over a shape parameter, one which does not require knowledge of the analytical form of such a derivative. To do so, we first construct an ensemble in which the shape is a thermodynamic variable, one that thermally fluctuates around some average. In such an ensemble, there will exist a derivative $\partial F/\partial\theta$, a derivative of the free energy with respect to the shape, and a thermodynamic conjugate to the shape parameter: $\mu_\theta N$. Next, we explain how to calculate this derivative; first outlining the basic scheme, then in a more advanced way that employs a biased ensemble to deal with sampling issues that might arise due to large entropic barriers. Following this explanation, we describe the methodology of these simulations: which step sizes to use and how to choose an appropriate bias. Finally, we apply our simulation to a simple cubic crystal of slanted cubes and look at the $\mu_\theta(\theta)$ shape equation of state.

5.1 Shape as a thermodynamic variable

If you know the free energy at one particular density ρ , you can obtain the free energy at another density by thermodynamic integration over the equation of state $p(\rho)$. To compute $p(\rho)$, one can use the well-known NVT or NPT ensembles, but these ensembles have fixed particle shapes. They have no “shape baths”. To integrate over the “shape equation of state”, $\partial F/\partial\theta = \mu_\theta(\theta)$, we must first define an ensemble in which the slant angle θ becomes a thermodynamic variable. Such an ensemble has been constructed in Ref. [61]. Let us apply their derivation to find our desired ensemble. Following Ref. [61], we define a shape-dependent Hamiltonian

$$\mathcal{H}(\theta) = \sum_{i=1}^N \left(\frac{\mathbf{p}_i^2}{2m} + \frac{1}{2} \mathbf{L}_i^{-1} \mathbf{I}_\theta^{-1} \mathbf{L}_i \right) + U_\theta(\mathbf{r}^N, \boldsymbol{\Omega}^N), \quad (52)$$

where \mathbf{p}_i are momenta, \mathbf{L}_i are angular momenta, \mathbf{I} is the moment of inertia tensor, and U is the interaction potential. From the Shannon-Jaynes entropy [63]

$$S = - \sum_{\sigma} \pi_{\sigma} \ln \pi_{\sigma} - \beta (\pi_{\sigma} \mathcal{H} - \langle E \rangle - \mu_{\theta} N (\pi_{\sigma} \theta - \langle \theta \rangle)) \quad (53)$$

and the normalization $\sum_{\sigma} \pi_{\sigma} = 1$, where σ labels the macrostates with probability π_{σ} and E is the total energy of the system, one can then derive the partition function:

$$Z = \sum_{\sigma} e^{-\beta(\mathcal{H} - \mu_{\theta} N \theta)}. \quad (54)$$

Filling in the Hamiltonian from Eq. 52 then yields

$$Z = \int d\mathbf{p}^N d\mathbf{L}^N d\mathbf{r}^N d\boldsymbol{\Omega}^N d\theta e^{-\beta \left(\sum_{i=1}^N \frac{\mathbf{p}_i^2}{2m} + \frac{1}{2} \mathbf{L}_i^{-1} \mathbf{I}_\theta^{-1} \mathbf{L}_i \right)} e^{-\beta U_\theta(\mathbf{r}^N, \boldsymbol{\Omega}^N)}. \quad (55)$$

We can do the Gaussian integration over the (angular) momenta as we do normally, reducing the expression to:

$$Z = \text{const.} \int d\mathbf{r}^N d\boldsymbol{\Omega}^N d\theta \det[\mathbf{I}_\theta]^{N/2} e^{-\beta U_\theta(\mathbf{r}^N, \boldsymbol{\Omega}^N)}, \quad (56)$$

where we drop the irrelevant constant in later use. For notational simplicity we denote the partition function from Eq. 54 as $e^{-\beta\phi}$. We then get:

$$\langle \theta \rangle = - \frac{1}{N} \left(\frac{\partial \phi}{\partial \mu_\theta} \right)_{N,V,T}. \quad (57)$$

Making a Legendre transform $F = \phi + \mu_\theta N \theta$ to obtain a free energy and taking the thermodynamic limit $\langle \theta \rangle \rightarrow \theta$ then gives us

$$\mu_\theta = \frac{1}{N} \left(\frac{\partial F}{\partial \theta} \right)_{N,V,T}. \quad (58)$$

The quantity μ_θ represents the infinitesimal free-energy increase per particle upon an increase of the slant angle θ . For one particular value of the slant angle, we already know what μ_θ has to be. With our definition of the

particle shape (recall Fig. 1), $\theta > 90^\circ$ is identical to $\theta < 90^\circ$, but with the reflex and obtuse angles swapped. If $\mu_\theta(90^\circ)$ were nonzero, it would imply a preference for either $\theta > 90^\circ$ or $\theta < 90^\circ$, which would break the symmetry. Ergo:

$$\mu_\theta(90^\circ) = 0. \quad (59)$$

We will use this as a consistency check.

5.2 Obtaining a free-energy difference by switching

Consider a system which has two macrostates between which we would like to know the free-energy difference. Specifically, assume that these two macrostates correspond to different slant angles θ , denoted θ_a and θ_b . The associated partition function for the combined system including both θ_a and θ_b is then (from Eq. 56):

$$\begin{aligned} \mathcal{Z}(N, V, T, \theta) &= Z(N, V, T, \theta_a) Z(N, V, T, \theta_b) \\ &= \int d\mathbf{r}^N d\mathbf{\Omega}^N \det[\mathbf{I}_{\theta_a}]^{N/2} e^{-\beta U_{\theta_a}(\mathbf{r}^N, \mathbf{\Omega}^N)} \int d\mathbf{r}^N d\mathbf{\Omega}^N \det[\mathbf{I}_{\theta_b}]^{N/2} e^{-\beta U_{\theta_b}(\mathbf{r}^N, \mathbf{\Omega}^N)}, \end{aligned} \quad (60)$$

where we use \mathcal{Z} to indicate the partition function for all θ , and Z to indicate the partition function for one particular value of θ . The probability of the system being in either macrostate, i.e. having either slant angle $\theta_i \in \{\theta_a, \theta_b\}$, is then:

$$P(\mathbf{r}^N, \mathbf{\Omega}^N | N, V, T, \theta_i) = \frac{\det[\mathbf{I}_{\theta_i}]^{N/2} e^{-\beta U_{\theta_i}(\mathbf{r}^N, \mathbf{\Omega}^N)}}{\mathcal{Z}(N, V, T, \theta)}. \quad (61)$$

As stated before, we use a hard potential that is based on the particle shape:

$$\beta U_\theta(\mathbf{r}^N, \mathbf{\Omega}^N) = \begin{cases} \infty & \text{when } \mathbf{r}^N, \mathbf{\Omega}^N \text{ is an overlapping configuration for the shape associated with } \theta. \\ 0 & \text{when } \mathbf{r}^N, \mathbf{\Omega}^N \text{ is nonoverlapping.} \end{cases} \quad (62)$$

The free-energy difference between the two possible states of θ can then be expressed as follows [57, 64]:

$$\begin{aligned} F(N, V, T, \theta_a) - F(N, V, T, \theta_b) &= -k_B T \ln Z(N, V, T, \theta_a) + k_B T \ln Z(N, V, T, \theta_b) \\ &= k_B T \ln \left[\frac{Z(N, V, T, \theta_b)}{Z(N, V, T, \theta_a)} \right] \\ &= k_B T \ln \left[\frac{\int d\mathbf{r}^N d\mathbf{\Omega}^N \det[\mathbf{I}_{\theta_b}]^{N/2} e^{-\beta U_{\theta_b}(\mathbf{r}^N, \mathbf{\Omega}^N)}}{\int d\mathbf{r}^N d\mathbf{\Omega}^N \det[\mathbf{I}_{\theta_a}]^{N/2} e^{-\beta U_{\theta_a}(\mathbf{r}^N, \mathbf{\Omega}^N)}} \right] \\ &= k_B T \left(\frac{N}{2} \ln \left[\frac{\det \mathbf{I}_{\theta_b}}{\det \mathbf{I}_{\theta_a}} \right] + \ln \left[\frac{\int d\mathbf{r}^N d\mathbf{\Omega}^N e^{-\beta U_{\theta_b}(\mathbf{r}^N, \mathbf{\Omega}^N)}}{\int d\mathbf{r}^N d\mathbf{\Omega}^N e^{-\beta U_{\theta_a}(\mathbf{r}^N, \mathbf{\Omega}^N)}} \right] \right) \\ &= k_B T \left(\frac{N}{2} \ln \left[\frac{\det \mathbf{I}_{\theta_b}}{\det \mathbf{I}_{\theta_a}} \right] + \ln \left[\frac{Q(N, V, T, \theta_b)}{Q(N, V, T, \theta_a)} \right] \right). \end{aligned} \quad (63)$$

