The kinetics of phase transitions in polydisperse systems

John Joseph Williamson
School of Physics and Astronomy, School of Mathematics
University of Leeds

Submitted in accordance with the requirements for the degree of
Doctor of Philosophy
February 2013

Declaration

The candidate confirms that the work submitted is his own, except where work which has formed part of jointly-authored publications has been included. The contribution of the candidate and the other authors to this work has been explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others.

Chapters 5 and 6 are based respectively on work in Refs. [Williamson and Evans, 2012] and [Williamson and Evans, 2013], co-authored with R. M. L. Evans (the supervisor), with the candidate as first author.

Work attributable to the candidate:

- Design, programming and deployment of simulation code
- Data analysis, creation of figures
- Interpreting and explaining results, supporting theoretical calculations
- Preparation of draft (full) manuscripts
- Incorporation of suggested changes to manuscripts.

Work attributable to R. M. L. Evans:

- Initial project design, including use of Kinetic Monte Carlo approach (Chapters 5 and 6) and strategy for template and simulation geometry for crystal growth (Chapter 6)
- Original development of perturbative theory of fractionation (Chapter 5) and kinetic theory of split interface growth (Chapter 6)
- Guidance in theoretical calculations
- Discussion and interpretation of results
- Critical reading of and suggestion of alterations to draft manuscripts.

This thesis is dedicated to my parents, of whom I am incredibly proud, in the hope that they will at least understand this page.

Acknowledgements

Thanks to my supervisor Mike Evans for showing me what science can be, and what a scientist can be. Mike's inspiration, guidance, and most importantly his example have been of more benefit to me than any thanks could express. But anyway, thanks Mike!

Thanks to a number of people in the Soft Matter community including Paul Bartlett, Adrian Baule, Robyn Cooke (née Moorcroft), Siobhan Liddle, Johan Mattsson, Wilson Poon, Paddy Royall, Peter Sollich, Patrick Warren, and Nigel Wilding for (variously) drinks, discussion, and inspiration. Discussions with Wilson and Nigel first helped to make me aware of the significance of some initial findings that formed the basis for the work in Chapter 5. I thank Johan for guidance in the interpretation of glassy intermediate scattering functions, discussed in Chapter 4. Nigel, Paul and Paddy's interest in the early results of the work in Chapter 6 encouraged me to pursue a paper on this topic and to test the hypothesis of dynamical facilitation of fractionation. Thanks, thereafter, to Nigel for suggesting the connection between my findings on local size correlations and his earlier work on similar correlations in the presence of spatial homogeneities, which I have cited in that chapter. Ongoing discussions with Siobhan and her supervisor Wilson have been very helpful in the preparation of Chapter 7, and in keeping me abreast of their excellent work on polydispersity. I thank them for providing some experimental data for comparison.

Thanks to Glenys Bowles, and more recently to Jeanne Shuttleworth, for help and support in all matters administrative.

Thanks to my friends at the University of Leeds, especially to the Leeds University Union Big Band for providing too much distraction. Thanks to Tom Welsh and Craig Hall for sharing an office with me and always having time to help, and remembering things I forgot. Thanks to Joe Pearson for first introducing me to Mercurial, which has been very useful in managing the simulation code. Thanks to Andy Chisholm and Tom West at Durham for supporting and inspiring me through my undergraduate degree.

Thanks to my family, for everything.

Abstract

The study of soft matter's phase behaviour is based on thermodynamics, originally developed to describe systems i) composed of identical particles, and ii) in their final equilibrium state. However, a practical understanding requires knowledge of how real systems do (or do not) *actually approach* equilibrium. This is especially difficult to achieve when, as often in soft matter, the constituents are *polydisperse*, i.e. comprise continuously non-identical particle species.

I present a wide-ranging simulation study of phase transition kinetics in the presence of polydispersity, in the context of model colloidal systems. After briefly exploring the structural and dynamical physics of polydisperse systems, I show that fractionation (the partitioning of a polydisperse property between phases) may be enacted in the very early stages of phase separation, and highlight the qualitative sensitivity of this effect to the details of inter-particle potentials.

I study the effects of metastable gas-liquid separation on crystal growth, finding a complex dependence on polydispersity which I explain with novel fractionation and local size correlation measurements.

I test a theory of fractionation against experimental data in a colloid-polymer mixture with small polymers, a regime in which the widely-used Mean-Field Asakura-Oosawa (MFAO) model becomes unphysical, and find that qualitative agreement can be obtained via a simple modification of the MFAO theory.

I precisely measure the composition of a diffusively-grown hard sphere crystal with small polydispersity. The results are agnostic about a prediction that diffusion *induces* nonequilibrium fractionation, but do show that equilibrium composition is not achieved: to within extremely small error bars, the crystal does not fractionate at all during growth.

I examine crystal growth on an epitaxial substrate composed of dual crystal templates.

Finally, I study the interdependent diffusion of particle size and concentration in a polydisperse hard sphere fluid, isolating the eigenmodes implied by the BMCSL polydisperse free energy.

Abbreviations and units

ϕ	Volume fraction
ρ	Number density
σ	Polydispersity
d	Particle diameter
Δ	Mean Brownian step size
k_B	Boltzmann's constant
k_BT	Thermal energy
L	Simulation box size
N	Number of particles
q	Wavenumber
T	Temperature
t	Time
BMCSL	Boublík-Mansoori-Carnahan-Starling-Leland free energy
C-P	Colloid-polymer (mixture)
C-S	Carnahan-Starling
FCC	Face-centred cubic
G-L, C-G etc.	Gas-liquid, crystal-gas etc.
HI	Hydrodynamic interactions
MC	Monte Carlo
MFAO	Mean-field Asakura-Oosawa
MSD	Mean-squared displacement
RCP	Random close packing
RDF	Radial distribution function

Note: these are the most commonly used abbreviations in this thesis. Others appear in particular chapters, but are always introduced and described at the point of use.

Note: scaled units, described in the text, are noted here for reference. The unit of distance is the mean hard-core particle diameter. The unit of time is the mean time for such a mean-sized particle to diffuse a distance equal to its diameter, in the dilute limit. The thermal energy k_BT is set to unity.

Contents

List of Figures

List of Tables

> “Truth springs from argument amongst friends.”
>
> Hume

> “To have spent so long learning so relatively little, and then to be menaced in every aspect of my life by people who already know everything, and who have all the information they need . . . ”
>
> Hitchens

> “Starnahan-Carling”
>
> Evans

Chapter 1

Introduction

1.1 Colloidal dispersions

Colloids consist of particles of mesoscopic size (between 10nm and 1μm) dispersed in some kind of fluid solvent. This broad definition covers a wide variety of natural and synthetic substances including milk, blood platelets, viruses, cosmetics and pharmaceuticals, paints and radioactive waste. Colloidal dispersions are a paradigmatic example of *soft condensed matter* – substances in which, relative to atomic or molecular systems, the constituents are large, slow and weakly held. Other examples include proteins and polymers: essentially, soft matter is anything which is in some sense 'squishy'.

Due to the relatively large length scales involved, colloids can be treated without recourse to quantum mechanics. (The term 'colloid' is properly used to refer to the system as a whole, but is often used to describe the individual particles, which are properly called 'latices'.) However, they are still small enough to be subject to thermal agitation from the solvent in which they are suspended. For this reason, thermodynamics and statistical mechanics are the natural theoretical tools to deal with colloids. These tools enable us to input the system parameters (particle interactions, density, temperature etc.) and to get out information about the equilibrium state of the system – the (macro)state it would eventually achieve if allowed to fully explore all possible microscopic configurations, or microstates.

The pertinent theoretical apparatus is familiar from simpler atomic and molecular substances, for which thermodynamics and statistical mechanics were originally developed. Thus, as well as having their own importance in e.g. industrial, technological and medical applications, colloidal systems are seen as a useful testing ground for general theories of fluid matter [Suresh, 2006; Wette *et al.*, 2009]. Given the fact that colloidal interaction potentials can be easily and precisely customised [Russel, 2003; Yethiraj, 2007; Yethiraj and Van Blaaderen, 2003], and the relatively forgiving length and time scales involved, it is easy to see why. However, it is important not to ignore the differences between colloids and simpler systems, because those differences may limit the generality of results obtained through the study of colloidal systems, and certainly do limit the theorist's license to apply well-understood theories of classical fluids to colloids, at least without careful consideration or modification of those theories. Colloidal dispersions – like soft matter in general – present their own challenges, as well as inheriting those of simpler systems.

Two of these challenges in particular are the focus of the present work: polydispersity (variation of properties between particles) and phase transition kinetics (the often complex kinetic pathways by which systems achieve or fail to achieve their equilibrium state).

1.2 Challenges

1.2.1 Polydispersity

In, for instance, a sample of pure water, it is true to say that each molecule of H_2O is, in a strict sense, identical. This is usually not the case in soft matter. A collection of mesoscopic particles will almost always exhibit continuous variation in properties such as size, charge, shape or chemical makeup. This variation, referred to as *polydispersity*, hugely complicates the study of such systems, introducing peculiar phenomena and necessitating a more sophisticated theoretical treatment than monodisperse (i.e. not polydisperse) systems.

1.2.2 Dynamics and phase transition kinetics

A calculation of the equilibrium state of a system is a vital prerequisite for predicting its phase behaviour, but is not sufficient alone. One must also consider the *kinetics* of how that equilibrium state might be achieved – from a dynamical point of view, how will the system go about arranging its constituents in such a way as to reach equilibrium? Systems may become trapped in a metastable (quasi-equilibrium) state and need to wait for a nucleation event before the equilibrium phase can grow; they may form kinetically frustrated glasses or gels, in which particles essentially cannot move and are therefore unable to explore their configuration space, never discovering the equilibrium state; they may be in a statistically steady state subject to driving by an external force. Nonequilibrium statistical mechanics, the adaptation of statistical mechanics to deal with these situations, is, comparatively speaking, in its infancy, with relatively few generally-applicable results [Evans, 2005; Jarzynski, 1997].

It is necessary to consider the kinetic pathway *towards* equilibrium, as well as the final state, in order to adequately describe and predict the behaviour of real systems. This task is even further complicated when polydispersity is present. The equilibrium phase diagram becomes more complex, as do the dynamics of many different species of particle undergoing their own subtly different stochastic motion – an accurate theoretical description of the behaviour observed can therefore become elusive.

1.3 Applications of colloid physics

To speak simply of the ‘applications of colloids’ would somewhat miss the point, implying that colloids are some newly-developed class of materials which reside in academic laboratories waiting to be put to use. In fact, a huge variety of everyday and not-so-everyday substances are colloidal in nature, so that an understanding of colloid physics is of direct relevance to the handling, manipulation and exploitation of those substances. In this section I discuss a few examples by way of motivation. This selection is not even close to being exhaustive, but serves to illustrate the range of disparate fields in which colloid physics is useful.

1.3.1 Medicine

The most obvious connection between colloids and medicine is that many biological substances are colloidal or are closely related to colloids, by virtue of their size, the environment they occupy, and the forces acting on them (being mesoscopic, suspended in solvents and subject to thermal fluctuations therein, with relatively weak interaction energies). Examples include viruses [Schmidt *et al.*, 2010] (in fact, the Tobacco Mosaic virus is often used in experiment as a model colloid) and globular proteins [Poon, 1997]. The crystallisation of proteins for the purposes of studying their structure is an important topic and one to which insights from colloid science are directly applicable, since the phase diagram of an appropriate colloid-polymer mixture can be made to closely resemble that of proteins [Anderson and Lekkerkerker, 2002]. The 'patchy' (i.e. directional) interactions often present in proteins are a key motivation behind a recent surge in interest in patchy colloids [Doye *et al.*, 2007].

As well as being applicable to substances already present in biology, colloids have applications in drug delivery [Shilpi *et al.*, 2007]. For example, colloidosome microcapsules, formed by the self-assembly of colloid particles encasing an emulsion droplet, have emerged as a way of encapsulating poorly-soluble drugs for intelligent delivery into the body [Dinsmore *et al.*, 2002].

1.3.2 Photonic crystals

Photonic crystals are those with lattice parameters such that they develop an 'optical bandgap' analogous to the electronic bandgap in semiconducting materials. Colloids, which are typically of the correct size to form crystals that scatter visible light – the reason opals display rainbow colours – are an ideal candidate for the production of such crystals, serving either as the photonic crystal itself or as a mould around which an inverted photonic crystal is created [Allard and Sargent, 2004; Colvin, 2001; Stein *et al.*, 2013]. It is hoped that these optical semiconductors will enable the creation of optical computers, whereby the electrons in traditional silicon chips are replaced with photons being manipulated by photonic crystals. In this application in particular, polydispersity is a vital consideration because its presence makes growing high quality crystals difficult. Even very mild polydispersity has been shown to significantly affect the properties of photonic crystals [Allard and Sargent, 2004].

1.3.3 Cosmetics, pharmaceuticals, sewage

Colloidal gels, creams and dispersions are common in the cosmetic [Tadros, 2008] and pharmaceutical industries [Tadros, 2007]. Developing, storing and processing these products on an industrial scale requires an understanding of their phase behaviour in order to ensure stability (i.e. shelf life) and efficacy. Less inviting but arguably even more important is the field of sewage treatment [Hiemenz and Rajagopalan, 1997], in which a number of biological and technological colloidal species must be removed during the purification process, for instance by electrophoresis of charged species.

1.4 Aims and layout

The main aim of the present work is to explore the kinetics of phase transitions in polydisperse systems, particularly as relevant to soft matter, by the study of model colloids. Dynamical simulations – those in which an effort is made to accurately represent the dynamics of the system – provide a 'virtual experiment' in which the bulk behaviour of real systems is reproduced but in which the precise properties of every particle at every time are accessible to the experimenter. Building on existing work on the equilibrium effects of polydispersity, I aim to investigate how polydisperse systems progress towards equilibrium, highlighting the importance of phase transition kinetics and dynamics in the phase behaviour of real systems.

The thesis is laid out as follows: in Chapter 2 I outline the basic thermodynamic and dynamical considerations relevant to colloidal systems. In Chapter 3 I introduce the concept of simulation in the study of condensed matter, with particular reference to the traditional Monte Carlo method from which my simulation method and many others are descended. This lays the ground for Chapter 4 in which calibration results to measure the basic structural and dynamical aspects of the simulation are presented. Chapter 5 is a study of fractionation during the gas-liquid phase separation of a polydisperse fluid. Chapter 6 concerns the effects of gas-liquid separation and polydispersity on the growth of a crystalline phase. Chapter 7 takes a detour into colloid-polymer mixtures, demonstrating the poor behaviour of a well-known model in the regime of small polymer size, improving the model, and comparing the consequent predictions for fractionation strength with some new experimental data. Chapter 8 concerns a precise characterisation of slightly polydisperse hard sphere crystals, performed in order to try and detect a signal from diffusion-induced fractionation. In Chapter 9 I study the effects on crystal growth of a templating strategy whereby two different and precisely controlled templates are placed side by side. Finally, in Chapter 10, I investigate theoretically the interdependent diffusion of fields describing local particle size and local concentration, within the BMCSL polydisperse hard sphere free energy. I conclude in Chapter 11.

1.5 Publications

The work presented in this thesis has resulted in two papers, co-authored with my supervisor R. M. L. Evans. Chapters 5 and 6 are based on these papers, Refs. [Williamson and Evans, 2012] and [Williamson and Evans, 2013] respectively. Details of the contribution of each of us to these papers is given in the Declaration at the start of this thesis.

Chapter 2

Thermodynamics and dynamics of colloids

I now briefly review the necessary thermodynamic concepts for discussing colloidal and soft matter phase behaviour, and how those concepts have been extended to account for polydispersity, with particular reference to the size-polydisperse hard sphere model. I then discuss some relevant dynamical considerations, including phase transition kinetics, frustration and arrest.

A large number of sources are available discussing the essentials of colloid physics. See e.g. Refs. [Russel, 1987; Russel *et al.*, 1992], which were particularly useful in writing this chapter.

2.1 Thermodynamics

2.1.1 Macroscopic state variables

The key principle of the thermodynamic approach is that a system comprising a thermodynamically large ($\sim 10^{23}$) number of identical particles, subject to stochastic thermal fluctuations, can be adequately described by *state variables* pertaining to bulk properties such as number density ρ, temperature T and pressure P, without knowledge of the individual trajectories of the particles involved. This 'black-boxing' of the system, whereby useful macroscopic properties are related to one another independently of the precise nature of the system's components, allows a highly general theoretical framework to be formulated. Indeed, the pioneers of this approach were able to usefully describe and predict the behaviour of systems such as steam engines long before anything was known about what the constituent molecules of steam actually were.

These state variables can be related to one another via *equations of state* such as the ideal gas law, $PV = nRT$ (in which V is the volume, n the number of particles and R a constant). This apparently simple equation embodies the power and utility of thermodynamics: even without knowing what particles a system is made of, if the system can be approximated as an ideal gas, the ideal gas law predicts how the pressure of the system will change if we increase its temperature by some amount, for example.

Thermodynamic variables may be classified into intensive and extensive variables. Intensive variables are those that do not change when two materials with the same values of that variable are brought together (e.g. temperature, density), while extensive variables do (e.g. energy, number of particles).