Note that one is free to distribute the mass as one wishes throughout the particle and hence one can, without loss of generality, choose $\mathbf{I}_{\theta_a} = \mathbf{I}_{\theta_b}$. In this case, Eq. 64 simplifies to:

$$F(N, V, T, \theta_a) - F(N, V, T, \theta_b) = k_B T \ln \left[\frac{Q(N, V, T, \theta_b)}{Q(N, V, T, \theta_a)} \right]. \quad (64)$$

We can evaluate this configurational term by devising a simulation in which the system is allowed to switch between these two states, and evaluating the two potentials simultaneously. In the spirit of Bruce and Wilding [57], we refer to this type of simulation as ‘‘shape-switch’’ simulations. Such a simulation is described in the following section.

5.3 Implementation of a shape-switch move

To sample the probability ratio from Eq. 64, we construct a simulation in an ensemble with a fixed number of particles N , a fixed volume V and a fixed temperature T , but where the system is free to change the shape (of all particles simultaneously) between θ_a and θ_b . This is done by defining an MC move that changes the shape. As with any Markov chain Monte Carlo step, we have to make these steps carefully to satisfy detailed balance if we want our simulation to converge to the equilibrium distribution. As we are only evaluating the configuration term of Eq. 64, the Metropolis criterion is:

$$\text{acc}(U_{\theta_a} \rightarrow U_{\theta_b}) = \min \left[1, e^{-\beta(U_{\theta_b} - U_{\theta_a})} \right] = \begin{cases} 1 & \text{when } U_{\theta_a} \rightarrow U_{\theta_b} \text{ causes no overlap.} \\ 0 & \text{when } U_{\theta_a} \rightarrow U_{\theta_b} \text{ causes overlap.} \end{cases} \quad (65)$$

which for a hard potential comes down to checking whether the change of shape would cause overlaps. The shape-switch move can be implemented as follows. We define a tracking parameter α which can take two values; $\alpha = 0$ for the θ_a state and $\alpha = 1$ for the θ_b state. When the shape-switch MC move is called, it goes through the following steps:

Monte Carlo shape-switch move:

1. The current shape state parameter α is checked:
If it corresponds to shape θ_a , the shape is changed to θ_b .
If it corresponds to shape θ_b , the shape is changed to θ_a .
2. The Metropolis criterion is evaluated; the current configuration with the new shape is checked for overlap. If overlap is found, the move is marked as rejected.
3. If the move is marked as rejected, the shape is changed back.
If the move is not marked as rejected, α is changed to correspond to the new shape.




Figure 15: A forbidden shape-switch move. When switching the slant angle of the particles from 90° to 66° , multiple overlaps are created.

For every MC cycle we now have information about the shape state α of the system. Because the Boltzmann weight of a hard potential is either 0 or 1, calculating the probability ratio of the two shape states becomes the trivial exercise of counting how often the two states are found in our simulation. Every so often (we use every 50 MC cycles) we sample the system. If α corresponds to θ_a , we add a count to our a variable n_a , and vice versa for θ_b with n_b . The configurational part of the free-energy difference is then estimated by

$$\Delta F_{conf.} = k_B T \ln \left[\frac{Q(N, V, T, \theta_b)}{Q(N, V, T, \theta_a)} \right] = k_B T \ln \left[\frac{n_b}{n_a} \right]. \quad (66)$$

5.4 Extended sampling / biasing

The effectiveness of the method described in the previous paragraph hinges on the assumption that the system is able to do enough shape switches to adequately sample the available phase space of both macrostates θ_a and θ_b . If we were to naively pick $\theta_a = 90^\circ$ and one of the angles we looked at earlier, say, $\theta_b = 72^\circ$, this assumption would be *very* wrong. Consider a variable M that denotes the number of overlapping particle pairs that are created if the shape of all particles is switched. We can track this value during a simulation, creating a histogram which gives us an estimate of $p(M)$: the probability of the system being in a microstate with a certain value M . See Fig. 16. Remember that to switch the shape, we require it to create no overlaps; the system must have $M = 0$. Already, we can see that this does not bode well for our shape-switch algorithm. Fitting the histogram to a Gaussian distribution (in accordance with the central limit theorem) and plotting it on a logarithmic scale, the problem is revealed in full (Fig. 16, right). The probability of the system visiting the required $M = 0$ state is so small ($\sim 10^{-500}$) that there is no way for our algorithm to obtain adequate sampling. This is unsurprising. It is nearly the exact same problem as Wilding et al. observed in Ref. [57] when attempting to switch between the face-centered cubic (FCC) and hexagonal close-packed (HCP) lattices of hard spheres. It should come as no surprise, then, that we can utilize a very similar solution.




Figure 16: (left) The probability histogram $p(M)$ and the Gaussian fit on a normal scale. Results were obtained from an unbiased shape-switch simulation attempting to switch from 90° particles (cubes) to 70° slanted particles. (right) The fit plotted on a (base 10) logarithmic scale. Here one can see the enormity of the "ergodic barrier", the trough in probability that prevents our system from performing switch moves.

Wilding et al.'s [57] solution was to bias all moves such that switchable configurations are visited more often. In practice, this would be done by modifying the probability distributions used in all MC moves as follows:

$$P(\mathbf{r}^N, \boldsymbol{\Omega}^N | N, V, T, \theta, \{\eta\}) = P(\mathbf{r}^N, \boldsymbol{\Omega}^N | N, V, T, \theta) e^{\eta(\mathcal{M}(\mathbf{r}^N, \boldsymbol{\Omega}^N))}, \quad (67)$$

where $\{\eta\} = \eta(\mathcal{M}(\mathbf{r}^N, \boldsymbol{\Omega}^N))$ is a set of weights that depends on \mathcal{M} , the difference in number of overlaps for the two states of the system, which is defined as

$$\mathcal{M}(\mathbf{r}^N, \boldsymbol{\Omega}^N) = M_{\theta_a}(\mathbf{r}^N, \boldsymbol{\Omega}^N) - M_{\theta_b}(\mathbf{r}^N, \boldsymbol{\Omega}^N). \quad (68)$$

Here, $M_\theta(\mathbf{r}^N, \boldsymbol{\Omega}^N)$ is the number of overlaps the system would have for a given configuration $(\mathbf{r}^N, \boldsymbol{\Omega}^N)$. Note that if the system is in the state θ_a (θ_b), we have that M_{θ_a} (M_{θ_b}) is guaranteed to be zero, as the simulation is not allowed to visit configurations with overlap. As can be seen in Fig. 15, configurations with the same spatial and orientational coordinates can have a different number of overlaps, depending on the particle shape. This set of weights η is not known *a priori* and needs to be determined before we can utilize it. The expression for the free-energy difference is also modified, as we have to unbiased:

$$\begin{aligned} \frac{\beta(F_{\theta_a} - F_{\theta_b})}{N} &= \frac{1}{N} \ln \left[\frac{P(\mathbf{r}^N, \boldsymbol{\Omega}^N | N, V, T, \theta_b)}{P(\mathbf{r}^N, \boldsymbol{\Omega}^N | N, V, T, \theta_a)} \right] \\ &= \frac{1}{N} \ln \left[\frac{\sum_{\mathcal{M}} P(\mathbf{r}^N, \boldsymbol{\Omega}^N | N, V, T, \theta_a, \{\eta(\mathcal{M})\}) e^{-\eta(\mathcal{M})}}{\sum_{\mathcal{M}} P(\mathbf{r}^N, \boldsymbol{\Omega}^N | N, V, T, \theta_a, \{\eta(\mathcal{M})\}) e^{-\eta(\mathcal{M})}} \right], \end{aligned} \quad (69)$$

which, since $\sum_{\mathcal{M} < 0} P(\theta_b | \mathcal{M}) = 0$ and $\sum_{\mathcal{M} > 0} P(\theta_a | \mathcal{M}) = 0$, reduces to

$$\frac{\beta(F_{\theta_a} - F_{\theta_b})}{N} = \frac{1}{N} \ln \left[\frac{\sum_{\mathcal{M} \geq 0} P(\mathbf{r}^N, \boldsymbol{\Omega}^N | N, V, T, \theta_b, \{\eta(\mathcal{M})\}) e^{-\eta(\mathcal{M})}}{\sum_{\mathcal{M} \leq 0} P(\mathbf{r}^N, \boldsymbol{\Omega}^N | N, V, T, \theta_a, \{\eta(\mathcal{M})\}) e^{-\eta(\mathcal{M})}} \right]. \quad (70)$$

Unless the difference between the states θ_a and θ_b is extremely small, the state $\mathcal{M} = 0$ will hardly ever be visited without biasing the sampling. If we unfold the bias, we can thus quite safely neglect the contribution of the $M = 0$ states, and simply write:

$$\frac{\beta(F_{\theta_a} - F_{\theta_b})}{N} = \frac{1}{N} \ln \left[\frac{\sum_{\mathcal{M} > 0} P(\mathbf{r}^N, \boldsymbol{\Omega}^N | N, V, T, \theta_b, \{\eta(\mathcal{M})\}) e^{-\eta(\mathcal{M})}}{\sum_{\mathcal{M} < 0} P(\mathbf{r}^N, \boldsymbol{\Omega}^N | N, V, T, \theta_a, \{\eta(\mathcal{M})\}) e^{-\eta(\mathcal{M})}} \right]. \quad (71)$$

There is one more trick we can use. The position of the peak of the (roughly) Gaussian distribution of \mathcal{M} -values the system takes depends on the magnitude of the difference between the two states θ_a and θ_b . For efficiency purposes, we have two requirements. First, the simulation must sample the values of \mathcal{M} corresponding to the unbiased distribution. Second, the simulation must visit the $\mathcal{M} = 0$ region often enough for the shape correlation time to exceed the sampling time interval. Looking back at Fig. 16, one can imagine these requirements as being quite tricky to fulfill. The position of the probability maximum is far away from $\mathcal{M} = 0$, and for every switch that is made the system must travel all the way back to $\mathcal{M} = 0$ for the next switch to occur, spending most of its time sampling \mathcal{M} -values of relatively low importance. Ideally, we would like to decrease

or eliminate this “dead space” in the \mathcal{M} -domain, by shifting the position of the probability maximum closer to $\mathcal{M} = 0$. This can be done by decreasing the size of the θ -step, making the two states θ_a and θ_b more similar. This has the advantage that we now also gather information about intermediate values of θ . However, it also has the unfortunate disadvantage that we now have to do multiple simulations, becoming worse the smaller the steps we take. As it turns out, for our code a reasonable compromise between accuracy and time used for one step was found for a step size of $\delta\theta = 0.15^\circ$ for $N = 1000$ particles, and $\delta\theta = 0.01^\circ$ for $N = 8000$ particles. To sample roughly the entire θ domain, that would imply $(90^\circ - 70^\circ)/0.15^\circ \approx 134$ or $(90^\circ - 70^\circ)/0.01^\circ = 2000$ steps, which are a wholly unreasonable number of steps. For reference, one such step required a simulation time of roughly 8 hours. Excluding the fact that we would need to repeat these simulations for various vacancy concentrations, even the shortest of such a simulation would take well over a month of CPU time. So we still have to do better.

If we can think of θ as a proper thermodynamic variable, we can do thermodynamic integration over its equation of state $\mu(\theta)$. Hence, instead of computing δF , we only need to compute

$$\mu_\theta = \frac{1}{N} \left(\frac{\partial F}{\partial \theta} \right)_{N,V,T} = \frac{1}{N} \lim_{\delta\theta \downarrow 0} \left(\frac{\delta F(\delta\theta)}{\delta\theta} \right)_{N,V,T}, \quad (72)$$

which is a quantity which actually benefits from a small step size $\delta\theta$. However, there are concerns to be addressed. Firstly, like with any form of thermodynamic integration, the integration path must not cross any first-order phase transitions. Secondly, in order to obtain accurate results for the free-energy-difference curve $\Delta F(\theta)$, the distance between sampling points must be small enough to properly cover the curvature of the $\mu_\theta(\theta)$ equation of state. Save for the value at $\theta = 90^\circ$, which we have derived earlier, the general shape of this equation of state is completely unknown. However, considering that we are looking at two very similar systems forming very similar SC phases, we can probably make the assumption that the free-energy difference $\Delta F(\theta)$ between a plastic cubic crystal comprised of cubes and a plastic cubic crystal comprised of slanted cubes with slant angle θ is a smoothly varying function of θ over the domain $70^\circ \leq \theta < 90^\circ$. This assumption empowers us to sample the derivative $\partial F/\partial\theta$ for a relatively small set of values of θ and from that, through thermodynamic integration, determine (an estimate for) the free-energy difference curve $\Delta F(\theta)$. Formally, we determine free-energy differences $\{\Delta F_{\theta_i \rightarrow \theta_{i+1}}\}$ for a fixed step size at various values of θ spaced, say, 1° apart. Then, we determine the derivative

$$\left. \frac{\partial F}{\partial \theta} \right|_{\theta_i + \Delta\theta/2} \approx \frac{\Delta F_{\theta_i \rightarrow \theta_{i+1}}}{\Delta\theta}. \quad (73)$$

This is where we have to use our assumption. It implies that the set of angles $\{\theta_i\}$ we sample the derivative at are spaced close enough that we can smoothly interpolate between them and obtain an accurate guess for the intermediate values.

5.5 Methodology of shape-switch simulations

Let us apply the shape-switch method to the region of phase space we want to sample. First, we must construct a path. This path should be in the region where the SC phase is the equilibrium phase, as this is the phase for which we have data to use as a reference state. Looking back at Fig. 14, we can see that a possible path to take would be from $\theta = 90^\circ$ to $\theta = 66^\circ$, at a constant density $\rho\sigma^3 = 0.58$. For each θ_i we wish to sample, we create nonoverlapping SC crystal configurations as initial configurations to speed up equilibration and to create as similar crystal lattices as possible. We also need to find proper values for the step size $\Delta\theta$ and the corresponding bias. The former can be found by simply looking at the \mathcal{M} -histograms for various step sizes and observing the distance between the two roughly Gaussian distributions. See Fig. 18. A good value for $\Delta\theta$ strikes a balance between three criteria:

1. The two distributions are clearly separated, such that the assumption we used to go from Eq. 70 to Eq. 71 holds.
2. The distance between the two peaks is as low as possible, such that the path back to $\mathcal{M} = 0$ is as short as possible. This keeps the correlation time of the shape low, and allows for faster sampling.
3. The step size $\Delta\theta$ is low compared to the curvature of μ_θ , such that the error from using a finite-difference approximation for the derivative is low.

We found that for $N = 1000$ particles, a good step size is $\Delta\theta = 0.10^\circ$, and for $N = 8000$ particles, it is $\Delta\theta = 0.01^\circ$.