2.1.2 Thermodynamic potentials

State variables may be combined into *thermodynamic potentials* which can then be used to find the equilibrium state of the system. The Second Law of thermodynamics states that, in the absence of constraints on extensive variables (such as V, T or the internal energy U), a system at equilibrium will maximise its entropy S. When some extensive variables *are* held constant, other potentials may be derived and extremised to find the equilibrium state. Each potential is appropriate to a certain *ensemble*, a statistical set of microstates in which some variables are held constant while others can vary. For our purposes, the most useful of these is the Helmholtz free energy, the appropriate potential for the canonical ensemble (in which V, T and the number of particles N are held constant), since the volume of the solvent in which colloidal particles are suspended fixes the volume of the system. The Helmholtz free energy F (hereafter referred to simply as 'free energy') is given by:

$$F = U - TS \tag{2.1}$$

and is minimised when a system in the canonical ensemble is at equilibrium. The value of any particular thermodynamic variable g at equilibrium may be found by minimisation of F with respect to that variable, $\partial F/\partial g = 0$.

2.1.3 Phases and coexistence

Substances can exist in different *phases*, distinct from one another by sudden qualitative changes in some order parameter when a control parameter (e.g. T) is varied. For instance, density ρ is an order parameter distinguishing gas from liquid. Crystalline and fluid phases can be distinguished also by the bond order parameter Q_6 (described in further detail in Chapter 3) which measures similarities in the local neighbour structure between particles. Phase transitions are classified as 'first-order' if they are accompanied by a discontinuity in the first derivative of F with respect to the control variable, or 'continuous' if the discontinuity is in higher derivatives. The phase transitions encountered in this thesis are first-order.

Systems often reach equilibrium not by switching wholesale from one phase to another, but by attaining a coexistence of two or more phases, if such a coexistence results in a lower free energy than can be achieved by any one phase alone. The phases (which need not necessarily be continuous) are then separated by interfaces, i.e. regions of space in which the order parameters associated with the phases change abruptly.

The way in which (for instance) gas-liquid phase coexistence may be thermodynamically desirable is easily visualised by plotting the free energy of an appropriate fluid as a function of ρ. We move to the free energy density $f = F/V$, an intensive variable, to simplify the discussion. FIG. 2.1 shows f as a function of ρ for a fluid which can separate into gas and liquid phases. This illustrative example is for a fluid of the type discussed in Chapter 5, in which particles consist of a hard core surrounded by an attractive square well. At equilibrium, temperature, being a *field variable*, must be the same throughout the system. In fact, in this example, the attraction

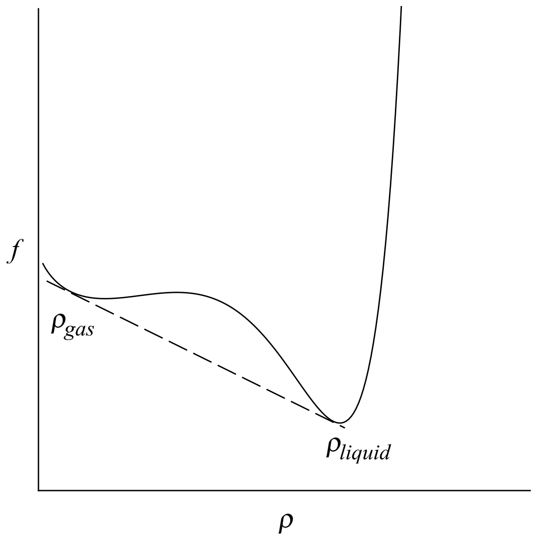

Figure 2.1: An example fluid free energy curve exhibiting double minimum structure. A common tangent (dashed) exists between ρ_{gas} and ρ_{liquid}, indicating a possible phase coexistence.

strength u sets an 'effective temperature' $T_{\mathit{eff}} = 1/u$ which is, by construction, the same for all particles in the system. The attraction, if strong enough, results in a 'double minimum' structure in the free energy curve, with cotangential points (i.e. points that can support a common tangent) at densities ρ_{gas} and ρ_{liquid}. A system whose 'parent', or overall, density ρ_p sits between these points can form some amount of gas and liquid phases while preserving its overall ρ_p according to the lever rule, which follows from conservation of material:

$$\rho_{gas}V_{gas} + \rho_{liquid}V_{liquid} = \rho_p V \tag{2.2}$$

where V_{gas} and V_{liquid} are the volumes of the gas and liquid phases formed. The shared tangent means that the slopes and f-axis intercepts of the curve at these points are equal. The slope, $\partial f/\partial\rho = \mu$ is the chemical potential of the phase, while the intercept gives the negative pressure $-P$. That is, the existence of the common tangent implies equality of μ and P across an interface between the two phases, so that they can coexist.

Labelling the corresponding free energy densities of the phases f_{gas} and f_{liquid}, the total free energy density $f_{total} = (f_{gas}V_{gas} + f_{liquid}V_{liquid})/V$ sits somewhere along the common tangent between the gas and liquid points. By definition, this tangent specifies a lower value of f than does the homogeneous free energy curve at any point between ρ_{gas} and ρ_{liquid}, so that the coexistence results in a lowering of the system's free energy and, if no further lowering is possible, is referred to as the equilibrium state of the system.

I noted that the presence of a double minimum structure in the fluid free energy requires an attractive component to the inter-particle potential. In FIG. 2.2, the effect on the free energy curve with decreasing u is shown. It is clear that as u is decreased (i.e. T_{eff} increases), the coexistence densities move closer together, eventually resulting in a *critical point* at which the distinction between gas and liquid disappears.

Phase diagrams may be plotted to show the behaviour of the coexistence densities as u varies.

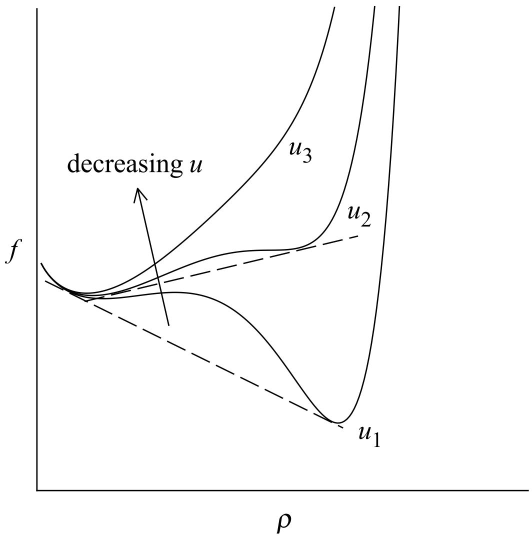

Figure 2.2: Behaviour of the fluid free energy as attraction strength u decreases. The cotangential densities move closer together, reaching a critical point where gas and liquid are no longer distinct phases.

The phase diagram corresponding to FIG. 2.2 is shown in FIG. 2.3. The gas and liquid phases are separated from the coexistence region by *binodals*. Any system prepared within the coexistence region can reduce its free energy by separating along the tie lines into gas and liquid regions whose densities are given by these binodals. Because u must be equal in the two phases, the tie lines are horizontal in this case.

It is worth noting in passing that adding a constant term, or one linear in ρ, to the free energy curve does not alter the resultant phase behaviour – the common tangent construction will give the same results. In calculations, such terms are often neglected, so the 'double minima' need not actually be minima. The only requirement for coexistence is the existence of cotangential points.

2.1.4 Metastability and instability

The fate of a system whose density sits between ρ_{gas} and ρ_{liquid} depends to a large extent on the shape of the free energy curve within that region. It is clear that in order for the cotangential points to exist, there must be a region where $\partial^2 f/\partial\rho^2 < 0$. If ρ_p falls within this region, the system is unstable to local fluctuations in ρ. This is because, when a region increases its density to $\rho + a$, causing a neighbouring region's density to fall to $\rho - b$, the combined free energy falls below the homogeneous curve. This is the case even if the fluctuation terms a and b are arbitrarily small. Such fluctuations are therefore thermodynamically favourable and are magnified, further lowering the free energy. Since these fluctuations naturally occur in the system anyway, phase separation begins immediately and takes place throughout the system, typically with some preferred initial length scale. This is called spinodal decomposition.

If a system is in a phase or phase coexistence from which the free energy can still be lowered,

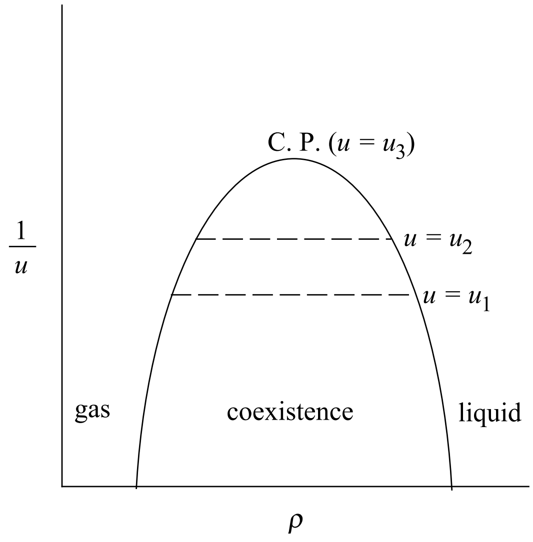

Figure 2.3: Phase diagram corresponding to the free energies in FIG. 2.2. The changing coexistence densities ρ_{gas} and ρ_{liquid} trace out binodals. C. P. indicates the critical point. Tie lines (dashed) indicate the densities of the phases into which a system prepared within the coexistence region will separate.

but is not unstable against arbitrarily small fluctuations in local density ($\partial^2 f/\partial\rho^2 > 0$), it is referred to as *metastable*. In this case, a nucleus of critical size must be formed to initiate phase separation. To see this, we can consider that the free energy contribution of a region of a new, equilibrium phase has two components – a volume free energy term (which would be the only one present if the new phase existed on its own and in bulk) and an interfacial term arising from the interface that must be created between this region and the metastable phase around it. Approximating the nucleus as spherical, so that volume scales as R^3 and surface (interfacial) area as R^2, we can write:

$$\Delta F_{nucleus} = -CR^3 + DR^2 \tag{2.3}$$

where C and D are positive constants, for the change in free energy brought about by a nucleus of radius R. The negative sign in the volume term means that the new phase, in bulk, has lower free energy than the metastable phase – that is why it is an equilibrium phase. However, it is counteracted by a positive term reflecting the fact that interfaces between different phases cost free energy. When R is small, the interfacial term wins: growing the nucleus costs free energy, so it will dissolve. At a critical value $R_{crit} = 2D/3C$, $\Delta F_{nucleus}$ has a maximum ($\frac{\partial \Delta F_{nucleus}}{\partial R} = 0$) so that further growth *reduces* the free energy of the nucleus and thus continues until the required amount of the new phase is formed. This competition between volume and interfacial terms also tells us that, when phase separation is initiated instead by spinodal decomposition, the initially small length scale of the phases will grow or 'ripen' until bulk phase coexistence is achieved. The points on the free energy curve separating these two regimes (i.e. where $\partial^2 f/\partial\rho^2 = 0$), when plotted on the phase diagram, trace out *spinodals*, which sit inside the binodals and enclose a region where spinodal decomposition is possible, as shown in FIG. 2.4.

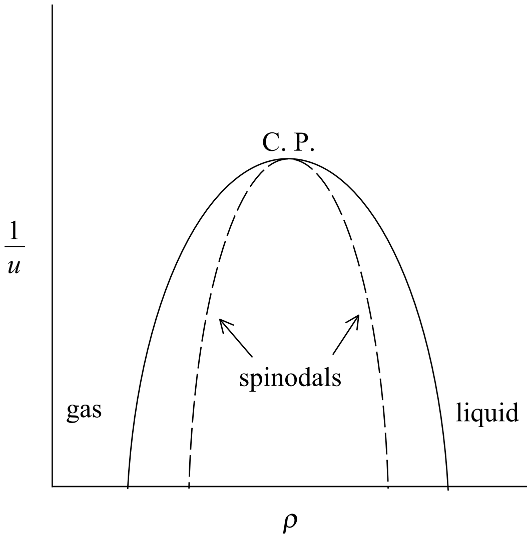

Figure 2.4: Phase diagram including spinodals (dashed), which trace out the densities where $\partial^2 f/\partial\rho^2 = 0$. Within the spinodals, phase separation may proceed by spinodal decomposition.

Metastable states play an important role in phase transition kinetics, and are discussed in depth in Chapter 6, while spinodal decomposition features in Chapter 5. It should be pointed out that the distinction between spinodal decomposition and nucleation is not sharp. In nucleation, the energy cost of producing a critical nucleus may be less than the thermal energy $k_B T$ – where k_B is Boltzmann's constant – so that nucleation events take place very easily. Therefore the boundary between spinodal-like and nucleation-like growth is gradual. Additionally, the discussion above is based on Classical Nucleation Theory which makes a number of rather drastic assumptions, including the assumption that 'bulk free energy' can be sensibly spoken of in a nucleus which may be only a few particles in size. Nevertheless, it serves to illustrate the kinds of phase transition kinetics we can expect to see, and how they relate to the free energy landscape of the system.

2.1.5 Application to colloids

The thermodynamical framework outlined so far was developed originally in the study of simple (i.e. molecular) fluids such as water, steam, air, and so on. By comparison, the application of thermodynamics to complex fluids such as colloids is, relatively speaking, new [Onsager, 1949]. The analogy permitting this borrowing of theory is that colloids, like molecules in a simple fluid, are subject to thermal fluctuations. Although far larger than simple molecules, the dispersed particles are small enough that they remain suspended in their solvent for appreciable times against the force of gravity, and in terms of their micro-dynamics are subject to random thermal 'kicks' from the surrounding solvent – they perform Brownian motion.

The solvent, since it is composed of particles far smaller than the colloids suspended in it (see Section 2.3), is usually treated as simply a background continuum through which the colloids move. As discussed in Section 2.3.2, this is not strictly valid – hydrodynamic forces exist between colloid particles, transmitted by the solvent. Since these interactions are velocity dependent, they do not influence the equilibrium phase diagram, but do influence particle dynamics. However, notwithstanding such complexities, colloidal fluids behave in a way more or less analogous to ordinary fluids, permitting the application of similar statistical mechanical and thermodynamic theories.

The structures exhibited in colloidal systems depend on the interplay of a wide variety of interaction potentials which can be present, including van der Waals, excluded volume and depletion potentials. The depletion attraction is of particular note. It is an effective attraction mediated by some small species (such as polymer coils) dispersed in the solvent along with the colloid [Lekkerkerker *et al.*, 1992], and as such can be very precisely tuned by controlling the properties of the depletant species [Russel, 2003; Yethiraj, 2007; Yethiraj and Van Blaaderen, 2003]. Strong enough attractions may lead to the formation of nonequilibrium gel-like networks, for example, while weaker attractions lead to equilibrium or metastable gas-liquid coexistence. Interaction energies are often weak enough that the internal energy and entropy terms in F can combine to yield complex phase behaviour sensitive to both factors. In fact, as we see in Section 2.2.1, even entropy alone can be enough to cause a phase transition.

The relative weakness of inter-particle bonding (along with relatively low number densities and therefore low energy densities) makes soft matter soft − in contrast to an atomic crystal such as diamond, colloidal crystals can often be ‘shear melted’ simply by shaking the container. Tunability of interactions, ease of rehomogenisation, large sizes and slow speeds all make colloids very attractive as an experimental system in which a wide variety of behaviour can be studied with comparative ease [Suresh, 2006; Wette *et al.*, 2009].

However, for a variety of reasons (e.g. their large size, the fact that the solvent is not a truly inert background) the physics of colloids is not directly comparable to that of molecules. Colloidal systems present unique challenges, the elucidation of which is, broadly speaking, the aim of this work. In the next section I discuss one in particular, polydispersity.

2.2 Polydispersity

A key feature distinguishing colloids from simple molecules is polydispersity. In order to introduce this I will, having discussed the concepts of phase transitions and coexistence with reference to the gas-liquid example, simplify matters (initially, at least) by moving to the famous ‘hard sphere’ model of fluids. In colloids and much of soft matter, some amount of polydispersity is unavoidable. So, even a system of colloidal particles which are nominally ‘the same’ will take their properties (e.g. particle size) from a probability distribution of finite width. The polydispersity σ of a distribution in some property is defined as the standard deviation of the distribution in units of its mean.

2.2.1 The hard sphere model

The hard sphere model captures the effect of the steric (excluded volume) interaction, which is extremely important in determining the structural behaviour of fluids. In the model, particles

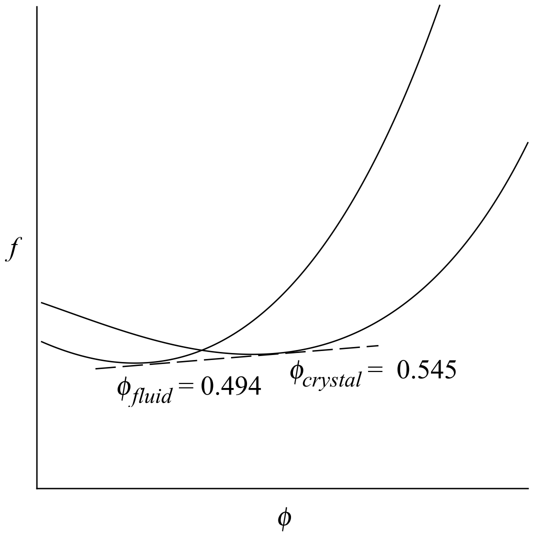

Figure 2.5: Free energy landscape for monodisperse hard spheres. In contrast to FIG. 2.1, there is no gas-liquid coexistence. Instead the crystal, which inhabits its own free energy branch, coexists with the fluid. Recall that adding to f constant terms or those linear in ϕ does not affect phase behaviour – only the positions in ϕ, not f, of the cotangential points are informative.

interact via a step-like repulsive potential which is infinite at contact and zero everywhere else:

$$V_{HS}(r) = \begin{cases} \infty & \text{if } r \leq d \\ 0 & \text{if } r > d \end{cases} \tag{2.4}$$

where d is the hard sphere diameter. Despite – and because of – its simplicity, the hard sphere model has found wide application in the study of e.g. the topological structure of amorphous fluids [O'Malley and Snook, 2005], crystal formation, and the glass transition [Zaccarelli *et al.*, 2009]. Furthermore, it is possible to synthesise colloids which behave as near-perfect hard spheres [Schöpe *et al.*, 2007], to which the hard sphere model is obviously a particularly close approximation.