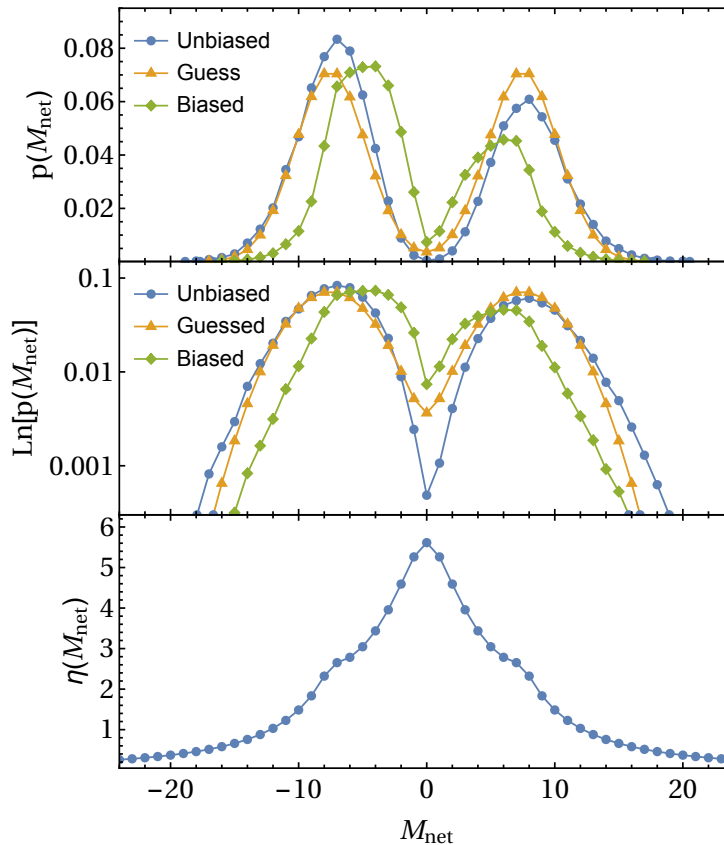


Figure 17: (Top) \mathcal{M} -distribution for the unbiased run in (blue circles), the fitted composite Gaussian function (orange triangles), and the resulting biased \mathcal{M} -distribution (green diamonds). (Middle) Same as above, but on a logarithmic scale. Note that $p(\mathcal{M} = 0)$ has increased by a factor of 10. (Bottom) The applied bias distribution $\eta(\mathcal{M})$.

Next, we should determine an appropriate bias. A general framework for this, along with suggestions on how to improve and automate it, can be found in Ref. [60]. For the step sizes $\Delta\theta$ we use, however, the ergodic block is quite small (on the order of 10^{-4}). This means we can get away with a fairly poor guess for the bias. Therefore, we shall use only the basics of the extended sampling framework, and use a single bias $\eta(\{\mathcal{M}\})$ for all θ_i .

We perform an unbiased simulation (where $\eta(\{\mathcal{M}\}) = 0$) in which we determine μ_θ at $\theta = 70^\circ$. To determine the bias, we look at the \mathcal{M} -histogram, and fit it to a symmetric two-Gaussian function

$$G(\mathcal{M}) = \frac{1}{2\sqrt{2\pi}\sigma_{\mathcal{M}}} \left(\exp \left[-\frac{1}{2} \left(\frac{\mathcal{M} - \bar{\mathcal{M}}}{\sigma_{\mathcal{M}}} \right)^2 \right] + \exp \left[-\frac{1}{2} \left(\frac{\mathcal{M} + \bar{\mathcal{M}}}{\sigma_{\mathcal{M}}} \right)^2 \right] \right). \quad (74)$$

The fact that the two Gaussians have almost the same expectation values $\bar{\mathcal{M}}$ is an empirical observation, though in a handwaving way we could argue it must be so based on the symmetry of how the shape changes when we change θ . The same goes for the standard deviation $\sigma_{\mathcal{M}}$. We also fix the amplitude of each Gaussian to $1/2$, since this amplitude ratio is what determines μ_θ , and will be different for each θ_i . The expectation value $\bar{\mathcal{M}}$ will also be different for each value of θ , but for our system this effect is small enough to not influence our choice of bias. Fig. 17 shows the \mathcal{M} -distribution for the initial unbiased run, the fitted composite Gaussian, and the resulting biased distribution on both a normal and a logarithmic scale, along with the applied bias. Here we have chosen to define the bias as

$$\eta(\mathcal{M}) = \begin{cases} -\ln G(\mathcal{M}) & \text{for } \mathcal{M} \leq \bar{\mathcal{M}} \\ -\frac{\ln G(\bar{\mathcal{M}})\bar{\mathcal{M}}^2}{\mathcal{M}^2} & \text{for } \mathcal{M} > \bar{\mathcal{M}} \end{cases}, \quad (75)$$

but there are many other possible choices. The main points are that $p(\mathcal{M} = 0)$ is now higher in the biased system than in the unbiased ensemble, and that we have not created any local maxima for the system to get stuck in. We are now ready to start our shape-switch simulation.

5.6 The result: a $\mu_\theta(\theta)$ equation of state

We use shape-switch simulations to determine μ_θ at 21 values of θ , spaced roughly 1° apart. The systems we consider are simple cubic crystals of $N = 1000$ or $N = 8000$ slanted cubes at a density $\rho\sigma^3 = 0.58$. Starting from an (N, V, T) equilibrated simple cubic lattice, the system is given 10^4 MC cycles to equilibrate in the new $(N, V, T, \theta, \{\eta\})$ ensemble, after which we sample \mathcal{M} every 50 MC cycles for another 1.5×10^6 MC cycles. A typical \mathcal{M} -histogram is shown on the left in Fig. 18. On the right in the same Fig., we show μ_θ for the sampled θ_i , along with a fifth-order polynomial fit. The reduced ($f = F/k_B T N$) free-energy difference $\Delta f(\theta) = f(\theta) - f(90^\circ)$ is shown as the shaded area. We can see that μ_θ does indeed go to zero at $\theta = 90^\circ$ as required. By combining the obtained $\Delta f(\theta)$ with separate free energies obtained by e.g. the Frenkel-Ladd method we can obtain the free energy $f_{SC}(\theta)$ for one density. By also calculating equations of state over the pressure-density regime, we obtain $f_{SC}(\theta, \rho)$. We will not list the free energies we obtain this way here, but will use them in the following chapter to update the phase diagram and extract information about the equilibrium vacancy behaviour.

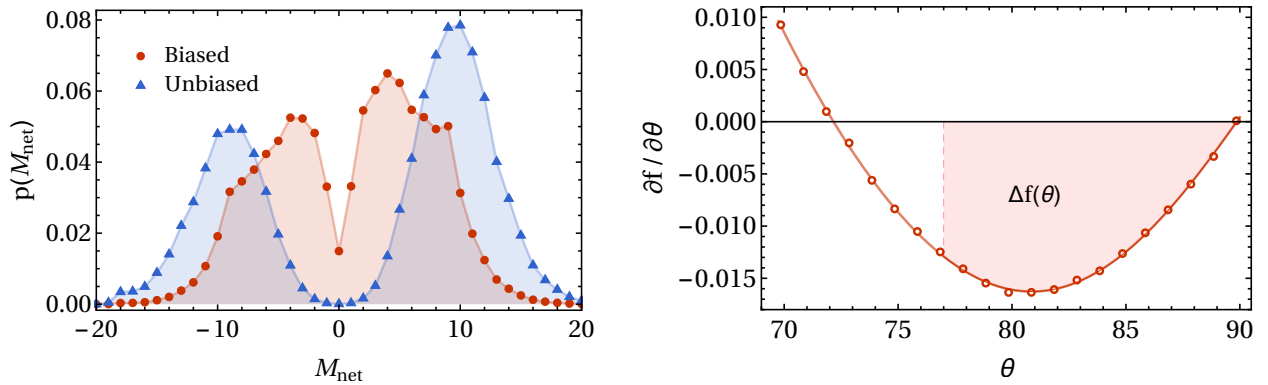


Figure 18: Left: the estimate for the probability distribution $p(\mathcal{M})$. The blue circles indicate the direct result of the biased simulation, the orange squares the results after unbiasing. We obtain the free-energy difference for this step from the sum of the unbiased data points $(\sum_{\mathcal{M}>0} p_{\mathcal{M}})/(\sum_{\mathcal{M}<0} p_{\mathcal{M}})$, and from there obtain the derivative $\partial f/\partial\theta$. Right: the derivative $\partial f/\partial\theta$ for all relevant values of θ . The solid orange line indicates the polynomial fit through these data points, and the shaded area then indicates the free-energy difference $\Delta f(\theta) = f(90^\circ) - f(\theta)$.