Due to the nature of the potential, all allowed configurations of a hard sphere system have zero potential energy. The (purely kinetic) internal energy U is therefore proportional to T, and $F = U - TS$ is minimised simply by maximising the entropy S. This means that the equilibrium state is determined purely by entropic effects and, for monodisperse (identically sized) spheres, the phase diagram has only one axis, namely the volume fraction $\phi = \pi d^3 \rho / 6$. Monodisperse hard spheres exist as a disordered fluid up to $\phi \approx 0.494$, then as coexisting fluid and FCC crystal phases up to $\phi \approx 0.545$, beyond which the equilibrium configuration is fully crystalline [Russel *et al.*, 1992]. This is because, as the fluid's volume fraction is increased, particles have less and less free volume entropy due to caging by their neighbours. Beyond $\phi \approx 0.494$, it becomes preferable for some of the particles to form a crystal: they lose configurational entropy (because the crystal lattice imposes an ordered arrangement of particles) but gain free volume entropy, because the crystal packs particles more efficiently, giving each particle more free volume than if it were in a highly concentrated fluid. The question of whether this purely entropy-driven fluid-crystal transition could exist was one of the first to be settled by molecular simulation [Alder and Wainwright, 1957].

Because there is no attractive part to the potential, there is no gas-liquid coexistence – the fluid branch of the free energy has no cotangential points. Rather the crystal phase, having a different structure to the fluid, inhabits its own free energy branch, as shown in FIG. 2.5. Accurate and widely used phenomenological free energies for the hard sphere system are given by Carnahan and Starling in the fluid [Carnahan and Starling, 1969] and Hall in the crystal [Hall, 1972]. However, we will see in Chapter 7 that the Carnahan-Starling expression in particular leaves something to be desired with respect to its behaviour at high fluid densities.

2.2.2 Size-polydisperse hard spheres

Polydispersity can be introduced into the hard sphere model by allowing the particle diameters d_i to vary. The ubiquity of this 'size-polydispersity' and the ease with which it can be incorporated into the hard sphere model mean that it is probably the most widely studied form of polydispersity – the present work follows in this tradition. However, much of the theoretical and conceptual apparatus is common also to polydispersity in charge, chemical makeup and so on [Evans, 2001]. The polydispersity can be made as low as $\sigma \approx 0.02$ with careful, specialised synthesis [Van Blaaderen and Vrij, 1992; Velikov and van Blaaderen, 2001], but is usually higher [Butter *et al.*, 2005; Hiramatsu and Osterloh, 2004]. Also, the details of production methods might limit to what degree the polydispersity can be reduced, assuming of course that such reduction is desired, which is not always the case.

Even the equilibrium physics of the hard sphere model is greatly complicated and enriched by polydispersity. For a finite mixture of M different species, the free energy landscape gains a new dimension for each corresponding species density variable ρ_m, where $1 \leq m < M$. For instance, in a binary mixture ($M = 2$) there are two density axes. The composition of coexisting phases must now be described in terms of both ρ_1 and ρ_2, which involves minimising f with respect to both variables. In a truly polydisperse system, each particle belongs, properly speaking, to its own species so that $M = \infty$. The free energy landscape therefore has an infinite number of axes, one for each size species' individual density. The free energy becomes a *functional* of the density distribution $\rho(d)$ – calculating this functional in the first place, so that the minimisation procedure can be performed, is in itself a difficult challenge.

2.2.3 Dealing with polydispersity

Various methods of overcoming the obstacles outlined above have led to qualitatively distinct predictions [Bartlett and Warren, 1999; Fasolo and Sollich, 2004; Sollich and Wilding, 2010]. Peter Sollich has written a detailed review of the various approaches to predicting polydisperse phase equilibria [Sollich, 2002]. Most methods involve reducing the dimensionality of the problem in order to get around the problem of the infinite number of particle species. For instance, an intuitively appealing approach is to discretise the distribution of particle sizes into a finite number of bins. However, there is no *a priori* way of knowing how the bins should be spaced or how wide they should be, and if a large number of bins is required the problem quickly becomes

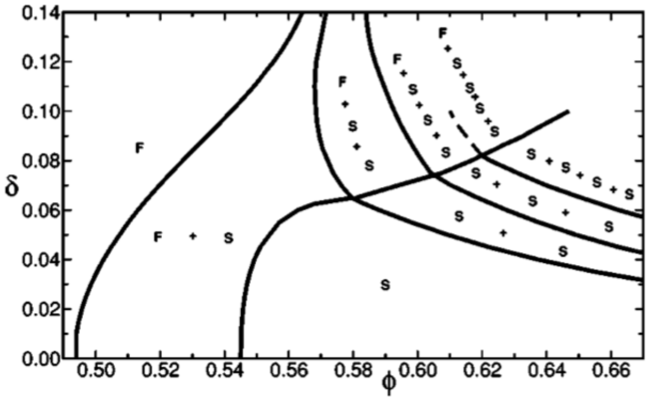

Figure 2.6: Hard sphere equilibrium phase diagram in volume fraction ϕ and polydispersity δ ($\equiv \sigma$) [Fasolo and Sollich, 2003], with fluid and solid (crystal) phases denoted by F and S. Destabilisation of the crystal narrows the 'F + S' coexistence region, and at high polydispersity, particles are distributed into multiple coexisting solid phases.

very complicated anyway. Any such approximation along these lines brings with it arbitrary uncontrolled errors [Sollich, 2002].

Perhaps the most successful and versatile approach has been to express the free energy of a polydisperse system as a function of a finite number of *moments* of the particle size distribution – the 'moment free energy' approach [Sollich and Cates, 1998; Warren, 1998]. Proceeding from this starting point, recent theoretical work by Sollich and Fasolo [Fasolo and Sollich, 2003] has produced a detailed equilibrium phase diagram for polydisperse hard spheres, reproduced in FIG. 2.6. Even at low to moderate polydispersity, there is a marked departure from the monodisperse case – note that 'near-monodisperse' colloidal suspensions typically have a polydispersity of at least $\sigma = 0.04$, usually more.

Polydispersity destabilises the crystal phase relative to the fluid, because particles of varying size cannot easily fit into a regular crystal lattice – the entropic benefit of increased free volume per particle in the crystal is therefore reduced as the particles are packed in a rather inefficient way given their distribution of sizes. This consideration, along with experimental observations, had led to much interest in the notion of a 'terminal polydispersity' of order $\sigma \approx 0.06 - 0.12$ beyond which a system cannot crystallise [Bartlett, 1997; Phan *et al.*, 1998] and the fluid phase is always stable. However, polydispersity is also associated with fractionation (demixing) of particle sizes between phases [Bartlett, 1998; Evans *et al.*, 1998; Wilding and Sollich, 2004], so that individual phases may reduce their polydispersity. The work of Sollich and Wilding has shown that polydispersity is only 'terminal' if one requires that all particles be incorporated into a *single* crystalline phase – the true equilibrium state at such high polydispersities is a coexistence of *multiple* crystal phases [Sollich and Wilding, 2011], as shown in FIG. 2.6, each with smaller polydispersity than the parent distribution.

Note that the dynamic pathway by which such a complex arrangement of phases might be achieved is another matter altogether. Indeed it appears that, in real systems, the required fractionation for the formation of multiple crystal phases may tend to be kinetically frustrated

[Liddle *et al.*, 2011]. Nonetheless, the work of Sollich and collaborators, which is generally accepted as the best current approach to dealing with polydispersity at equilibrium, highlights the complexity of polydisperse thermodynamics and the pitfalls that can be encountered in approximating or ignoring its effects.

R. M. L. Evans has developed a perturbative approach for small polydispersities [Evans, 2001], in which the thermodynamic properties of polydisperse systems are expanded about the monodisperse limit. The theory provides, among other things, a general way of describing the fractionation of polydisperse properties between coexisting phases, becoming exact in the limit of small polydispersity, and has strong experimental support [Erné *et al.*, 2005; Evans and Fairhurst, 2004; Evans *et al.*, 1998]. It is the method used in the present work when such predictions are required. For fluid-crystal coexistence in polydisperse hard spheres, the theory predicts that the higher density crystal phase should contain on average larger particles than the fluid, the difference in mean particle size between the two phases being proportional to the *variance* of the parent distribution. This prediction can be understood heuristically as being due to the greater positional entropy that is made available to the fluid particles by removing larger particles into the crystal.

It should be clear from this discussion and from the results in subsequent chapters that polydispersity plays a key role in the physics of most soft matter. This is particularly frustrating from the experimental point of view. Difficulties in accurately quantifying a system's polydispersity, and the already very complex variety of equilibrium and nonequilibrium behaviour on display, mean that polydispersity often arrives as a mysterious catch-all explanation for any remaining weirdness once the other results have been explained. Some experiments, e.g. [Erné *et al.*, 2005; Evans and Fairhurst, 2004; Evans *et al.*, 1998; Liddle *et al.*, 2011] have sought to rectify this, but as our theoretical understanding of polydispersity is almost entirely restricted to equilibrium, there is difficulty in its application to these real – and hence dynamical – systems. I hope to show that dynamical simulation of polydisperse phase behaviour can help to fill in this gap, elucidating how polydisperse phase equilibria may be approached in real systems and helping to make the understanding of polydispersity and its effects more practically applicable.

2.3 Dynamics

Many of the practical differences between colloids and simple fluids arise from their differing dynamics. In this section I outline a few dynamical features of colloidal suspensions by way of preparation for the calibration results in Chapter 4.

2.3.1 Brownian particles

In the absence of other particles (or in a suspension dilute enough that particles are well-separated) and of external fields (e.g. gravity), a hard sphere-like colloidal particle is subject to two forces from its solvent: viscous drag (proportional to the particle's velocity) and a fluctuating noise force, resulting from random collisions from the solvent molecules. On extremely short time scales, the motion of a colloidal particle, like that of a molecule, is ballistic. For times longer than a characteristic inertial timescale, however, the particle's velocity is viscously damped. This timescale is given by:

$$t_i = \frac{m}{6\pi\eta a} \tag{2.5}$$

where $a = d/2$, m is the particle's mass, and η is the viscosity of the solvent. Beyond this timescale, the particle's motion is *diffusive*. For an isolated particle the appropriate diffusion coefficient is the ideal Stokes-Einstein diffusion coefficient D_0:

$$D_0 = \frac{k_B T}{6\pi\eta a} \ . \tag{2.6}$$

The timescale associated with diffusive motion on the scale of the particle itself (i.e. for a particle to diffuse approximately its own diameter) is $t_d = a^2/D_0$ and, in a colloidal suspension, is orders of magnitude longer than t_i (~ 1 ns versus ~ 1 ms). So, on the timescale on which most interesting colloidal phenomena occur, the colloidal particles' inertia is unimportant and they move diffusively via short Brownian hops. The $t > t_d$ regime is called the long-time Brownian regime [Segrè *et al.*, 1995], while that between t_i and t_d is the short-time Brownian regime.

2.3.2 Hydrodynamic interactions

As well as interacting directly, colloidal particles may interact with one another via hydrodynamic interactions (HI), i.e. those mediated through the solvent. This is one respect in which colloidal suspensions are very different to molecular fluids. These forces are intrinsically many-body in nature and are notoriously complex, exact theoretical results being exceedingly rare, and restricted to the dilute limit (e.g. [Batchelor and Green, 1972]). On the simulation side, John Brady and collaborators have made significant progress using the method of *Stokesian dynamics* [Brady and Bossis, 1988], but the computational expense of this and other methods of simulating HI [Cates *et al.*, 2004] severely limits the size of simulation which can be performed. For this reason, the simulations presented in this thesis do not incorporate HI, in common with many comparable Molecular and Brownian Dynamics simulations in the literature.

Discussion of the potential impact of neglecting HI in the context of particular results is given as and when necessary, but the overarching theme is that the measurement of 'solvent-independent' phenomena (i.e. those without HI) constitutes at least an important baseline. This is particularly true given that the relative importance of HI depends on the details of the system in question; for example, HI can be screened to some extent by the electrostatic interaction in charged colloidal suspensions [Riese *et al.*, 2000].

2.3.3 Other forces, external fields

Beyond Brownian fluctuations and HI, many other forces can affect both the thermodynamics and dynamics of particular colloidal systems. Inter-particle forces include dispersion forces, such as the ubiquitous van der Waals interaction, electrostatic forces, and depletion interactions (whereby exclusion of some smaller dispersed species from the space between two colloid particles approaching one another leads to an overall osmotic pressure bringing the colloids closer together). External fields often play a role as well. Gravity, although by definition weak in comparison to Brownian fluctuations in colloidal systems, will typically lead to sedimentation on some timescale unless the colloids are density-matched to the solvent. Electric fields can be applied in order to induce electrophoresis of charged species, thereby separating a mixture of particles.

Clearly, the variety of possible interaction potentials and external fields is huge. For this reason, simulation and experimental work on comparatively simple model systems such as hard spheres or square well particles (as used in this thesis) is extremely valuable in determining general principles as a baseline from which the behaviour of more complex systems can be predicted and explained. For instance, experiments seeking to discover the behaviour of these simple model systems will typically aim for low polydispersity, density-matching between particles and solvent or even microgravity conditions [Bailey *et al.*, 2007], refractive index-matching to reduce van der Waals attractions, etc.

2.3.4 Self diffusion and collective diffusion

Self diffusion describes the translational motion of individual colloidal particles. Since diffusive motion is random, the average displacement $\langle r(t) \rangle$ of a particle at time t is always zero in an isotropic system. Instead, we are interested in the mean squared displacement (MSD), $\langle r(t)^2 \rangle$, which varies linearly with time for diffusive motion. In three dimensions, the dependence is given by:

$$\langle r(t)^2 \rangle = 6D_0 t \tag{2.7}$$

for a dilute suspension. When colloid-colloid interactions become important the self diffusion coefficient becomes a function of ϕ, $D_{self}(\phi)$. On further increasing ϕ, two regimes are resolved: the MSD increases diffusively at short times, subdiffusively on intermediate times due to 'caging' of particles by their neighbours, and diffusively again on long timescales, the unit of motion then being a stochastic 'hop' of a particle from one cage to the next. Therefore, distinct short- and long-time self diffusion coefficients D_{self}^{short} and D_{self}^{long} are defined and measured in order to describe this range of timescales [Segrè *et al.*, 1995]. Since the presence of other particles inhibits the motion of any given particle, both these self diffusion coefficients decrease with increasing ϕ. The latter decreases due to difficulty in moving through cages, whereas the decrease in D_{self}^{short} due to concentration is a hydrodynamic effect caused by the effect of nearby particles on the hydrodynamic environment of the particle [Medina-Noyola, 1988].

Collective diffusion is a somewhat different concept. It concerns the collective redistribution of particle matter as opposed to the displacement of individual particles. The collective or 'gradient' diffusion coefficient appears in Fick's law, which describes the diffusion of matter down a concentration gradient:

$$J = -D_{coll} \nabla \phi \tag{2.8}$$

where J is the concentration flux. D_{coll}, in contrast to the self diffusion coefficient, increases with ϕ. It can also be defined at short and long times.

Both self and collective diffusion coefficients may also depend on the wavenumber q, where $2\pi/q$ is the characteristic length scale on which the corresponding diffusive rearrangements take place [Segrè *et al.*, 1995]. Except when explicitly noted, I will be talking about the low-q macroscopic limit.

Further consideration of these dynamical quantities is included in Chapter 4, where I present calibration simulations used to confirm the presence of realistic structural and dynamical physics. In Chapter 8, diffusion plays a key role in a theory of dynamically-induced fractionation which I attempt to test. In Chapter 10, an approach to polydisperse diffusion that takes advantage of a moment-based approach (as is used by Sollich and others to simplify equilibrium calculations) is introduced.

2.3.5 Phase transition kinetics

One of the most important things to understand about a colloidal system, or indeed any sample of condensed matter, is when and how it will change phases. We have seen how phase transitions and coexistences may be predicted from thermodynamics. A full picture requires also a consideration of the dynamics of the system, since it is only by the collective re-organisation of the constituent particles that phase transitions actually take place. The most fundamental connection to make is that it is the Brownian motion of particles that allows the system to explore its microstates and renders thermodynamics relevant at all.

The equilibrium states predicted by thermodynamics may be achieved in practice. However, for a variety of reasons, they may not be, giving rise to the important possibility of nonequilibrium states. *Dynamical arrest* describes a situation in which particles become locked into amorphous states with diffusion slowing dramatically. For instance, colloids with sufficiently short-ranged and strong interactions form gels, cluster- or string-like networks of particles which may subsequently collapse under the force of gravity [Zaccarelli, 2007]. At very high densities, repulsive particles form glasses if they are for some reason unable to crystallise (perhaps due to high polydispersity or fast compression). Glassy systems are the subject of great interest due to their wide variety of applications and inherently nonequilibrium statistical mechanics. The glass transition is extremely general and seems not be accompanied by the growth of any structural correlation length, but by the growth of a dynamical one [Berthier, 2011].

Even if equilibrium is approached, the kinetics of phase ordering may be nontrivial. Metastable states, described in Section 2.1.4, can play an important role. For instance, metastable gas-liquid separation can enhance crystal nucleation [Anderson and Lekkerkerker, 2002; Fortini *et al.*, 2008] as the likelihood of nucleation is increased inside the high density liquid phase. The role of metastable states in crystal growth is the subject of Chapter 6.

Finally, there is the possibility of frustration (as opposed to arrest), i.e. even when a system is not arrested in the way that a gel or glass is, the required equilibrium arrangement may be so complex that is not achieved on a reasonable timescale. This is thought to be the case for some of the complex phase diagrams in polydisperse systems [Liddle *et al.*, 2011], in which the path to equilibrium may rely on relatively tortuous fractionation processes. In contrast to their equilibrium aspects, the dynamics and phase transitions kinetics of polydisperse systems remain largely unexplored from a theoretical point of view [Evans and Holmes, 2001; Warren, 1999]. This gap in knowledge is one of the key obstacles in the path towards a more practical understanding of the effects of polydispersity in real systems.

2.4 Chapter summary

In this chapter I have briefly introduced the thermodynamics of phase behaviour, before explaining how it is that this familiar apparatus is applicable to colloidal systems. I explained the complicating effects of polydispersity on the thermodynamics and equilibrium phase diagrams, with reference to the hard sphere model. I then outlined the key dynamical features of colloidal systems and introduced the idea of phase transition kinetics, an important consideration in applying equilibrium theory to real systems. I hope this overview has served to illustrate the generality of the physics displayed by colloidal systems, the complexity of their specific nuances, and the considerable progress remaining to be made with respect to usefully predicting and explaining the behaviour of polydisperse systems, particularly from a dynamical point of view.

Chapter 3

Simulation

In this chapter I introduce and motivate the use of simulation in the study of condensed matter. I then outline the simulation approach used in my work and the details of particular techniques I have employed.

3.1 Simulation in soft matter

Condensed matter physics is typically concerned with the collective behaviour of large numbers of particles. Statistical mechanics attempts to determine this collective behaviour by starting from a knowledge of the microscopic dynamics of the particles in question – equations of motion, interaction potentials and so on – and proceeding analytically, introducing assumptions or approximations where necessary to make the calculations tractable. The bulk properties thus derived can be related to one another using thermodynamics, so statistical mechanics and thermodynamics complement each other in studying the microscopic and macroscopic properties of condensed matter. Solving by hand the equations of motion of a thermodynamically large number of particles is not a practical possibility, so the strength of analytical results depends on the reliability of the approximations and expansions used.

Computer simulation attacks the problem from another angle. It provides a way of numerically solving the equations of motion in a large system, for a model specified by the simulator. Therefore, insofar as the simulation is bug-free and does what the simulator claims it does, it accurately reproduces the collective behaviour of the model being simulated. The usefulness of the simulation then depends on e.g. how well the specified particle properties mimic those of the physical system the simulation is intended to represent or, if the model is intentionally rather abstract and generic, the value of the physical insights gained from studying its behaviour.

3.1.1 Equilibrium

In large class of simulations – those intended to study equilibrium behaviour – one is in fact rather free to choose the details of the microscopic dynamics of the simulation, within some constraints. For a classical system described by the canonical ensemble at equilibrium the probability distribution of microstates is described by the Boltzmann distribution. A microstate is a particular configuration of particle positions and velocities which we can label X. If a microstate is associated with an energy $E(X)$, then the probability of the system being in that microstate is proportional to the Boltzmann factor $e^{-E(X)/k_BT}$, where k_B is Boltzmann's constant. Each *macrostate* (a state described by macroscopic observables, such as a crystal or fluid phase) is associated with a certain subset of all the possible microstates. A system at equilibrium explores its microstates with the probabilities given by the Boltzmann distribution, thereby attaining the macrostate whose microstates dominate the distribution.

To achieve equilibrium in a simulation, therefore, we need to correctly sample the Boltzmann distribution. This places a constraint on the rates of transitions between microstates X and Y which is called *detailed balance*:

$$\frac{\omega_{XY}}{\omega_{YX}} = \frac{e^{-E(Y)/k_BT}}{e^{-E(X)/k_BT}} \tag{3.1}$$

where ω_{XY} is the rate of transitions from X to Y and ω_{YX} is that of the reverse transition. The rate ω_{XY} may be defined probabilistically as the probability per unit time to move into state X given that we start in state Y.

3.1.2 Monte Carlo simulation

Monte Carlo (MC) describes a simulation strategy in which computer-generated pseudo-random numbers are used to sample the configuration space of the system. That is, transitions between states are proposed stochastically and accepted in such a way as to satisfy detailed balance. A popular way of doing this is to use the Metropolis algorithm. In this algorithm trial transitions, having been proposed, are accepted according to:

$$P_{accept} = \min(1, e^{-\Delta E/k_BT}) \tag{3.2}$$

where P_{accept} is the acceptance probability and $\Delta E = E_{new} - E_{old}$ is the energy change involved in the transition – detailed balance is therefore respected. Implementing this algorithm in a simulation of hard spheres is particularly simple: any moves which cause particles to overlap are rejected (as $\Delta E = \infty$) and any others are accepted with $P_{accept} = 1$.

The efficiency of an equilibrium Monte Carlo simulation, i.e. how quickly it reaches equilibrium, depends primarily on how clever the simulator is in choosing what trial moves to use. Consider a simulation of polydisperse hard spheres in which one is trying to equilibrate crystal-fluid coexistence. As discussed in Section 2.2, such a coexistence involves some degree of fractionation of particle sizes between the phases and even, if the polydispersity is high enough, *multiple* crystalline phases coexisting with one another with individually small polydispersities. As I mentioned there, such complex fractionation and coexistence may be rather difficult to achieve in a real system, so that experiments do not actually reach equilibrium [Liddle *et al.*, 2011; Sollich and Wilding, 2011]. This is particularly the case when crystal phases are involved, because particles in a crystal are essentially locked in to their positions – the long distance diffusion required for particle exchange between phases is mediated mainly by the presence of defects and is extremely slow in comparison to that in a fluid. Therefore, if the trial moves chosen in the Monte Carlo simulation are only those which are dynamically realistic, i.e. mimic the tiny

Brownian hops of colloidal particles, equilibration with respect to phase composition will be unfeasibly slow. It is worth emphasising that simulations of even moderately sized systems like this are far slower than real time – if an experimental system does not reach equilibrium on a reasonable timescale, a simulation which mimics the dynamics of that system certainly will not!

Equilibrating such a system requires more cunning. In the case of polydisperse hard spheres, the simulations performed by Nigel Wilding to verify Sollich's equilibrium calculations are an excellent example [Sollich and Wilding, 2010]. Recognising that, in an equilibrium simulation, we are not interested in the microscopic dynamics of the particles (as long as they satisfy detailed balance), Wilding's simulations use a sophisticated set up in which particles are allowed to spontaneously appear and disappear, change size and so on. At the cost of being, dynamically speaking, emphatically unphysical, the simulation explores its configuration space much more efficiently than a real system. Because detailed balance is respected overall (notwithstanding the further technique of *biased sampling*, in which some bias is introduced to the sampling to increase efficiency and then 'rolled out' of the results to recover detailed balance afterwards), it gives access to the equilibrium phase diagram of the system and can be used to verify equilibrium calculations. The value of this combined approach of equilibrium theory and equilibrium simulation should be clear: for any system that can be defined and simulated efficiently, the equilibrium physics can be determined, compared with predictions, and used to determine whether the observed behaviour of corresponding experimental systems is equilibrium or nonequilibrium in origin.

3.1.3 Kinetic Monte Carlo

Equilibrium simulation, like equilibrium theory, does not tell the whole story – it cannot, and is not designed to, tell us about kinetics: how real systems actually go about reaching their final steady states, which may not be equilibrium anyway. Nonequilibrium or *dynamical* simulations are concerned with addressing this question. An example is Molecular Dynamics (MD), in which the force fields describing the system are specified and the equations of motions integrated to obtain dynamical trajectories. This method is often used in simulating biological processes, which are typically nonequilibrium in nature [Scholz *et al.*, 1987], and is very computationally expensive for large systems (e.g. a single protein molecule in water) so that it is often only feasible to simulate timescales corresponding to nanoseconds or microseconds of real time.

Kinetic Monte Carlo is a somewhat cheaper method which is popular in soft matter simulation. Kinetic (or 'dynamic') MC is distinct from equilibrium MC methods in that it is restricted to dynamically realistic trial moves, i.e. small stochastic 'hops' representing the movement of particles performing Brownian motion. Unlike equilibrium MC, it seeks to replicate a real system's dynamical progress through configuration space. For small step sizes, dynamical results from Kinetic MC simulation compare well with those of Brownian Dynamics (a form of MD for systems where inertia is negligible, such as colloids) [Sanz and Marenduzzo, 2010]. The simplicity of Kinetic MC and its natural suitability for representing inertia-free stochastic motion have resulted in a number of applications in simulating hard spheres and related particles [Auer and Frenkel, 2001, 2005; Sanz *et al.*, 2007]. The first such application I am aware of is Ref. [Cichocki and Hinsen, 1990].

3.2 Simulation method

The simulation method used in my work is based on a Kinetic MC algorithm. In this section, I describe the details of my implementation and how I have applied it to simulate the phase behaviour of colloids.

3.2.1 Particle dynamics

As I have mentioned, the defining feature of Monte Carlo simulation is the use of pseudo-random numbers to propose trial moves which are then accepted probabilistically. The aim of my work is to simulate the Brownian motion of particles, where inertia is negligible due to viscous damping by the solvent (see Section 2.3), so the trial moves used must be physically realistic from this point of view. The Brownian steps are therefore random in direction, with magnitudes taken from a Bates (pseudo-Gaussian) distribution centred on zero[1].

In each iteration of the dynamical loop, a particle is chosen at random, and a trial move proposed. As we saw in Equation 2.6, the Stokes-Einstein diffusion coefficient for a particle of radius a depends inversely on a. This is incorporated into the simulation by scaling the Bates distribution so that the mean step size for particle i is $\Delta\sqrt{\langle d\rangle/d_i}$, with particle diameter $d = 2a$ and $\langle d\rangle \equiv 1$ the mean diameter, so as to give the correct dependence of the short-time diffusion coefficient upon particle size. In general, I set the step parameter to $\Delta = 0.02$, then check the results presented to ensure that they are not affected by reducing Δ. Comparison with other systems is facilitated by choosing the natural time unit t_d, the time taken for a free particle of diameter the mean $\langle d\rangle$ to diffuse a distance equal to its own diameter. With this choice of Δ, each t_d corresponds to 2500 Monte Carlo Steps. A Monte Carlo Step is a single ‘sweep’ of the system in which each particle is, on average, visited once (remember that the selection of which particle undergoes a trial move is stochastic).

The Metropolis acceptance criterion, Equation 3.2, is popular for equilibrium MC simulations. It assigns an equal probability (unity) to all moves which do not increase the system’s energy, even those which *decrease* it. For a dynamical simulation in which particles may have an attractive component to the potential, it is more natural to represent their dynamics with the alternative acceptance probability [Martin, 1977; Novotny, 2001]:

$$P_{accept} = \frac{\exp(-E_{new})}{\exp(-E_{new}) + \exp(-E_{old})} \tag{3.3}$$

Moves that conserve energy are now accepted with probability 0.5, so that moves that bring a particle into the reach of attractive potentials can be assigned a higher probability, representing more rapid movement down a potential gradient.

The magnitude of the Brownian hops essentially defines the short-time self diffusion coefficient D_{self}^{short} which, at finite concentration, is different to the long-time coefficient D_{self}^{long}. In simulations which do not include hydrodynamics, D_{self}^{short} does not in itself depend on ϕ, as this dependence

---

[1] The Bates distribution is a sum of n random numbers uniformly distributed on the unit interval – in my implementation $n = 4$. The Bates distribution is computationally inexpensive to sample from, hence its use in the central dynamics loop of the simulation.

is a purely hydrodynamic effect [Medina-Noyola, 1988]. D_{self}^{long} depends on ϕ via the increased structural frustration (caging) by neighbouring particles, as is shown in Section 4.4.

3.2.2 Initialisation

An important question in a simulation of hard spheres or related particles is how to place the particles in the simulation box in the first place. In all the work presented in this thesis, the appropriate initialisation is a random packing of particles at the required volume fraction ϕ. For volume fractions approaching the maximum that can be achieved in a random packing ($\phi \approx 0.64$), rather sophisticated algorithms such as that in Ref. [Jodrey and Tory, 1985] are required in order to generate the packing while avoiding crystallisation. Most are descended from original Lubachevsky-Stillinger algorithm, in which particles perform Brownian or Newtonian motion while growing from point particles to particles of the size required to give the required ϕ. I use a generalisation of this for the polydisperse case in which each particle has a 'seed' radius drawn from a polydisperse distribution, with the actual radius given by a scaling factor which increases from 0 initially to 1, when the particles reach the required ϕ.

3.2.3 Finite-size effects, periodic boundaries

Simulations must contain a finite number of particles, but one is typically interested in the behaviour of a thermodynamically large number of particles, as would exist in a macroscopic sample of whatever substance it is that is being simulated. In this case, 'taking the thermodynamic limit' is important, i.e. checking whatever results are found for dependence on the simulation size and, if necessary, extrapolating to infinite system size if system size dependence cannot be eliminated. Finite-size effects must also be considered at the other end of the size spectrum: the step size parameter Δ, in order to accurately represent Brownian motion, must be small enough that particles perform a large number of random steps in between encountering one another (colliding, or otherwise interacting). These considerations must be offset against the fact that an infinitely large simulation in which particles take infinitesimally short steps is not feasible. Chapter 4 includes some calibration results illustrating possible finite-size effects, forming the basis for decisions about what system and step sizes to use in general for efficient but physical simulations. Additionally, the particular results presented in later chapters are checked for finite-size effects on each result, in order to verify that any discoveries or findings are not simply artefacts.

Another important consideration is the edge effects encountered when particles interact with the walls of the simulation box. Although physical and interesting in their own right, such effects are typically not desirable if one is trying to study phenomena that would occur in the bulk of a real system. The standard way to get around this problem is to use periodic boundary conditions, under which particles which disappear from one side of the simulation box simply reappear on the other, without encountering a wall (as in the *Snake* mobile phone game). There are therefore no edges in the simulation. Even with this solution, finite-size effects may also come into play if any correlation length in the system is longer than half the side length of the periodic simulation box – see Section 4.5 for discussion of this in terms of a dynamical correlation length in glassy systems. Unless mentioned, all simulations presented in this thesis use standard periodic boundary conditions.

3.2.4 Crystal templating

The work on crystallisation in this thesis makes use of a crystal templating method suggested by R. M. L. Evans, which allows the growth process of a single crystal to be simulated on a physically relevant timescale. Crystal nucleation is a rare event and therefore difficult to simulate in that one may waste a large amount of simulation time just waiting for nucleation to happen – specialised techniques such as forward flux sampling are required to simulate such rare events [Filion *et al.*, 2010]. In order to be able to focus on crystal growth without waiting for nucleation, I induce crystal growth with the use of a template, positioned before the simulation starts.

The template is positioned so as to cover one end of the simulation cell (in the $y-z$ plane), so that a crystal is induced to grow along the x axis. It consists of a layer of immobilised particles arranged in a face of the FCC lattice, which is the equilibrium crystal structure for hard spheres and the closely related systems studied in this thesis. I have used the (100) and sometimes (110) faces, in order to avoid the complicating effects of *stacking faults* – these may occur in the (111) direction, where planes can switch randomly between ABC... and ABAB... registrations, resulting in a mixture of HCP and FCC structures [Amini and Laird, 2006].

The y and z dimensions are periodic, with lengths L_y and L_z such that the template fits into the periodic boundary conditions. The x axis length L_x is much larger than the other two so that the crystal can grow many layers over a long period of time, reproducing the physical situation in which nuclei are well separated from one another.

To create the template, each lattice point in the template is filled in turn by the leftmost (lowest-x) particle from the bulk. Then, in order to preserve the local volume fraction, the template is moved forwards by an amount corresponding to the mean x coordinate of the particles which were used to create it. The template is then moved in the y and z directions so as to minimise overlaps between template particles and any others nearby. For high bulk concentrations, some small overlaps usually remain, of order 5% of a particle volume per template particle at the highest bulk concentrations studied. The simulation then begins. While any overlaps with the template exist, the particles overlapping with the template are only allowed to perform moves which reduce their overlap. Using this algorithm, any overlaps disappear after only a few t_d. Notwithstanding complicating factors such as polydispersity, the template induces the particles close to it to match its crystalline order, causing a crystal to grow through the system. An illustrative snapshot shortly afterwards is shown in FIG. 3.1.

If the simulation remained periodic in the x direction also, a crystal would grow from both sides of the template, which is not desirable as I am aiming to study the growth of a single crystal in one direction. In order to avoid this, the bulk (amorphous) particles whose x coordinate is greater than $L_x - 1.5d_{max}$ (where d_{max} is the diameter of the largest particle) are immobilised, along with any particles that subsequently stray into this region. This amorphous wall prevents any particles reaching the back side of the template and does not induce nucleation in the way that a flat wall would.