6 Phase diagram and vacancy behaviour of a simple cubic crystal of slanted cubes

Let us now return to the questions we posed at the beginning of this thesis: what is the equilibrium phase diagram and, more importantly, how does the vacancy behaviour of slanted cubes differ from that of cubes? Using the shape switch method we developed in Chapter 5, we can now obtain the full free-energy profile of the simple cubic phase. That is to say, for all slant angles θ , densities ρ , or vacancy fractions α for which the SC phase is stable, we can obtain the free energy $F_{SC}(\theta, \rho, \alpha)$. From the qualitative analysis in Chapter 4 we have seen that the SC phase behaviour seems to be quite independent of the slant angle. But *how* independent is it? We have seen that the density-pressure equations of state are different. For the lower densities, the system shows signs of vacancy subsaturation just like the cubes do. One possible quantitative measure we can compute is the equilibrium vacancy concentration as a function of the slant angle and the density: $\alpha_{eq}(\theta, \rho)$. What is this value? The fact that there is a nonzero equilibrium vacancy concentration is the result of two contributions to the free energy: namely the free-energy cost to create vacancies, and the combinatorial entropy gained by creating them. In other words:

$$F = F_{perfect} + F_{cost} + F_{comb}, \quad (76)$$

where $F_{perfect}$ describes the free energy of a lattice without any vacancies, F_{cost} the free-energy cost to create a vacancy and F_{comb} the combinatorial free energy. The functional form of F_{cost} is unknown and has to be determined from simulations, but F_{comb} can be found by simply considering the number of ways one can arrange N particles on N_L lattice sites:

$$F_{comb} = -k_B T \ln \left[\frac{N_L!}{N!(N_L - N)!} \right]. \quad (77)$$

From the work of Smallenburg et al. [31] we know that in the case of cubes the vacancies hardly interact with each other. Let us assume that this is also the case for slanted cubes, and check the validity of this assumption afterwards. If the vacancies are ideal (that is, there is no vacancy-vacancy interaction), the free-energy cost incurred by having vacancies can be described by

$$F_{cost} = \epsilon_{vac} N_{vac}, \quad (78)$$

where ϵ_{vac} is the free energy cost of creation one vacancy and N_{vac} is the number of vacancies. Let us rewrite this to a dimensionless free energy per particle, and in terms of α and N_L . Using $N_{vac} = \alpha N_L$ and $N = (1 - \alpha)N_L$:

$$f_{cost} = \frac{F_{cost}}{k_B T N} = \frac{1}{k_B T} \frac{\alpha \epsilon_{vac} N_L}{(1 - \alpha)N_L} = \frac{\alpha}{1 - \alpha} \frac{\epsilon_{vac}}{k_B T} \quad (79)$$

where $f_{perfect}$ describes the free energy of the lattice without any defects. We now use all our available tools to construct a path through (θ, ρ, α) phase-space to obtain a free-energy profile $f(\theta, \rho, \alpha)$ for the SC phase. See Fig. 19 for a visual representation. The first step of this path is to calculate the free energy of an SC lattice at some particular density/slant angle. The starting point is the data from Smallenburg et al., who calculated the free energies of an SC lattice of cubes for different vacancy fractions and densities. This gives us $f(\alpha)|_{\theta=90^\circ, \rho=0.56}$. Next, we calculate the $p(\rho)$ equations of state for the SC lattice of various slant angles using either a modified version of Smallenburg's EDMD code or our own anisotropic NPT code. This will give us the density dependence of the free energy $f(\rho)|_{\theta, \alpha}$. Instead of repeating the multiple Einstein integrations for every slant angle we want to study, we use the shape switch simulations described in Chapter 5 to obtain the $\mu_\theta(\theta)$ profile. This completes the picture, giving us $f(\theta)|_{\rho, \alpha}$.

6.1 Shape invariance of vacancy formation cost

Through a combination of methods we now have access to the full free-energy profile $f(\theta, \rho, \alpha)$ of the simple cubic phase. It is now possible to calculate the equilibrium vacancy concentration of the crystal. To do so, we subtract the free energy of the perfect lattice from the free energies we have calculated for the defective lattices, thereby obtaining $f - f_{perf} = f_{cost} + f_{comb}$, which is the contribution to the free energy solely of the vacancies. This quantity is shown on the left in Fig. 20. To extract the equilibrium vacancy concentration from this data, we now only have to determine the vacancy concentration for which this vacancy contribution to the free energy is at a minimum. These are the black dots in the left plot of Fig. 20. A good check to do at this point is to also subtract the combinatorial free energy, and check whether the free-energy cost of vacancy creation is indeed linear with the number of vacancies, as was the case for the perfect cubes. We take the free energies we have calculated, subtract the free energy of the perfect lattice and the combinatorial contribution, and obtain the data shown on the right in Fig. 20.

This plot reveals both something expected as well as something quite unexpected. The former is that the fits are quite good, indicating that indeed the vacancies do not interact with one another. This ideal behaviour

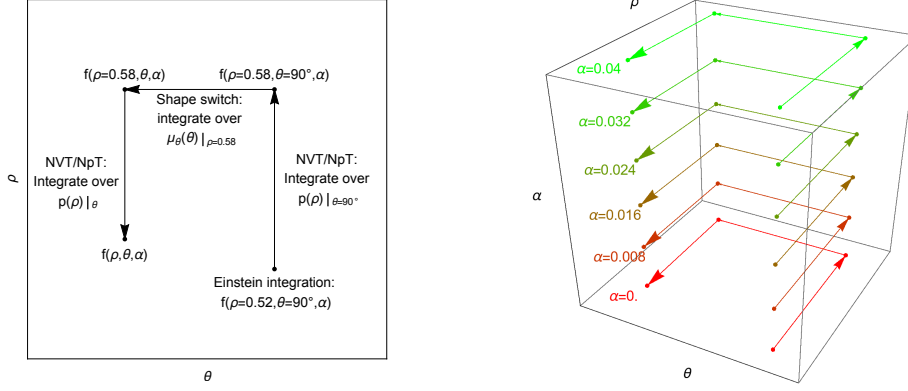


Figure 19: (left) The path through (θ, ρ) phase-space we use to obtain the free energies of the plastic cubic phase for a given vacancy concentration. (right) The paths we use to obtain free energies of the system for multiple vacancy concentrations.




Figure 20: (Left) Fits through the vacancy contributions to the free energies of an SC crystal of slanted cubes with slant angle $\theta = 72.5424^\circ$. The estimated free energy of the lattice without defects, f_{perf} , has been subtracted to plot all curves in the same graph. Empty points indicate the data, solid lines the fits, and filled points the minimum of these fits: the equilibrium vacancy concentration. (Right) Total vacancy free-energy cost (per particle). Shown are data and fits to Eq. 79 for $\theta = 90^\circ$, ($\theta = 78.4630^\circ$ and $\theta = 72.5424^\circ$). The colors here indicate various packing fractions in the range $0.53 \leq \eta \leq 0.61$, in steps of 0.02.

was also found for simple cubic crystals of parallel rounded cubes in Ref. [29] as well as for simple cubic crystals truncated cuboctahedra in Ref. [30]. The latter, unexpected result is that the vacancy cost seems to be independent of the slant angle of the particles, and only depends on the packing fraction. We can see this even more clearly if we plot the equilibrium vacancy concentration as a function of the packing fraction for various slant angles, as shown in Fig. 21.