The flat crystal template can be considered as representing the limit of a large nucleus – on a local scale, such a nucleus appears flat. Alternatively, one can choose to interpret the template

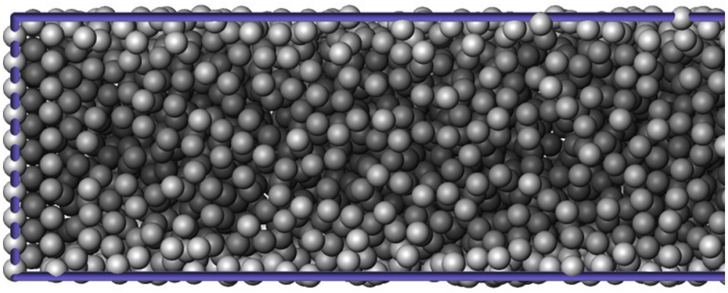

Figure 3.1: A simulation snapshot showing the template positioned at $x = 0$. At this stage, 1 or 2 layers have subsequently grown.

(more literally) as representing an actual template as might be realised in an experimental system. From that point of view, the simulation becomes one of *colloidal epitaxy* [van Blaaderen *et al.*, 1997], a process in which the properties of a template are chosen in such a way as to control the properties of the resulting crystal. Simulations in this direction are presented in Chapter 9.

In all chapters apart from Chapter 9, the particles used to create the template retain their size upon being repositioned so reflect the polydispersity (if any) of the bulk particle size distribution. This is to ensure that any effects of a particular size distribution in the template are averaged out over the multiple independent simulations used. In Chapter 9, the focus is on precise control of the template (now considered as an epitaxial substrate), so the particles are made monodisperse on being moved into the template, in an attempt to actively induce selectivity.

When crystallisation is not being studied, the simulation box is cubic instead of cuboidal, with periodic boundaries in all dimensions.

3.2.5 Flagging crystalline particles

In my simulations, as in most comparable ones in the literature, crystalline particles are identified by constructing bond order vectors $\mathbf{q}_6$ from the spherical harmonics Y_{6m} [Williams *et al.*, 2008; ten Wolde *et al.*, 1996]. A local order parameter is defined:

$$\mathbf{q}_{6m}(i) \equiv \frac{1}{N_b(i)} \sum_{j=1}^{N_b(i)} Y_{6m}(\widehat{\mathbf{r}}_{ij}) \ , \tag{3.4}$$

where the sum is over the N_b ‘near neighbours’ of particle i and $\widehat{\mathbf{r}}_{ij}$ is a unit vector pointing from the centre of particle i towards the centre of particle j. Following [Ackland and Jones, 2006], I define a near neighbour as a particle that is within $\sqrt{1.45 r_0^2}$ of particle i, where r_0 is the average separation between particle i and its closest 6 neighbours. For each particle can then be constructed a normalised 13-component complex vector $\mathbf{q}_6(i)$, the components of which are determined by the 13 spherical harmonics Y_{6m} summed over the angles of the imagined ‘bonds’ connecting it to its near neighbours.

Clearly, if two particles share a similar environment (as in a periodic solid), their vectors $\mathbf{q}_6$ will have similar components (importantly, this similarity holds irrespective of the coordinate system chosen), and the scalar product $\mathbf{q}_6(i) \cdot \mathbf{q}_6(j)$ will be nonzero and positive. Since $\{\mathbf{q}_6\}$ are normalised, the maximum possible value of such a scalar product is 1. I follow [Williams *et al.*, 2008] in classifying a particle as crystalline if $\sum_{j=1}^{N_b(i)} \mathbf{q}_6(i) \cdot \mathbf{q}_6(j) > 8.5$.

3.2.6 Visualisation

I have written my own visualisation code for the purposes of bug checking while developing the simulation. As the number and size of simulations have increased, so that they are usually run on remote computer clusters, I have moved to a third-party visualiser which is used to view the simulation trajectories after they have been run. In this respect I have been very fortunate to discover OVITO, which was used to produce the snapshots presented in this thesis [Stukowski, 2010].

3.3 Chapter summary

In this chapter I have introduced the idea of computer simulation as a method for studying condensed matter, with particular reference to the Monte Carlo (MC) method in which random numbers are used to sample configurations. I explained the need for sophisticated trial moves in equilibrium MC simulations in order to ensure configuration space is explored efficiently. I then moved from equilibrium to Kinetic MC, in which trial moves are limited to dynamically realistic ones, so that dynamical processes can be studied. Finally, I introduced the essential components of the simulation apparatus used in the present work.

Chapter 4

Results I – Calibration, structure and dynamics

4.1 Introduction

This chapter concerns simulations of the hard sphere system performed to confirm realistic structural and dynamical behaviour and to check for finite-size effects thereon. The simplicity and popularity of the hard sphere system means there is a large literature of theoretical, simulation and experimental results available. A viable simulation protocol must therefore successfully reproduce known structural features and, since we are concerned here additionally with the kinetics of phase transitions, exhibit sensible dynamical behaviour.

I will go through in turn various measurable quantities, with reference to the literature, introducing and discussing the corresponding simulation results produced by my code. All results in this section are taken from isotropic simulations which are cubic, i.e. $L_x = L_y = L_z \equiv L$. When polydispersity is introduced, the Bates (pseudo-Gaussian) distribution described in Chapter 3 is used.

Except where explicitly noted (e.g. Sections 4.3.2 and 4.5), I have been able to run the simulations for long enough that the averaging window for whatever quantity is being measured, and the time elapsed before measurement begins, are much longer than any structural correlation time. It should be pointed out that this does not necessarily ensure equilibration, if the system is in a metastable state such that a rare event (like a crystal nucleation) must take place for it to reach equilibrium. It does ensure that any spurious effects due to initial conditions are removed, and that statistical noise is averaged out. In the case of e.g. a stable fluid (which does not require rare events to reach equilibrium), this protocol should mean that measured quantities approach their equilibrium values.

I will discuss, as they occur, cases where long or divergent correlation times make the data susceptible to dependence on initial conditions, and the resultant deterioration in data quality. In such cases, I have used multiple starting configurations in order to ensure at least that one does not rely *entirely* on time averaging to remove dependence on initial conditions. In general, however, equilibration of phase coexistence, or proper mitigation of divergent relaxation times, requires specialised simulation techniques, outside the scope of the present work.

4.2 Static structure factor

The *static structure factor* (often just 'structure factor') is a quantity describing the presence of spatial correlations in a system. It is defined:

$$S(\mathbf{q}) \equiv \frac{1}{N}\langle\sum_{i,j} e^{-i\mathbf{q}\cdot(\mathbf{r}_i-\mathbf{r}_j)}\rangle \ , \tag{4.1}$$

in which $\mathbf{r}_n$ are particle positions, and the average is over the ensemble of microstates – if a system is ergodic, i.e. is fully exploring its configuration space, this is approached by a time average over long enough times. The structure factor depends on the wavevector $\mathbf{q}$, but in an isotropic system we are usually interested instead in $S(q)$, found by averaging over all wavevectors of magnitude q. $S(q)$ contains information about the magnitude of spatial correlations on a length scale $\sim 2\pi/q$. Experimentally, the structure factor is measured by performing scattering experiments [Segrè and Pusey, 1996], in which photons, neutrons or X-rays are directed at a macroscopic sample of the substance under investigation.

4.2.1 Monodisperse hard spheres

It is possible to extract $S(q)$ from simulation by explicitly performing the sum in Equation 4.1 and averaging through time. The finite side length of the simulation box means that the q-axis is discrete: $q_x = 2\pi/L, 4\pi/L, ...$. In order to avoid any artefacts involved in binning on the q-axis, I plot a data point for every unique q magnitude, averaging only over vectors of identical magnitude, e.g. $\mathbf{q} = (4\pi/L, 4\pi/L, 0)$ and $\mathbf{q} = (0, 4\pi/L, 4\pi/L)$. Additionally, I neglect at least the lowest q value in each dimension as it represents correlations on the length scale of the simulation box, so would be subject to artefacts from the periodic boundary conditions.

The structure factor of a hard sphere fluid is fitted reasonably well by the Percus-Yevick approximation [Hansen and McDonald, 2006]. In order to check that the simulations reproduce the structural behaviour of the hard sphere fluid, FIG. 4.1 shows measured structure factors at volume fractions $\phi = 0.3$ and $\phi = 0.5$ (just above the crystallisation volume fraction $\phi = 0.494$), along with a Percus-Yevick calculation for $\phi = 0.5$, for a monodisperse system of $N = 256$ hard spheres.

Thanks to the lack of binning in the q-axis, the data are very detailed. The low noise indicates good sampling, i.e. the simulations are behaving ergodically and have run for long enough that the time average approaches an ensemble average. The first peak in $S(q)$ in a homogeneous fluid corresponds to structuring due to the fact that particles cannot overlap – an inter-particle (centre to centre) spacing of approximately d is therefore extremely common, leading to a peak corresponding to this wavelength. As ϕ is increased, this peak shifts to higher q, i.e. shorter wavelength, and becomes stronger, indicating that particles are moving closer together on average and becoming more structured as they are more tightly packed. A peak at lower q than this would indicate spatial inhomogeneity on some larger length scale, as seen during spinodal decomposition in Chapter 5.

Although $\phi = 0.5$ lies just within the crystal-fluid coexistence region, no crystallisation was observed on the simulated timescale – the structure factor at this volume fraction is therefore that of a metastable fluid. It is known that the Percus-Yevick calculation slightly overestimates the

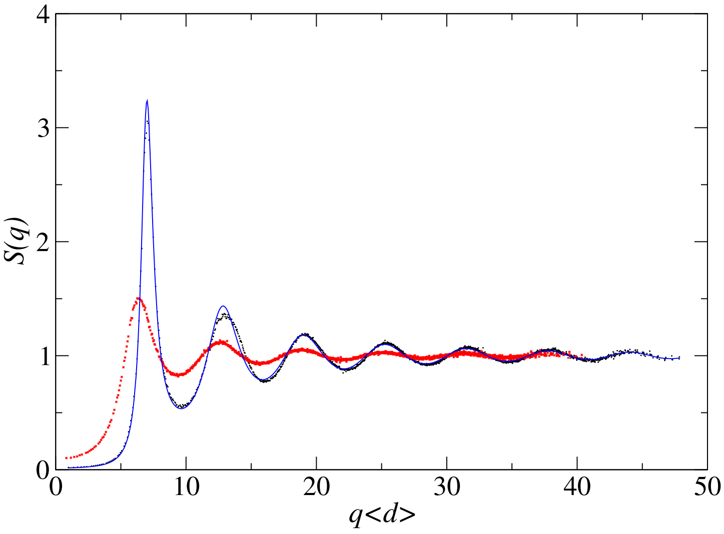

Figure 4.1: Measured structure factors from simulations of monodisperse hard spheres at $\phi = 0.3$ (red), $\phi = 0.5$ (black). The blue line is a Percus-Yevick calculation for $\phi = 0.5$. The q-axis is nondimensionalised by multiplication by $\langle d \rangle \equiv d \equiv 1$.

peak height for concentrated monodisperse systems [Scheffold and Mason, 2009], as is indeed observed in FIG. 4.1. Overall, though, the simulation data correspond well to the Percus-Yevick approximation and display the expected dependence on ϕ, indicating that equilibrium and metastable hard sphere fluid structural behaviour is being correctly reproduced.

4.2.2 Polydisperse total structure factor

Introducing polydispersity into the hard sphere model of course creates the possibility of much more complex structuring and, correspondingly, more measurable quantities. To begin with, I show the effects of polydispersity on the total structure factor, i.e. the one defined by Equation 4.1. FIG. 4.2 shows the monodisperse data from FIG. 4.1 along with that from corresponding simulations of polydispersity $\sigma = 0.15$, with all other parameters kept constant.

Zooming in on the first few oscillations, it is possible to resolve differences. Compared to the monodisperse case, the polydisperse structure factors show a first peak at slightly lower q, indicating the dominance of larger particles over the structuring signal (discussed in Section 4.2.3) which, due to their size, are less close together on average. Further, the height of the peak is lower. This indicates not so much *weaker* structuring as the fact that structuring is distributed over many slightly different length scales due to the particles' size-polydispersity, resulting in a greater 'smearing-out' of the signal over a range of q values. Both these features are known effects of size-polydispersity [Ginoza and Yasutomi, 1999].

4.2.3 Polydisperse partial structure factors

In order to gain more insight into the structural behaviour of a polydisperse fluid, it is necessary to go beyond its effects on the total structure factor and look also at *partial* structure factors.

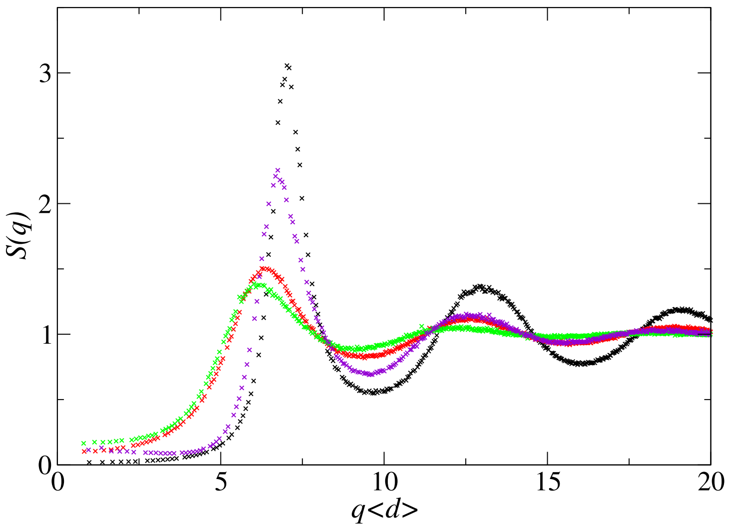

Figure 4.2: Measured structure factors from simulations of $\sigma = 0.15$ polydisperse hard spheres at $\phi = 0.3$ (green), $\phi = 0.5$ (purple). The monodisperse data from FIG. 4.1 are also shown for comparison.

The partial structure factor $S_{AB}(q)$ is defined:

$$S_{AB}(q) \equiv \frac{1}{N}\langle\sum_{i,j} e^{-i\mathbf{q}\cdot(\mathbf{r}_i-\mathbf{r}_j)}\rangle, i \in A, j \in B \tag{4.2}$$

where the directional average to give dependence on scalar q is implied as before. A and B denote size species – in a polydisperse system, particles are binned according to size in order that we end up with a finite number of partial structure factors. These structure factors measure spatial correlations between particles of particular sizes. Although the particle labels i and j can no longer be exchanged in the bracket $(\mathbf{r}_i - \mathbf{r}_j)$, the difference between the vectors is as likely to occur as its negative since the system is isotropic. This guarantees that the partial structure factors still take real values.

Having shown the overall effect on the structure factor of a relatively strong polydispersity in Section 4.2.2, I now move to a lower polydispersity of $\sigma = 0.06$ and a larger $N = 768$ in order to ensure good sampling of the particle size bins. FIG. 4.3 shows data from simulations with $\phi = 0.3, 0.4$. The partial structure factors result from binning moderately-sized particles into two bins either side of $\langle d\rangle$. As expected according to the definition [Weysser *et al.*, 2010], the off-diagonal element oscillates around zero while the diagonal ones oscillate around some finite value. The curves are broadly similar in shape to one another, as the particles considered are of roughly similar size. It is clear however that the larger particles in $S_{11}(q)$ have greater typical separation, as evidenced by the lower q position of the first peak compared to $S_{00}(q)$. The total structure factor is obtained simply by summing the four partial elements. The deconstruction reveals that the reduced peak height due to polydispersity, observed in FIG. 4.2, is indeed the result of out-of-phase contributions from different particle sizes, in agreement with Ref. [Weysser *et al.*, 2010] (there, such partial structure factors serve a key function as the input parameters

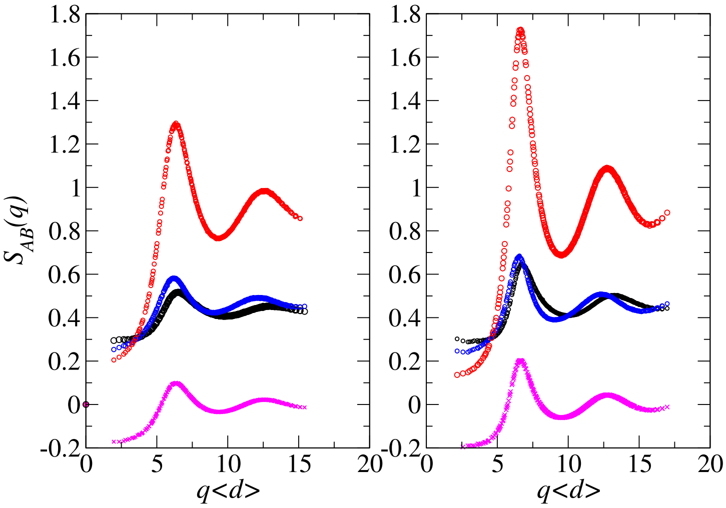

Figure 4.3: Measured structure factors from a simulation of $\sigma = 0.06$ polydisperse hard spheres at $\phi = 0.3$ (left) and $\phi = 0.4$ (right), with $N = 768$. The particles are binned into $0.9 < d < 1$ ('Bin 0') and $1 < d < 1.1$ ('Bin 1'). The total structure factor is shown in red. $S_{00}(q)$ is shown in black, $S_{11}(q)$ in blue and the off-diagonal element $S_{01}(q) = S_{10}(q)$ in pink.

for the Mode Coupling Theory of the glass transition.)

4.3 Intermediate scattering function

In addition to the static structure factor, a dynamical structure factor can be defined to give information about dynamical structural correlations in the system, i.e. the timescale on and manner in which structural correlations relax due to thermal motion. Experimentally, this is accessed by resolving the energy as well as the direction of the photons, neutrons etc. scattered by a substance during a scattering experiment: just as wavenumber and distance are related by a Fourier transform, so are time and energy. The most useful quantity is the intermediate scattering function $F(q,t)$, defined as:

$$F(q,t) \equiv \frac{1}{N}\langle \sum_{i,j} e^{-i\mathbf{q}\cdot(\mathbf{r}_i(t)-\mathbf{r}_j(0))} \rangle \ , \tag{4.3}$$

where again the dependence on scalar q is achieved by averaging over wavevectors. In taking and analysing these data I referred extensively to the work of Peter Pusey and collaborators – see for instance Refs. [Segrè *et al.*, 1995; Segrè and Pusey, 1996]. Ref. [Segrè *et al.*, 1995] in particular contains a clear explanation of the underlying theory and collects together extensive experimental and simulation data.

In Table 4.1, the various diffusion coefficients extracted from the intermediate scattering functions in the following discussion are collated, along with appropriate reference values. I now go through each regime in turn, making detailed comparison with reference values and explaining any discrepancies.

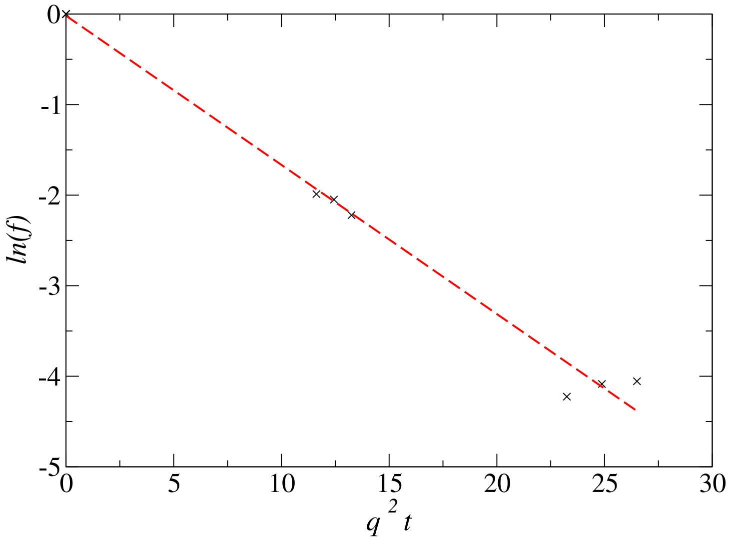

Figure 4.4: Initial decay of $f(q,t)$ in a non-interacting system on a log-linear plot – the negative of the slope therefore defines a diffusion coefficient D. Data points are combined from $q = 6.82, 7.05, 7.28$. The red line is a fit of negative gradient 0.165 ± 0.005, in agreement with the Stokes-Einstein diffusion coefficient $D_0 \equiv 1/6$.

4.3.1 Decay of the first peak

It is convenient to focus on the decay through time of $F(q,t)$ for a particular value or range of q. A feature often chosen is the first peak of the structure factor – the decay of this peak gives us information about the rearrangement of particle density on the length scale of the inter-particle spacing, and has the added advantage of being the region in which the signal is strongest. The data in this subsection concern systems of $N = 768$ with polydispersity $\sigma = 0.06$ (chosen to approximate the polydispersity of an experimental 'near-monodisperse' system).

One usually works in terms of the normalised intermediate scattering function $f(q,t) \equiv F(q,t)/F(q,0)$. In the ideal case of non-interacting particles, it should decay diffusively as:

$$f(q,t) = e^{-Dq^2 t} \tag{4.4}$$

where D is the collective diffusion coefficient which, if the particles are non-interacting, is just the Stokes-Einstein diffusion coefficient D_0.

Non-interacting system

In order to confirm the correctness of my implementation, FIG. 4.4 shows the decay of $f(q,t)$ in a non-interacting system of particles. Three q values in the vicinity of where the first peak would be in an interacting system are combined – note however that to within finite-size effects the decay should be independent of q anyway as the particles are non-interacting. It can be seen that $f(q,t)$ is decaying exponentially, defining a diffusion coefficient $D_{\textit{non-interact}}$ according to Equation 4.4. In the absence of interactions, $D_{\textit{non-interact}}$ is equal within error to $D_0 \equiv 1/6$, as

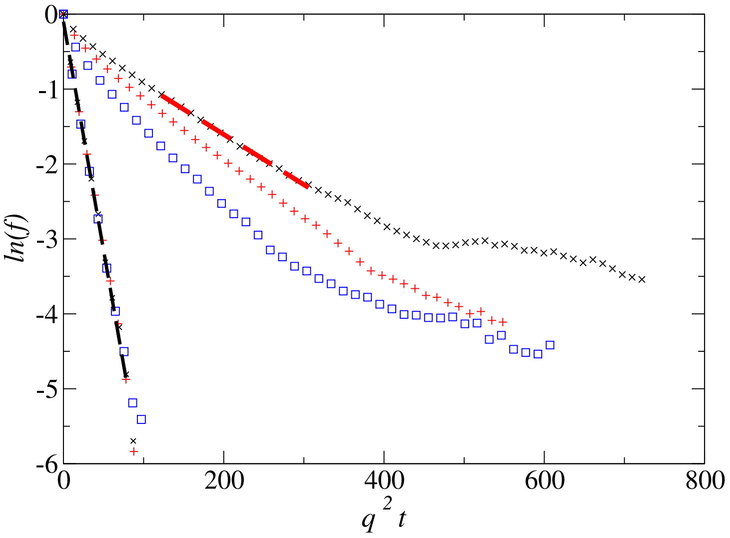

Figure 4.5: Decay of $f(q,t)$ at finite concentration. The faster decaying trio of series are for $\phi = 0.3$, the slower decaying ones for $\phi = 0.5$. Black crosses, red plusses and blue squares indicate respectively increasing q values around the first peak in $S(q)$: $q = 5.90, 6.24, 6.57$ for $\phi = 0.3$ and $q = 7.00, 7.40, 7.79$ for $\phi = 0.5$. The black dashed line is a fit to the $\phi = 0.3$ data giving $D = 0.0610 \pm 0.0007$. The red dashed line is a fit to the $\phi = 0.5, q = 7.00$ data over the range spanned by the line, giving $D = 0.00670 \pm 0.00008$.

expected. As well as illustrating the behaviour of $f(q,t)$ in the ideal case, this serves to give confidence in the substantial analysis code written to determine $f(q,t)$.

Finite concentration

I now move on to the effects of finite concentration at finite-q. It is known that in this regime, the collective diffusion coefficient controlling the decay of $f(q,t)$ becomes a function of q and ϕ, splitting into distinct short- and long-time coefficients at high concentrations.

Data for systems at $\phi = 0.3, 0.5$ are shown in FIG. 4.5 (with all other parameters as for the non-interacting system). At $\phi = 0.3$ the data from different q values around the first peak still overlay reasonably well, and define a diffusion coefficient substantially less than D_0, illustrating that particle rearrangement on the scale of the inter-particle separation is slowed by the hard sphere interaction. The ratio $D_0/D_{firstpeak}(\phi = 0.3) \approx 2.73$ compares well with the value of ~ 2.7 from corresponding experimental data for the short-time collective diffusion coefficient in Ref. [Segrè *et al.*, 1995].

Increasing the concentration to $\phi = 0.5$, some q dependence is clear even in the values immediately around the first peak, and decay is much slower again, in agreement with Ref. [Segrè *et al.*, 1995]. Short- and long-time decay regimes seem to be apparent – by comparison with Ref. [Segrè and Pusey, 1996] the crossover appears to occur at around $q^2t = 100$, although the time-resolution of the data is not sufficient to properly observe the short-time regime at this volume fraction. The extracted long-time diffusion coefficient at the peak of $S(q)$, $q = 7.00$,

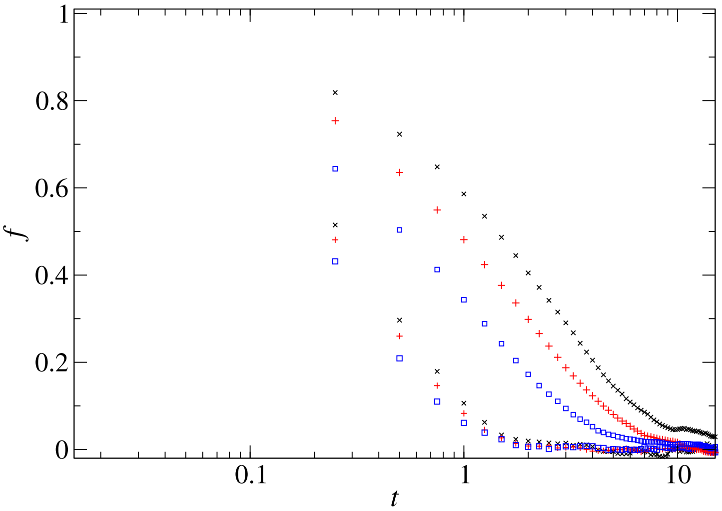

Figure 4.6: As FIG. 4.5, on a linear-linear plot with the time axis no longer scaled by q^2.

gives $D_0/D_{L,\,firstpeak}(\phi = 0.5) \approx 24.87$, rather far from the experimental value of ≈ 30 in Ref. [Segrè and Pusey, 1996] at $\phi = 0.465$. Since D at the first peak decreases with ϕ, D_0/D increases – therefore the discrepancy is even larger taking account the slightly higher volume fraction in my simulation. To explain this, it is worth discussing the hydrodynamic factor $H(q)$. This is related to the *short-time* diffusion coefficient $D_S(q)$ via $D_S(q) = D_0H(q)/S(q)$, and the long-time coefficient $D_L(q)$ is found empirically to scale approximately with $D_S(q)$ [Segrè and Pusey, 1996]. $H(q)$ is known, in hard sphere-like systems, to be substantially less than 1 and a decreasing function of ϕ, whereas it is equal to 1 in the absence of hydrodynamics. Therefore, one would expect at high concentration that D would be artificially high in the absence of hydrodynamics, as is found to be the case comparing my simulation with the experiments in Ref. [Segrè and Pusey, 1996].

The data also become somewhat more noisy, particularly after $q^2t \approx 300$ where the signal from correlations remaining after this time is very weak relative to statistical noise, as shown in FIG. 4.6. A significant exacerbation is the fact that the slower dynamics means fewer microstates are sampled in a given simulation time, so that the time average upon which the measurement relies departs from the true ensemble average. This makes getting good statistics in glassy systems particularly difficult, as seen in Section 4.5.

It should be noted that although the approximate agreement with experiment is reassuring, the lack of hydrodynamics in the simulations precludes more precise comparisons. Nevertheless, the data presented so far pertaining to dynamical correlations are physically convincing, internally consistent and in good agreement with experimental data. A full exploration of the wavenumber dependence of $D(q)$ is outside the scope of the present calibration measurements – for such information, Ref. [Segrè *et al.*, 1995] is an extremely lucid and comprehensive starting point.

4.3.2 The low-q limit

In the limit of low-q, the decay of the intermediate scattering function describes relaxation on macroscopic length scales. It is known that, in practice, this requires $q\langle d\rangle < 0.4$ at most [Russel *et al.*, 1992]. By measuring $D(q \rightarrow 0)$, one is in fact measuring the macroscopic collective or 'gradient' diffusion coefficient D_{coll} that appears in Fick's law, Equation 2.8. Reaching this limit in simulation is not trivial: from Equation 4.4, the characteristic correlation time t_{corr} scales as q^{-2}. Therefore, measurements separated by times less than this are not independent – statistical noise remains correlated on this timescale so that simulations many times longer than t_{corr} are needed to bring out the signal. Additionally, simulations must be larger (and therefore slower!) in order that low enough q can be reached.

For these reasons, the data are sufficient only for a qualitative illustration of the effect of moving towards the low-q limit. I consider the volume fractions $\phi = 0.1, 0.3$ and investigate the decay of $f(q,t)$ around the first peak in $S(q)$ and at the lowest q values attainable in simulations of $N = 5000$ particles, down to a minimum of $q = 0.37$.

FIG. 4.7 compares the results in these two regimes. As we saw in Section 4.3.1, decay around the first peak is slower at higher concentration. For the low-q regime, it is unclear whether the data have yet reached the limiting behaviour of $q \rightarrow 0$: D defined by the $\phi = 0.1$ data is $D_{coll}(\phi = 0.1) \approx 0.3$, somewhat higher than the expected first-order result for hard spheres, $D_0(1 + 1.45\phi) = 0.19$ [Russel *et al.*, 1992]. To further complicate matters, as q decreases, hydrodynamic effects play an increasingly important role [Cichocki and Felderhof, 1991], and are neglected in the simulations. Therefore, as discussed in 4.3.1, we would expect again an increased value of D_{coll} due to the neglect of hydrodynamics.

However, an important qualitative feature of the limit is apparent in FIG. 4.7: namely, the diffusion coefficient for low-q *increases* when ϕ is increased. This captures the fact that the macroscopic collective diffusion coefficient increases with concentration, as discussed in Section 2.3.4. This opposite dependence of the collective and self diffusion coefficients on ϕ is an important concept in Chapters 5 and 6 because bulk phase separation is achieved by collective rearrangement of particle matter, whereas fractionation processes rely on self diffusion. Although the neglect of hydrodynamics and difficulty in reaching the low-q limit mean that direct comparison with experiment in this regime is difficult, it is heartening that the qualitative features of structural relaxation in this limit are present.

4.3.3 Finite-size effects

The intermediate scattering function provides a way of measuring finite-size effects on the collective dynamics. For now, I will focus on the effect of reducing the step parameter Δ – corresponding to the mean size of a Brownian hop – from 0.02 (as is used unless otherwise mentioned) to 0.01. The possibility of system size effects on structural relaxation is explored in relation to glassy behaviour in Section 4.5.

For the $N = 768$ simulations, I have run simulations at volume fractions $\phi = 0.3, 0.5$ with the same parameters as before but with $\Delta = 0.01$. FIG. 4.8 shows the effects on the decay of the first peak. At $\phi = 0.3$, the data are in good agreement. At $\phi = 0.5$, decay with the smaller Δ appears to be slightly faster, for each value of q around the peak (note that, as discussed in

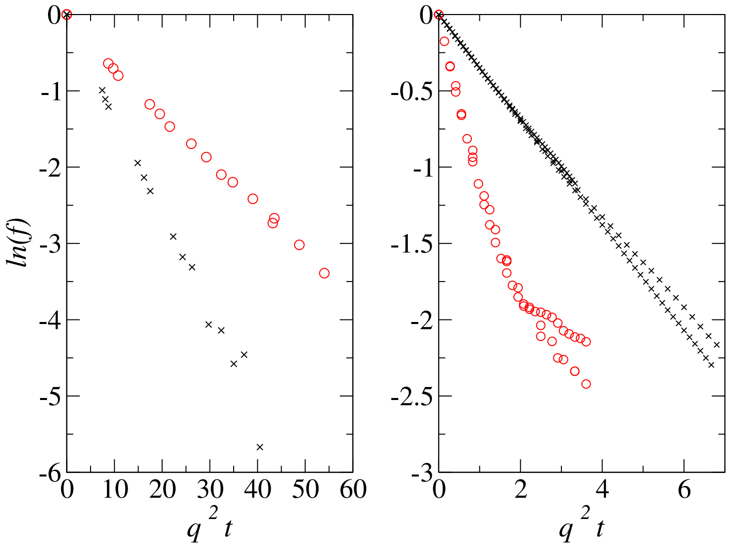

Figure 4.7: The decay of $f(q, t)$ for volume fractions $\phi = 0.1$ (black crosses) and $\phi = 0.3$ (red circles). The left pane shows data for $q = 5.45$ to 6.57 (around the first peak in $S(q)$), and the right pane shows $q = 0.37$ to 0.91.

Diffusion coefficient	Extracted value	Comparison	Comments
$D_{non\text{-}interact}$	$D = 0.165 \pm 0.005$	$D = D_0 = 0.16\dot{6}$	D_0 defined by choice of time units.
$D_{firstpeak}(\phi = 0.3)$	$D_0/D \approx 2.73$	$D_0/D(\phi = 0.3) \approx 2.7$ [Segrè *et al.*, 1995]	
$D_{L,\ firstpeak}(\phi = 0.5)$	$D_0/D \approx 24.9$	$D_0/D(\phi = 0.465) \approx 30$ [Segrè and Pusey, 1996]	No hydrodynamics in simulations.
$D_{coll}(\phi = 0.1)$	$D \approx 0.30$	$D \approx 0.19$ [Russel *et al.*, 1992]	Comparison from first-order in ϕ theory.

Table 4.1: Diffusion coefficients extracted from measured intermediate scattering functions, compared with appropriate values from the literature, as discussed in the text.

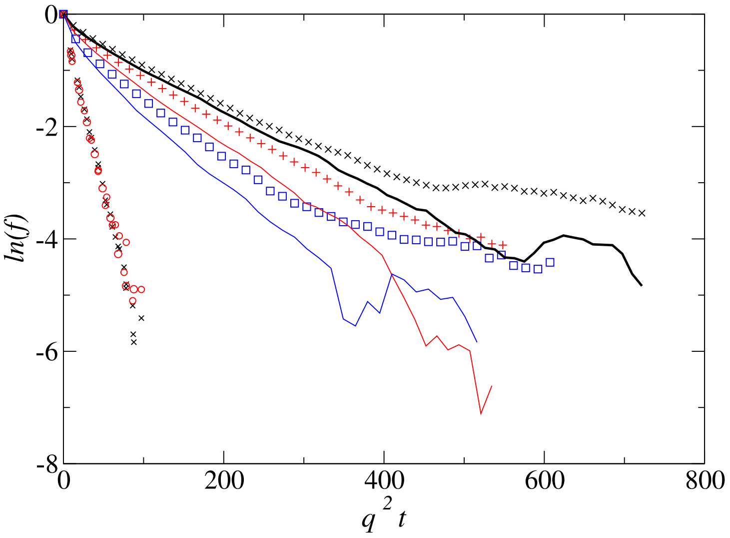

Figure 4.8: As FIG. 4.5, with data from simulations with smaller Δ included. At $\phi = 0.3$, all q overlay so are shown as black crosses, with red circles indicating the $\Delta = 0.01$ data. At $\phi = 0.5$, the symbols are as in FIG. 4.5 with the corresponding coloured lines indicating the $\Delta = 0.01$ data.