It appears that the equilibrium vacancy concentration of the SC lattice is independent of the slant angle of the particles. Is this shape invariance also found in other systems? Comparing again with the results from Ref. [29] and Ref. [30], we find that such invariance is also found for the parallel rounded cubes, but not for the truncated cuboctahedra. Apparently, the vacancy behaviour is invariant to both slanting as well as rounding of the particles (though for the latter, it should be checked for rotating rounded cubes as well), but not to truncation to an octahedron. We find this to be a rather puzzling result. Somehow, the entropic cost of creating a vacancy is unaffected by slanting of the particles. We would think that effectively elongating the particles would hinder the formation of delocalized vacancies, but this is apparently not the case. This result offers some interesting opportunities for further research. Sadly, they fall beyond the scope of this thesis. We would direct the interested reader to the work of Berend van der Meer, who is studying these defects in greater detail. We, however, will return to a promise we made in Section 4.4, and update our phase diagram of the slanted cube system now that we have information on the free energies of the simple cubic phase.

6.2 Updated phase diagram

Now that we have free energies for the fluid, simple cubic and sheared cubic phases, we can accurately estimate the coexistence densities using common tangent constructions. Doing so for the slant angles $\theta = 72^\circ$ and $\theta = 78^\circ$, we find that the coexistence densities hardly shift at all with respect to our previous estimates: the difference is merely $\Delta\eta \sim 0.01$. The updated phase diagram is shown in Fig. 22. Not much has changed.




Figure 21: (Left) The equilibrium vacancy concentration α_{eq} as a function of the packing fraction η . For the three slant angles we have looked at, the curves collapse on top of each other. The dashed lines indicate the coexistence densities.

The coexistence densities determined by this method are much more accurate, and so we can display smaller errorbars, but the locations of the coexistence densities stay much the same. Using the information from Section 6.1, we also represent the equilibrium vacancy concentration of the simple cubic phase with a colour map. Our investigation into the phase and vacancy behaviour of slanted cubes, as much as could be part of the current work, is now complete.

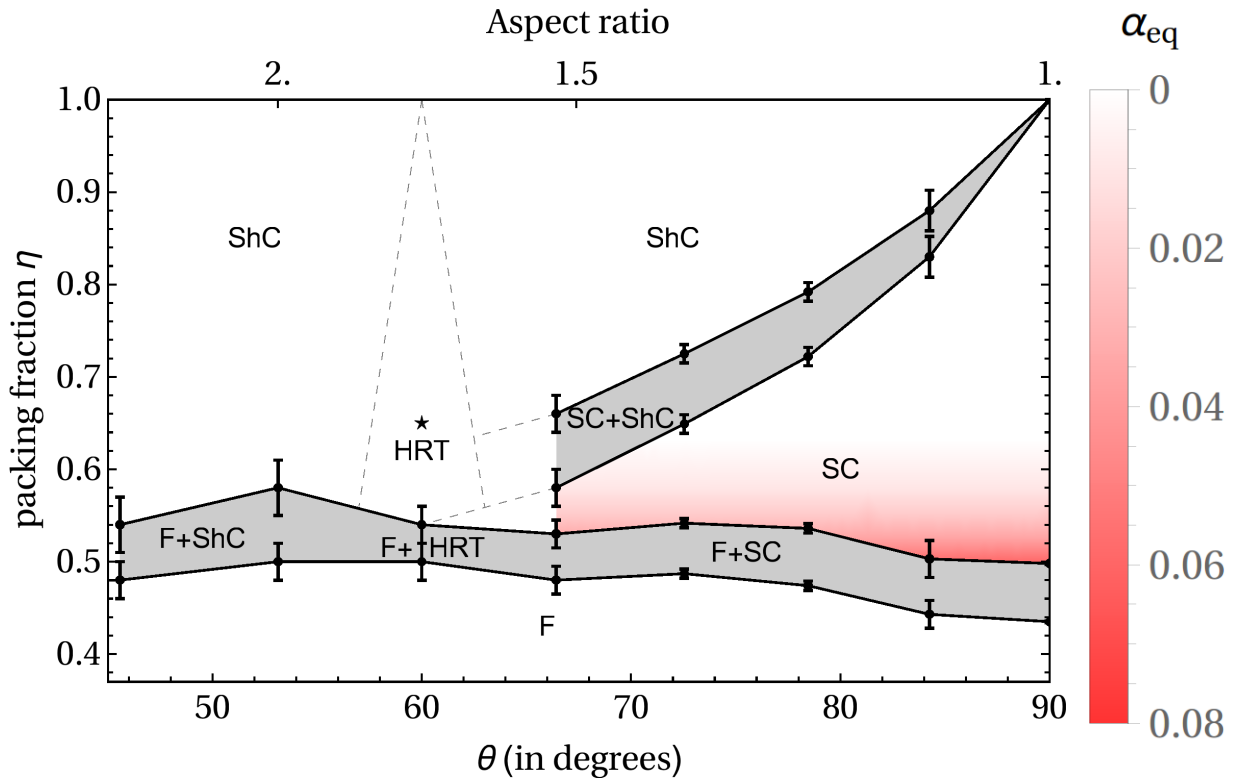


Figure 22: Phase diagram cutout of the $\theta - \eta$ plane for the slanted cube system. Indicated are the fluid (F), simple cubic (SC), sheared cubic (ShC), and hexagonal random tiling (HRT) phases. The shaded areas denote coexistence regions, the (red) colouring represents the equilibrium vacancy concentration in the SC phase, the star indicates the highest density for which the HRT phase was observed, and the dashed lines indicate our guess for the region of stability for the HRT phase. The coexistence points with a large errorbar were obtained from looking at snapshots and equations of state, while those with small errorbars were obtained from common tangent constructions after calculating the free energies of the various phases.

7 Summary and outlook

Using MC and EDMD simulations, we have partially mapped the phase diagram of the slanted cube system (Fig. 22). For slant angles of $66^\circ \leq \theta \leq 90^\circ$ we found a simple cubic crystal lattice at intermediate densities and a sheared cubic crystal at high densities. For $\theta = 60^\circ$, a phase with layers of hexagonal random tilings was found for intermediate densities. Because this tiling can fill space, it is also a candidate for the high density phase. We calculated the free energies of the sheared cubic and alternating cubic crystal lattices, both possible space-filling high density crystal lattices, for particles with a slant angle $\theta = 72^\circ$. We found that the free-energy difference between these two is $\sim 0.01 k_B T/N$ (roughly the accuracy of our free energy calculations), but consistently in favor of the sheared cubic lattice. We also calculated the free energies of the sheared cubic and rhombille crystals for particles with a slant angle of $\theta = 60^\circ$. Here, the free-energy difference was $0.004 k_B T/N$ in favor of the sheared cubic lattice. For the simple cubic lattice, we used a different approach to compute the free energies. Inspired by the lattice switch method of Bruce and Wilding [57, 60] and recent work by Van Anders et al [61], we calculated the free-energy difference between a simple cubic crystal of hard cubes and a simple cubic crystal of slanted cubes by thermodynamic integration over the slant angle θ . This was done by creating a simulation in which a system could switch between two different slant angles θ and, by use of the expression $\Delta F = F_{\theta_a} - F_{\theta_b} = k_B T \ln[P_{\theta_b}/P_{\theta_a}]$, calculating the free-energy difference corresponding to a change $\theta_a \rightarrow \theta_b$. By choosing a suitably small step size $\Delta\theta$ one could then treat the quantity $\Delta F/\Delta\theta$ as a generalized chemical potential μ_θ , conjugate to θ . A $\mu_\theta(\theta)$ “equation of state” can then be constructed and integrated over, using the normal restriction of thermodynamic integration that the integration path does not cross any first order phase transitions. We used these shape-switch calculations to obtain the free energy of the simple cubic phase for multiple vacancy concentrations, and looked at how the equilibrium vacancy concentration changes when the shape changes. To our surprise, we found this equilibrium concentration to be independent of the slant angle of the particles, depending only on the packing fraction of the simple cubic lattice. Finally, using all information gathered by anisotropic MC-NPT simulations and free-energy calculations, we drew a $\theta - \eta$ phase diagram of the slanted cube system.