Section 4.3.1, the data beyond $q^2 t \approx 300$ are dominated by noise, so discrepancy beyond this point should not be attributed to finite-size effects). Although only a quantitative effect (i.e. the order of decay speeds with q is the same, as is the qualitative difference between $\phi = 0.3$ and $\phi = 0.5$), this illustrates the intuitively reasonable idea that finite-size effects in Δ would become more important at higher volume fraction: in the dilute limit, there is no length scale with which to compare Δ as particles are infinitely well separated, so its value is unimportant. The results show that, at high concentrations, small quantitative finite-size effects may come into play in the collective particle dynamics. Simulation findings in that regime in particular must therefore be treated cautiously and checked for dependence on Δ.

4.4 Self diffusion

I now move from collective dynamics to studying individual particle dynamics in terms of self diffusion.

4.4.1 Mean squared displacement and finite step size

According to Equation 2.7, the self diffusion coefficient can be measured by observing the variation of the mean squared displacement (MSD) of particles through time. In the long-time limit this is governed by the long-time self diffusion coefficient D_{self}^{long}. As discussed in Section 3.2.1, the short-time coefficient D_{self}^{short} is a simply a defined property of each particle, depending on the individual particle's size. It can only depend on ϕ via hydrodynamic interactions with its neighbours, which are not included in the simulation.

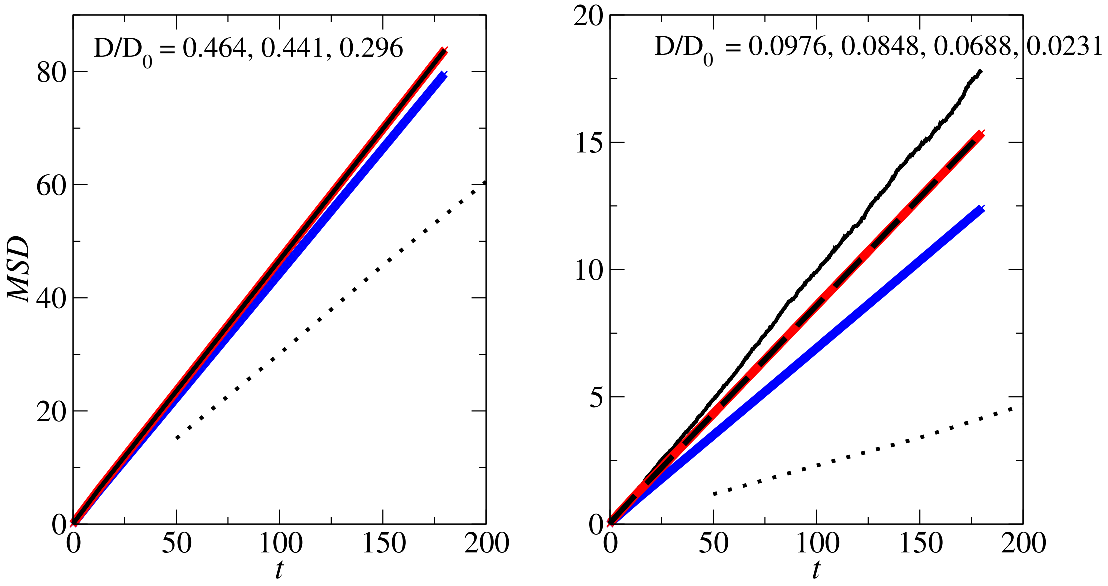

Figure 4.9: Mean squared displacement at $\phi = 0.3$ (left) and $\phi = 0.5$ (right). My data with $\Delta = 0.02$ is shown in blue, $\Delta = 0.01$ in red. The solid black lines are from the Brownian Dynamics simulations of Ref. [Sanz and Marenduzzo, 2010]. The dashed black line is Kinetic MC data with $\Delta = 0.01$ from Ref. [Sanz and Marenduzzo, 2010], and the dotted line is for $\Delta = 0.1$ from the same source. At $\phi = 0.3$ the dashed black line overlays the Brownian Dynamics data so is not shown for clarity. The extracted long-time diffusion coefficients (fitted after $t = 100$) of the curves with distinct gradients, in decreasing order, are shown on the plots.

I first measure the overall MSD in systems again with $N = 768$, $\sigma = 0.06$ and $\phi = 0.3, 0.5$. I also take this opportunity to run simulations at each volume fraction with the mean step size $\Delta = 0.01$ instead of 0.02, in order to check the effect of step size on self diffusion, prompted by Ref. [Sanz and Marenduzzo, 2010] in which comparison was made between Kinetic Monte Carlo (MC) and Brownian Dynamics simulations for different values of the MC step size. For my measurements, the MSD up to $t \approx 200$ is measured, averaging together results from 40000 starting points (i.e. definitions of $t = 0$ throughout the simulation) spanning 10000 time units.

The degree of agreement between Kinetic MC and Brownian Dynamics methods depends on the use of small enough MC step sizes that most moves are accepted (empirically, around 70% or more) [Romano *et al.*, 2011]. It has also been found that rescaling time by the acceptance probability improves agreement further [Romano *et al.*, 2011; Sanz and Marenduzzo, 2010]. In this work I do not attempt such a rescaling procedure – it is more suitable where acceptance probabilities (therefore concentrations) are uniform throughout the system, whereas I will be concerned with e.g. gas-liquid separation in which large differences in local concentration exist, which could not be accounted for with an overall time rescaling. (Besides, I am investigating the features of phase transition kinetics in a self-consistent model system, as opposed to trying to precisely match experimental diffusion coefficients or other dynamical quantities, which in any case would require the inclusion of hydrodynamics.) For this reason it is important that Δ

is sufficiently small to achieve reasonably close agreement with Brownian Dynamics in terms of diffusion rates at different ϕ, without rescaling by acceptance probability – i.e. it is not possible to *locally* rescale diffusion rates in order to improve data from a too-large step size.

To this end, in FIG. 4.9 I compare MSD data from my simulations at $\phi = 0.3, 0.5$ with Brownian Dynamics and Kinetic MC data from Ref. [Sanz and Marenduzzo, 2010] (not rescaled by the acceptance probability). Good agreement at $\phi = 0.3$ is observed, with diffusion slightly slower in the $\Delta = 0.02$ case. At $\phi = 0.5$, as for the intermediate scattering function, step size dependence is increased, with the Brownian Dynamics and Kinetic MC data splitting even for the smallest MC step size. For equivalent step size, my data coincides with that from the Kinetic MC implementation of Ref. [Sanz and Marenduzzo, 2010]. It is interesting to note that the polydispersity in my simulations, $\sigma = 0.06$, along with the extremely steep but not quite hard sphere potential used by Sanz, does not seem to strongly affect agreement between our simulations.

The data concerning overall MSD confirm the expected picture of much slower self diffusion as concentration increases, and agree with the finding in Section 4.3.3 that finite step size effects are more influential at higher concentration. Due to the good agreement at $\phi = 0.3$ and relatively good agreement at $\phi = 0.5$, and the fact that the simulation speed is proportional to $1/\Delta^2$, I choose $\Delta = 0.02$ as a compromise between speed and precision in general. This compromise seems reasonable, given that in many cases I will not be dealing with fluid concentrations greater than $\phi \approx 0.4$, and that any findings of interest are subsequently checked for dependence on step size anyway.

4.4.2 Mean squared displacement with square well interaction

In addition to the hard sphere interaction, this thesis also includes work on the square well system in which particles are surrounded by an attractive well of range λ and constant depth u. Further details appear in Chapters 5 and 6. In order to gauge step size effects when this interaction is present, FIG. 4.10 shows mean squared displacement in a square well system of $\lambda = 1.15$ and $u = 1.54$ (in units of k_BT) at volume fraction $\phi = 0.3$, with step sizes $\Delta = 0.01, 0.02$. Also shown are the corresponding hard sphere cases. As discussed in Chapter 6, these parameters place the system in the metastable fluid region of the phase diagram, above the critical effective temperature for gas-liquid separation – that is, barring crystallisation (which is not observed on the simulated timescale), the system is in a steady state so that MSD data can be time-averaged. As can be seen in the figure, the square well attraction slightly slows diffusion relative to the hard sphere case. Further, it can be seen that the relative discrepancy in MSD between the smaller and larger step sizes is magnified compared to the hard sphere case.

It is known that the Kinetic MC algorithm, when attractive interactions are present, can result in kinetic ‘trapping’ whereby the collective movement of a mutually attracting cluster of particles is suppressed because trial moves are applied to each particle *individually.* If attractions are strong (in comparison to k_BT) and short ranged, each particle on its own has a very small chance of moving an appreciable distance from the cluster, and as particles are moved one-by-one in the Kinetic MC algorithm, the cluster as a whole becomes unphysically trapped. Further details and a method for mitigating the problem are given in Ref. [Whitelam and Geissler, 2007]. However, such effects are less important as attraction strength is decreased, and in Ref. [Romano

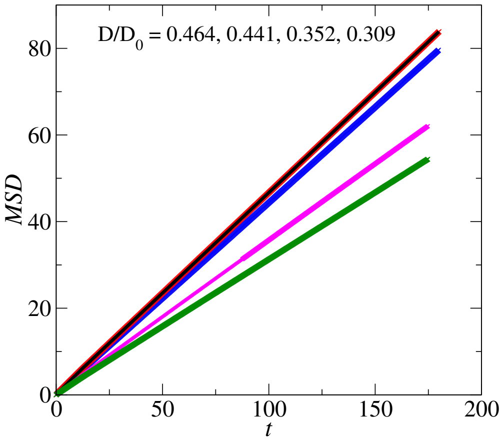

Figure 4.10: Mean squared displacement at $\phi = 0.3$ with a square well interaction of range $\lambda = 1.15$ and depth $u = 1.54$, as used in Chapter 6. Pink shows data for $\Delta = 0.01$, green $\Delta = 0.02$. The corresponding hard sphere data (i.e. $u = 0$) from the left panel of FIG. 4.9 are also shown. The extracted long-time diffusion coefficients (fitted after $t = 100$), in order of decreasing gradient, are shown on the plot.

et al., 2011] it is shown that Kinetic MC agrees reasonably well with Brownian Dynamics up to $u \approx 6.6$, far higher than is used in the present work. I therefore expect any artefact of this nature to be negligible, especially in comparison to e.g. the neglect of hydrodynamic interactions in simulations.

4.4.3 Self diffusion variation with particle size

As can be seen from Equation 2.6, the diffusion coefficient for a particle depends inversely on its size. This dependence may be thought of as the most basic hydrodynamic effect of a particle's size – although my simulations ignore the solvent degrees of freedom (i.e. hydrodynamic interactions between particles are neglected), this basic interaction between the solvent and each individual particle *is* incorporated by adjusting the mean step size for each particle based on its diameter. This was discussed in Section 3.2.1. This gives the correct dependence, in the absence of hydrodynamic interactions between particles, of D_{self}^{short} upon particle size.

As was already seen, the overall long-time diffusion coefficient D_{self}^{long} is reduced on increasing ϕ. It is reasonable to expect that this effect would also depend on particle size, since the slowing effect of increasing ϕ is due to collisions with neighbours, which may be better avoided by smaller particles. This is in addition to the fundamental dependence of particle step size upon diameter. To measure this effect I use the simulations from Section 4.4.1, with an additional simulation at $\phi = 0.1$ (the step parameter is $\Delta = 0.02$, which should be taken for granted henceforth unless specified). I split particles of diameters $0.95 < d < 1.05$ into five equal bins. FIG. 4.11 shows the results. The dependence of diffusivity on inverse size remains approximately linear, but it is clear that the strength of the dependence increases with concentration, illustrating that size-dependent geometric effects enhance the size-dependence of diffusion, over and above the programmed-in effect on a particle's step size.

The effects of polydispersity upon the *overall* diffusion rates of the system [Yiannourakou *et al.*, 2010] seem to be more studied in the literature than that on the individual diffusivities of size species. One exception, for a Molecular Dynamics simulation of a Lennard-Jones fluid, is Ref. [Murarka and Bagchi, 2003]. There, for low temperature near-glass parameters, non-linear dependence of diffusivity on inverse size is found, and it is argued that this is due to the dominance of the cage-hopping diffusion mode – smaller particles would be able to hop more frequently in addition to being intrinsically 'faster' due to their size, leading to the non-linearity observed. This is perhaps the limiting case of the trend I have observed, in which geometric effects become gradually more important with increasing concentration but the dependence on inverse diameter remains approximately linear. For further information, I refer the interested reader to Ref. [Murarka and Bagchi, 2003] and references therein.

Note that, as for phase equilibria, theoretical treatment of diffusion is made complicated by polydispersity, if one wishes to retain information about its size dependence. This is clear in Chapter 8, where the dependence of diffusion on particle size is the central component of a dynamical theory of fractionation but requires a drastic approximation to be made. In Chapter 10, an approach which simplifies a well-grounded treatment of polydisperse diffusion is presented.

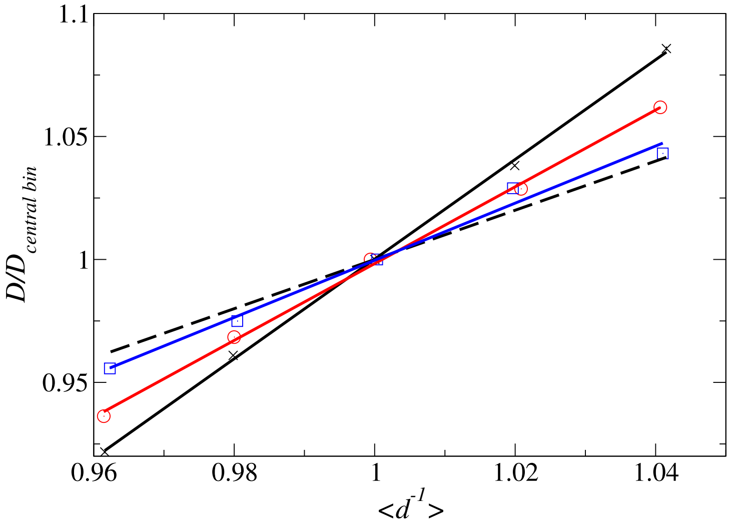

Figure 4.11: Long-time self diffusion coefficient in particle size bins plotted as a function of the mean inverse diameter of particles in the bin, normalised by the coefficient in the central bin. The size bins cover d from 0.95 to 1.05 in five equal intervals. Black crosses, red circles and blue squares correspond to $\phi = 0.5, 0.3, 0.1$ respectively and lines are linear fits, with gradients of 2.03 ± 0.03, 1.56 ± 0.04 and 1.16 ± 0.07. The dashed black line shows the dilute limit, i.e. the dependence resulting from the variation in step size with diameter, in the absence of hard sphere interactions, and has a gradient of unity by definition.

4.5 Dynamical heterogeneity and finite system size

The behaviour of model colloids at very high concentrations is relevant to the study of glassy behaviour, i.e. the structural and dynamical properties of amorphous, dynamically arrested systems [Scheffold and Mason, 2009]. Although glassy systems behave mechanically as solids, the transition from metastable fluid to glass (occurring around $\phi \approx 0.58$) does not seem to be accompanied by a growing structural correlation length like the long-range order in a crystalline solid. Instead, it seems that the glass transition is accompanied by the growth of a *dynamical* correlation length – regions of the system become correlated in terms of their local relaxation rates, and these regions grow in size with increasing density [Berthier, 2011]. Polydispersity is intimately related to this area, since it inhibits the formation of crystals [Schöpe *et al.*, 2007], thus facilitating the formation of high density amorphous states [Auer and Frenkel, 2001; Weysser *et al.*, 2010]. Glasses of low enough polydispersity do however seem to be able to crystallise via cooperative local rearrangements in which individual particles only move on average less than their own diameter [Zaccarelli *et al.*, 2009].

In my work, I have not focused on glassy behaviour. Rather, I have been more interested in how polydispersity affects the kinetics of transitions between non-frustrated phases, particularly in terms of fractionation which, like all dynamical processes, would tend to be frustrated in glasses. Furthermore, at extremely high densities, the Kinetic MC method becomes less suitable than more usual Molecular Dynamics methods [Middleton and Wales, 2004], not least because the decreased inter-particle spacing in a near-jammed system would necessitate a very small MC step size [Sanz and Marenduzzo, 2010] – the trend in this direction was indeed apparent as concentration was increased in Section 4.4.1.

Nevertheless, and bearing these caveats in mind, I now present results illustrating an apparent link between glassy behaviour and finite system size effects. I acknowledge helpful discussions with J. Mattsson regarding the physics described in this section [Mattsson, 2012].

A simulation with periodic boundaries can be used to find the structure factor $S(q)$ for $2\pi/q > L/2$ where L is the side length of the simulation cell. That is, when the length scale corresponding to the value of q becomes larger than half the size of the cell, measuring correlations over that length scale no longer makes sense because correlations of wavelength $\sim L$ are spuriously introduced by the periodic boundary conditions.