We would like to take this opportunity to address the possible applications of the (shape) switch method. The most obvious application, and the one most similar to the one in this work, would be use it to do thermodynamic integration over a different shape parameter: the truncation of cuboctahedra in Ref. [30], the rounding of superballs in Ref. [29] or the aspect ratio of some elongated particle such as rods or dumbbells. As long as the particles are hard, one can even use the exact same biasing and sampling scheme based around the number of overlaps created upon switching. The only restriction to the method is the same as that imposed on all thermodynamic integration paths: that the derivative of the free energy is continuous i.e. that you do not cross (first-order) phase boundaries. However, this method need not be restricted to shape. We think that with some minor adjustments, it could also be adapted to use other “input” parameters for integration. One notable example of such a parameter would be the size ratio of hard spheres in a binary mixture, as studied in e.g. Refs. [65, 66].

To conclude this thesis, we would like to address some possible avenues of further research. The most pertinent subject, in our view, is that of the shape (in)dependence of the vacancies in a simple cubic crystal. Despite some discussion, we have been unable to find a satisfying explanation for this phenomenon. Several questions can be asked whose answers might hold the key to understanding this shape invariance, questions such as “is this shape invariance of vacancies unique to the simple cubic crystal?”, “Why does the truncation to octahedra matter, but not the slanting?”, “Is this behaviour also found for interstitials?”, “Which shape transformations is the vacancy behaviour invariant to?” or “Are the vacancies truly the same for every slant angle, or just the number of vacancies?”. Again we direct the reader to the work of Berend van der Meer and Laura Filion, who may soon be able to answer some of these questions.

Another intriguing possible research subject is the behaviour of the slanted cubes at or near $\theta = 60^\circ$. The three-dimensional behaviour of this system appears to be quite distinct from that of the two-dimensional case of rhombic platelets. Whereas we see no sign whatsoever of the random tiling phase for slant angles $\theta \leq 53^\circ$ or $\theta \geq 66^\circ$, there are signs of it for two-dimensional systems of rhombic platelets with slant angles of $\theta = 72^\circ$ (Ref. [18]). Two-dimensional rhombic networks of attractive patches also have vacancies (Ref. [67]) that could be investigated. While the dynamics of these networks are probably quite slow, the dynamics of an entropy-driven system of hard slanted cubes are likely to be quite fast. This crystal lattice is not very accommodating for delocalized vacancies, however, so we expect the equilibrium vacancy concentration in this random tiling crystal to be quite small.

References

1. Van Melsen, A. G. *From Atomos to Atom*. ISBN: 0486495841 (Dover Publications, New York, 1952).
2. Lloyd, G. *Early Greek Science: Thales to Aristotle* (Chatto and Windus, 1970).
3. Boyle, R. *A Defence of the Doctrine Touching the Spring and Weight of the Air* (1662).
4. Brown, R. *A brief account of microscopical observations made in the months of June, July and August 1827, on the particles contained in the pollen of plants; and on the general existence of active molecules in organic and inorganic bodies*. *Philosophical Magazine* **2**, 161–173 (1828).
5. Einstein, A. *Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen*. *Annalen der Physik* **322**, 549–560. ISSN: 00033804 (1905).
6. Perrin, J. *Mouvement brownien et réalité moléculaire*. *Annales de Chimie et de Physique* **18**, 5–104 (1909).
7. Perrin, J. *Les atomes* (Librairie Félix Alcan, Paris, 1913).
8. Dardo, M. *Nobel Laureates and Twentieth-Century Physics*. 114–116. ISBN: 0521540089 (Cambridge University Press, 2004).
9. Maxwell, J. C. *On the Motions and Collisions of Perfect of Elastic Spheres*. *Philosophical Magazine* **19**. doi:10.1080/14786446008642818 (1860).
10. Gibbs, J. W. *On the equilibrium of heterogeneous substances*. *American Journal of Science* **96**, 441–458 (1878).
11. Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., *et al.* *Equation of State Calculations by Fast Computing Machines*. *The Journal of Chemical Physics* **21**, 1087. ISSN: 00219606 (1953).
12. Walther, A. & Müller, A. H. E. *Janus particles*. *Soft Matter* **4**, 663. ISSN: 1744-683X (2008).
13. Pawar, A. B. & Kretzschmar, I. *Fabrication, Assembly, and Application of Patchy Particles*. *Macromolecular Rapid Communications* **31**, NA–NA. ISSN: 10221336 (2010).
14. Rossi, L., Sacanna, S., Irvine, W. T. M., *et al.* *Cubic crystals from cubic colloids*. *Soft Matter* **7**, 4139–4142. ISSN: 1744-683X (2011).
15. Su, C., Wang, H. & Liu, X. *Controllable fabrication and growth mechanism of hematite cubes*. *Crystal Research and Technology* **46**, 209–214. ISSN: 02321300 (2011).
16. Zhang, Y., Lu, F., van der Lelie, D., *et al.* *Continuous Phase Transformation in Nanocube Assemblies*. *Physical Review Letters* **107**, 135701. ISSN: 0031-9007 (2011).
17. Quan, Z. & Fang, J. *Superlattices with non-spherical building blocks*. *Nano Today* **5**, 390–411. ISSN: 17480132 (2010).
18. Zhao, K. & Mason, T. G. *Twinning of Rhombic Colloidal Crystals*. *Journal of the American Chemical Society* **134**, 18125–18131. ISSN: 0002-7863 (2012).
19. Ye, X., Chen, J., Engel, M., *et al.* *Competition of shape and interaction patchiness for self-assembling nanoplates*. *Nature Chemistry* **5**, 466–473. ISSN: 1755-4330 (2013).
20. Ahmadi, T. S., Wang, Z. L., Green, T. C., *et al.* *Shape-Controlled Synthesis of Colloidal Platinum Nanoparticles*. *Science* **272**, 1924–1925. ISSN: 0036-8075 (1996).
21. Frenkel, D. *The simulation of entropic phase transitions*. *Journal of Physics: Condensed Matter* **6**, A71–A78. ISSN: 0953-8984 (1994).
22. Damasceno, P. F., Engel, M. & Glotzer, S. C. *Predictive Self-Assembly of Polyhedra into Complex Structures*. *Science* **337**, 453–457. ISSN: 0036-8075 (2012).
23. Agarwal, U. & Escobedo, F. A. *Mesophase behaviour of polyhedral particles*. *Nature Materials* **10**, 230–235. ISSN: 1476-1122 (2011).
24. Marechal, M., Cuetos, A., Martinez-Haya, B., *et al.* *Phase behavior of hard colloidal platelets using free energy calculations*. *The Journal of Chemical Physics* **134**, 094501. ISSN: 00219606 (2011).
25. Avendaño, C. & Escobedo, F. A. *Phase behavior of rounded hard-squares*. *Soft Matter* **8**, 4675. ISSN: 1744-683X (2012).
26. De Graaf, J., Filion, L., Marechal, M., *et al.* *Crystal-structure prediction via the Floppy-Box Monte Carlo algorithm: Method and application to hard (non)convex particles*. *The Journal of Chemical Physics* **137**, 214101. ISSN: 00219606 (2012).
27. Qi, W., Graaf, J. D., Qiao, F., *et al.* *Ordered Two-Dimensional Superstructures of Colloidal Octapod-Shaped Nanocrystals on Flat Substrates*. *Nano Letters* **12**, 5299–5303. ISSN: 1530-6984 (2012).

28. Batten, R. D., Stillinger, F. H. & Torquato, S. *Phase behavior of colloidal superballs: Shape interpolation from spheres to cubes*. Physical Review E **81**, 061105. ISSN: 1539-3755 (2010).
29. Marechal, M., Zimmermann, U. & Lowen, H. *Freezing of parallel hard cubes with rounded edges*. The Journal of Chemical Physics **136**, 144506. ISSN: 00219606 (2012).
30. Gantapara, A. P., de Graaf, J., van Roij, R., *et al.* *Phase Diagram and Structural Diversity of a Family of Truncated Cubes: Degenerate Close-Packed Structures and Vacancy-Rich States*. Physical Review Letters **111**, 015501. ISSN: 0031-9007 (2013).
31. Smallenburg, F., Filion, L., Marechal, M., *et al.* *Vacancy-stabilized crystalline order in hard cubes*. Proceedings of the National Academy of Sciences **109**, 17886–17890. ISSN: 0027-8424 (2012).
32. Singh, S. P. & Das, S. P. *A density functional model for the binary crystal of hard spheres with vacancies*. The Journal of Chemical Physics **126**, 064507. ISSN: 00219606 (2007).
33. Oettel, M., Görig, S., Härtel, A., *et al.* *Free energies, vacancy concentrations, and density distribution anisotropies in hard-sphere crystals: A combined density functional and simulation study*. Physical Review E **82**, 051404. ISSN: 1539-3755 (2010).
34. Bates, M. a. & Frenkel, D. *Influence of polydispersity on the phase behavior of colloidal liquid crystals: A Monte Carlo simulation study*. The Journal of Chemical Physics **109**, 6193. ISSN: 00219606 (1998).
35. Frenkel, D. *Order through entropy*. Nature Materials **14**, 9–12. ISSN: 1476-1122 (2014).
36. Van den Broeke, J. *Colloidal Systems of Hard Tilted Cubes*. (Unpublished bachelor's thesis). <<http://dspace.library.uu.nl/handle/1874/316776>> (2015).
37. Sugimoto, T. & Sakata, K. *Preparation of monodisperse pseudocubic α -Fe₂O₃ particles from condensed ferric hydroxide gel*. Journal of Colloid and Interface Science **152**, 587–590. ISSN: 00219797 (1992).
38. Zhao, K., Bruinsma, R. & Mason, T. G. *Entropic crystal-crystal transitions of Brownian squares*. Proceedings of the National Academy of Sciences **108**, 2684–2687. ISSN: 0027-8424 (2011).
39. Nolte, D. D. *The tangled tale of phase space*. Physics Today **63**, 33–38. ISSN: 0031-9228 (2010).
40. Frenkel, D. & Ladd, A. J. C. *New Monte Carlo method to compute the free energy of arbitrary solids. Application to the fcc and hcp phases of hard spheres*. The Journal of Chemical Physics **81**, 3188. ISSN: 00219606 (1984).
41. Wertheim, M. S. *Exact Solution of the Percus-Yevick Integral Equation for Hard Spheres*. Physical Review Letters **10**, 321–323. ISSN: 0031-9007 (1963).
42. Frenkel, D. & Smit, B. *Understanding molecular simulation: from algorithms to applications*. ISBN: 0-12-267351-4 (Academic press, 2011).
43. De Leeuw, S. W., Perram, J. W. & Smith, E. R. *Simulation of Electrostatic Systems in Periodic Boundary Conditions. I. Lattice Sums and Dielectric Constants*. Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences **373**, 27–56. ISSN: 1364-5021 (1980).
44. Bet, B., Boosten, G., Dijkstra, M., *et al.* *Efficient shapes for microswimming: from three-body swimmers to helical flagella*. arXiv [cond-mat.soft], 1–6 (2016).
45. Jacobs, W. M. & Frenkel, D. *Self-Assembly of Structures with Addressable Complexity*. Journal of the American Chemical Society **138**, 2457–2467. ISSN: 0002-7863 (2016).
46. Frenkel, D. *Simulations: The dark side*. The European Physical Journal Plus **128**, 10. ISSN: 2190-5444 (2013).
47. Hastings, W. K. *Monte Carlo sampling methods using Markov chains and their applications*. Biometrika **57**, 97–109. ISSN: 0006-3444 (1970).
48. Gilbert, E., Johnson, D. & Keerthi, S. *A fast procedure for computing the distance between complex objects in three-dimensional space*. IEEE Journal on Robotics and Automation **4**, 193–203. ISSN: 08824967 (1988).
49. Gottschalk, S. *Separating axis theorem* (1996).
50. Gottschalk, S, Lin, M. C., Manocha, D, *et al.* *OBBTree: A Hierarchical Structure for Rapid Interference Detection*, 171–180 (1996).
51. Parrinello, M. *Polymorphic transitions in single crystals: A new molecular dynamics method*. Journal of Applied Physics **52**, 7182. ISSN: 00218979 (1981).
52. Filion, L., Marechal, M., van Oorschot, B., *et al.* *Efficient Method for Predicting Crystal Structures at Finite Temperature: Variable Box Shape Simulations*. Physical Review Letters **103**, 188302. ISSN: 0031-9007 (2009).

53. Martoák, R., Laio, A. & Parrinello, M. *Predicting Crystal Structures: The Parrinello-Rahman Method Revisited*. Physical Review Letters **90**, 075503. ISSN: 0031-9007 (2003).
54. Vega, C. & Noya, E. G. *Revisiting the Frenkel-Ladd method to compute the free energy of solids: The Einstein molecule approach*. The Journal of Chemical Physics **127**, 154113. ISSN: 00219606 (2007).
55. Vega, C., Sanz, E., Abascal, J. L. F., *et al.* *Determination of phase diagrams via computer simulation: methodology and applications to water, electrolytes and proteins*. Journal of Physics: Condensed Matter **20**, 153101. ISSN: 0953-8984 (2008).
56. Noya, E. G., Conde, M. M. & Vega, C. *Computing the free energy of molecular solids by the Einstein molecule approach: Ices XIII and XIV, hard-dumbbells and a patchy model of proteins*. The Journal of Chemical Physics **129**, 104704. ISSN: 00219606 (2008).
57. Bruce, A. D., Wilding, N. B. & Ackland, G. J. *Free Energy of Crystalline Solids: A Lattice-Switch Monte Carlo Method*. Physical Review Letters **79**, 3002–3005. ISSN: 0031-9007 (1997).
58. Paik, T., Ko, D.-K., Gordon, T. R., *et al.* *Studies of Liquid Crystalline Self-Assembly of GdF₃ Nanoplates by In-Plane, Out-of-Plane SAXS*. ACS Nano **5**, 8322–8330. ISSN: 1936-0851 (2011).
59. Whitelam, S., Tamblyn, I., Beton, P. H., *et al.* *Random and Ordered Phases of Off-Lattice Rhombus Tiles*. Physical Review Letters **108**, 035702. ISSN: 0031-9007 (2012).
60. Bruce, A. D. & Wilding, N. B. *Computational strategies for mapping equilibrium phase diagrams*. arXiv [cond-mat.soft], 45 (2002).
61. Van Anders, G., Klotsa, D., Karas, A. S., *et al.* *Digital Alchemy for Materials Design: Colloids and Beyond*. ACS Nano **9**, 9542–9553. ISSN: 1936-0851 (2015).
62. Singer, S. J. & Mumaugh, R. *Monte Carlo study of fluidplastic crystal coexistence in hard dumbbells*. The Journal of Chemical Physics **93**, 1278. ISSN: 00219606 (1990).
63. Jaynes, E. T. *Information Theory and Statistical Mechanics. II*. Physical Review **108**, 171–190. ISSN: 0031-899X (1957).
64. Bennett, C. H. *Efficient Estimation Of Energy Differences from Monte Carlo Data*. Journal of Computational Physics **22**, 245–268. ISSN: 00219991 (1976).
65. Dijkstra, M., van Roij, R. & Evans, R. *Phase diagram of highly asymmetric binary hard-sphere mixtures*. Physical Review E **59**, 5744–5771. ISSN: 1063-651X (1999).
66. Fillion, L. *Self-assembly in colloidal hard-sphere systems*. Thesis. <<http://igitur-archive.library.uu.nl/dissertations/2011-0119-200314/UUindex.html>> (2011).
67. Blunt, M. O., Russell, J. C., Gimenez-Lopez, M. d. C., *et al.* *Random Tiling and Topological Defects in a Two-Dimensional Molecular Network*. Science **322**, 1077–1081. ISSN: 0036-8075 (2008).