This is illustrated in FIG. 4.12 for hard sphere fluids of $\sigma = 0.06$, $\phi = 0.51, 0.59$, for two system sizes $N = 768$ and $N = 3000$. I have discarded the first two q values in each direction. At each volume fraction, the first peaks of the structure factor and a significant region before them are well out of range of system size artefacts, as the peaks correspond to correlations on the scale of the particle diameter d – the two system sizes therefore produce the same results, and would only deviate if we were measuring correlations on length scales approaching L in the smaller simulation. This is the case even in the glassy $\phi = 0.59$ system, although the statistics here are worse due to the slowing dynamics, as discussed in Section 4.3.1.

FIG 4.13 shows the decay of $f(q, t)$ at the first peak in $S(q)$. In the 'normal', non-glassy fluid at $\phi = 0.51$, the *decay* of the first peak in $S(q)$ is also unaffected by changing between the two simulation sizes, just as was the structure factor as a whole. However, at $\phi = 0.59$, the decay of the first peak *is* affected by system size, even though the structure factor is not. More precisely, the decay at $N = 3000$ is far slower than at $N = 768$. To check that this is not an artefact of

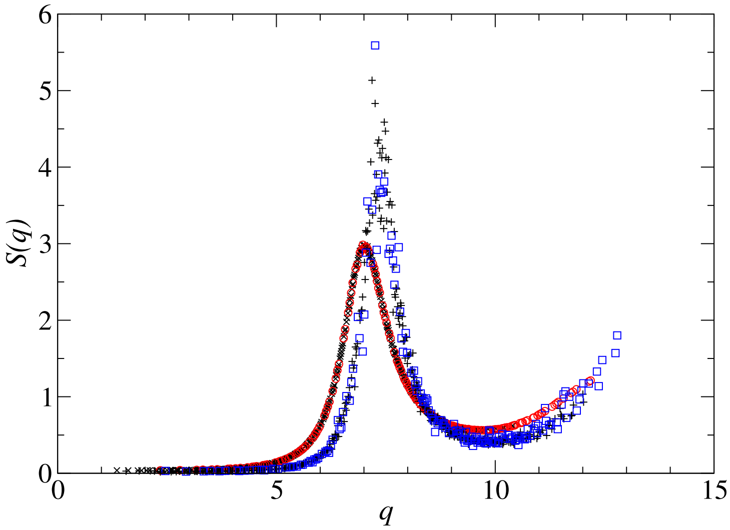

Figure 4.12: Structure factors at $\sigma = 0.06$ for ($\phi = 0.51, N = 768, 3000$) (black crosses and red circles respectively) and ($\phi = 0.59, N = 768, 3000$) (black plusses and blue squares). At each value of ϕ, the two system sizes produce the same results, because the first peak in $S(q)$ corresponds to wavelengths much smaller than either system size.

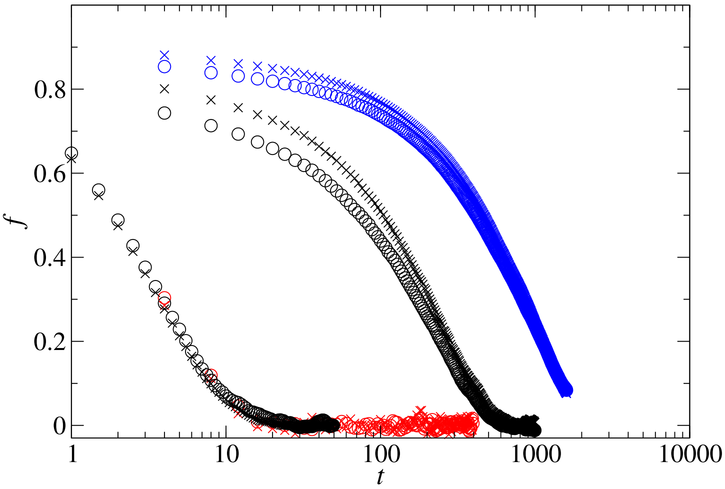

Figure 4.13: Decay of the first peak in $S(q)$ at $\sigma = 0.06$. Colours indicate system size N, as in FIG. 4.12. The faster decaying, overlaying curves are for $\phi = 0.51$ while the slower four series are for $\phi = 0.59$. For each system, two closely spaced q values around the first peak are shown as circles and crosses – circles are higher in q than crosses by ~ 0.3. At $\phi = 0.59$, there is a marked dependence on system size which is absent at $\phi = 0.51$, even though the first peak in $S(q)$ corresponds to wavelengths far smaller than the system size in both cases.

particular trajectories, the data at $\phi = 0.59$ are averaged over at least two independent runs of times ~ 30 times the decay length presented, and the qualitative dependence on system size is checked for repeatability between individual runs.

This suggests that, at $\phi = 0.59$, the decay of correlations on a length scale $\sim d$ is slaved to some dynamical correlation length which, in the case of the smaller system, becomes comparable with the simulation box size. The difference between the two volume fractions therefore seems to be a manifestation of dynamical heterogeneity, the finite-size effects arriving at higher volume fraction because heterogeneities approach the simulation size. Since correlated regions involving up to ~ 100 particles are possible in glassy systems [Berthier, 2011] and $L \approx 9$ in the $N = 768$ simulation, this seems a plausible explanation.

These results show that, in fluids up to around $\phi = 0.51$, the system size does not affect the structural or dynamical behaviour on length scales shorter than the system size, as expected. However, if the volume fraction is raised so that glassy behaviour and dynamical heterogeneities come into play then, even while static structural behaviour is unaffected by system size, the decay of the first peak in $S(q)$ appears to become coupled to longer-scale dynamical correlations, which set a lower limit on how small the system can be before the periodic boundaries spuriously affect the decay of the first peak. This is a means of estimating the extent of dynamical heterogeneities via the exploitation of finite-size effects, in addition to the four-point correlation functions more commonly used [Berthier, 2011].

4.6 Conclusions

In this chapter I have presented a wide variety of structural and dynamical measurements in order to calibrate my simulations with other work and briefly overview important features. I first measured the static structure factor $S(q)$, showing the increased structuring and closer packing associated with higher volume fractions and observing good agreement with the well-known Percus-Yevick approximation. I then showed the effect of a moderate degree of polydispersity on $S(q)$, and generalised to partial structure factors to show that the apparently reduced structuring of polydisperse systems is due to destructive interference between peaks on different length scales, caused by the spread of particle sizes. These measurements, agreeing with existing work in the literature, confirm that the simulations are correctly exploring configuration space so as to reproduce the structural features of polydisperse metastable and equilibrium fluids.

I then moved to the intermediate scattering function, which measures the decay of structural correlations through time. I focused first on the decay of the first peak in $S(q)$, checking the accuracy of the measurement code in the ideal non-interacting system. Moving to finite concentration, I observed qualitative agreement with experimental and simulation data from the literature, noting that, unlike static structural features, these dynamical aspects depend on hydrodynamic interactions, which are not included in my simulation. I moved to the low-q (macroscopic) limit, qualitatively resolving the fact that the macroscopic collective diffusion coefficient D_{coll} increases with ϕ, whereas relaxation speed at smaller length scales decreases with increasing ϕ. I finally checked the dependence of the decay of the first peak on the MC step size, finding that finite step size effects increase with ϕ, as expected due to the decreasing inter-particle spacing.

I next investigated self diffusion, measured by observing the evolution of mean squared displacement through time, particularly in relation to finite step size effects. I found exact agreement between my Kinetic MC algorithm and a similar one from the literature for equivalent step size. As for the collective dynamics, finite step size effects in self diffusion increase with ϕ. I settled on a step parameter $\Delta = 0.02$, in view of the relatively small magnitude of finite size effects at $\phi = 0.5$, the upper end of the fluid concentrations I have typically studied. Also bearing on this compromise is the fact that the simulation speed scales as $1/\Delta^2$, and that individual findings of interest will be checked as necessary for dependence on step size anyway.

I then measured size-dependent self diffusion by binning particles according to size, observing that the fundamental Stokesian dependence on diameter is enhanced by geometric (cage-like) effects on increasing ϕ.

I finally presented a very brief study of a near-glass state point ($\phi = 0.59$), finding that the decay of the first peak in $S(q)$ in this regime may become coupled to dynamical relaxation on a longer length scale, causing finite system size effects when this longer length exceeds the system size. This coupling was not present at $\phi = 0.51$, and at both volume fractions the *static* structure factor around the first peak was independent of system size. I argued that this is a manifestation of dynamical heterogeneity in near-glass states, referring to studies in the literature, and speculated that exploitation of this finite-size effect could be used to estimate the extent of dynamical heterogeneities, in addition to the usual methods involving four-point correlation functions.

The structural and dynamical physics of my simulations seem to be in broad agreement with the literature, this agreement being quantitative where it can reasonably be expected to be so. I have delved into more detailed measurements relevant to polydisperse systems, observing the expected results.

Chapter 5

Results II – Spinodal fractionation

This chapter concerns simulations of gas-liquid separation in polydisperse systems. For the first time, fractionation is observed during the very early stages of phase separation of a polydisperse fluid; some selectivity with respect to phase composition is possible even as the system is rapidly evolving. Additionally, I find that the direction of fractionation can depend qualitatively on the choice between two very similar inter-particle potentials which, in the monodisperse case, would be identical.

5.1 Introduction

Gas-liquid (G-L) phase separation is possible in systems for which there is a sufficient attractive component to the inter-particle potential. From one point of view, it is the simplest form of phase separation in colloids, in that the gas and liquid phases are described by the same 'fluid' branch of the free energy, whereas a crystal phase (with its significantly different qualitative structure) occupies its own branch.

As discussed in Chapter 2, the majority of existing work on polydispersity concerns its effects on the equilibrium properties of coexisting phases, whereas the dynamical processes by which polydisperse systems progress to equilibrium are less well understood. Moreover, in such studies as do exist, polydispersity is often introduced only to inhibit crystallisation so that high density amorphous states can be studied [Auer and Frenkel, 2001; Weysser *et al.*, 2010].

Fractionation is thought to be involved in the nucleation of crystals from polydisperse fluids [Martin *et al.*, 2003, 2005; Schöpe *et al.*, 2007], with frustration of this mechanism held responsible for the experimental inaccessibility of some complex equilibrium phase diagrams [Liddle *et al.*, 2011]. In practical applications, fractionation of particle species between separating phases may be intended, for example to purify a mixture, or unintended. In either case, it is important to consider consequent effects on related properties (composition, crystal lattice parameter etc.) in the resultant phases.

In this chapter I address the issue by dynamically simulating G-L spinodal decomposition in a size-polydisperse colloidal fluid. I observe fractionation phenomena on simulation timescales,

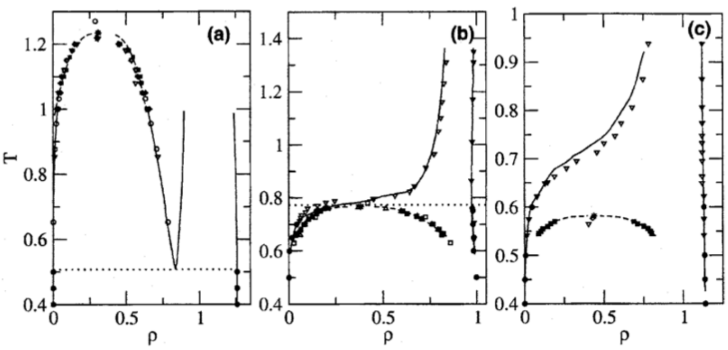

Figure 5.1: Phase diagrams reproduced with permission from Ref. [Liu *et al.*, 2005] for (a) $\lambda = 1.5$, (b) $\lambda = 1.25$ and (c) $\lambda = 1.15$ in the density-temperature plane. With my choice of units, $T = T_{\mathit{eff}} \equiv 1/u$ and $\rho = 6\phi/\pi$. As λ decreases, the G-L binodal becomes metastable with respect to crystal-solid separation. The relevant diagram for my system is (c).

and uncover a surprising example of the sensitivity of such phenomena to the details of particle interactions, which I explain by appealing to a perturbative equilibrium theory. In Section 5.5 I extend the study further by using a Schulz particle size distribution and going to higher polydispersities, finding that the skew of the Schulz distribution causes fractionation in diameter *variance*, in addition to its mean, between the phases.

5.2 Simulation setup

The fundamentals of the simulation method are outlined in Chapter 3. The specifics are as follows. The system studied in this work consists of N spherical particles whose diameters d_i are drawn from some size distribution of polydispersity σ. The simulation cell is cubic with periodic boundaries. I set the length unit as the mean hard core diameter $\langle d \rangle$. In order to mimic the qualitative features of any attractive interaction in the colloid, the hard particle cores are surrounded by a pairwise additive square well potential of (mean) range $\lambda = 1.15$ and depth u, where the energy unit is $k_B T$. In the monodisperse limit, these parameters and the overall parent volume fraction $\phi_p = 0.229$ place the system in the middle of the G-L coexistence region as shown in FIG. 5.1 (c), where phase separation would be expected to proceed by spinodal decomposition (see Section 2.1). For the value of λ used here, the resulting G-L coexistence is metastable with respect to fluid-solid phase separation which, on the timescale simulated here, I do not observe. The time unit is t_d. The step parameter Δ is set to 0.02 as normal – at the gas and liquid concentrations under investigation finite size effects should be small, as discussed in Section 4.4.1 – but I have also checked the fractionation measurements while reducing Δ and found no effect within error.

Note that there is a choice to be made in precisely how the specific attraction range between any two particles i and j should depend upon their size. I have simulated two possibilities. The 'scalable' dependence is motivated by its simplifying effect on the relevant theory viz. the equilibrium effects of fractionation [Evans, 2001]:

$$V_{scal}(r) = \begin{cases} \infty & \text{if } r \leq d_{ij} \\ -u & \text{if } d_{ij} < r \leq \lambda d_{ij} \\ 0 & \text{if } r > \lambda d_{ij} \end{cases} \tag{5.1}$$

where $d_{ij} = (d_i + d_j)/2$. The 'non-scalable' dependence is chosen for its similarity to depletion attractions [Lekkerkerker *et al.*, 1992], in which the attraction range depends on the hard core diameter via an additive constant, e.g. the radius of gyration of a dissolved polymer:

$$V_{non\text{-}scal}(r) = \begin{cases} \infty & \text{if } r \leq d_{ij} \\ -u & \text{if } d_{ij} < r \leq d_{ij} + (\lambda - 1) \\ 0 & \text{if } r > d_{ij} + (\lambda - 1) \end{cases} . \tag{5.2}$$

In the monodisperse case, $d_{ij} = 1$ for all particle pairs so that the two definitions are strictly identical. Even in the polydisperse case the *mean* attraction range is unaffected by the choice of dependence, but it will nevertheless turn out that this easily-overlooked detail can make a radical difference to the fractionation of the system.

The Kinetic Monte Carlo algorithm used allows large systems to be simulated on a physically relevant timescale, necessary for measurement of the required statistics for fractionation. However, as in comparable Molecular or Brownian Dynamics algorithms, hydrodynamic interactions (HI) are neglected. Given the importance of HI in spinodal decomposition, this requires some comment. In an approximate theoretical scheme, for example, HI have been shown [van de Bovenkamp and Dhont, 1994] to enhance longer-wavelength decomposition while suppressing shorter wavelengths. Measuring the effect of such corrections on fractionation in a polydisperse system would require for instance explicit treatment of the solvent via Molecular or Stokesian Dynamics, or a combined Brownian Dynamics/lattice Boltzmann method, considerably increasing computational expense.

Nevertheless, I will show that the phase separation observed here is at least qualitatively representative of spinodal decomposition in a real system and expect that, particularly given that HI can be screened (e.g. by direct interactions in dense systems [Riese *et al.*, 2000]), it constitutes a relevant and important baseline in the low-HI limit. Given this, and since the objective is to collect good statistics for the observation of fractionation, I leave the incorporation of HI to future work.

5.3 Spinodal decomposition

The system is initialised as an amorphous fluid, whereafter the square well attraction is turned on (defining an instantaneous quench at $t = 0$) and the system allowed to evolve according to the Kinetic Monte Carlo algorithm. In addition to direct observation through time (FIG. 5.2), the development of a low-q peak in the structure factor $S(q)$ is used to track the G-L phase separation (FIG. 5.3). As shown in FIG. 5.3, this peak has already started to shift to lower q after 2 or 3 t_d, indicating coarsening of the G-L texture. To confirm that the phase separation is

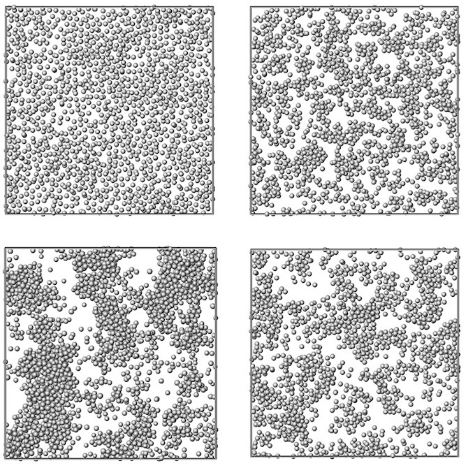

Figure 5.2: Illustrative simulation snapshots at (clockwise from top-left): $t = 0$, $t = 4$, $t = 64$ and $t = 760$ for a slice 3 units wide through the system, showing the coarsening spinodal texture.

spinodal in nature, FIG. 5.4 shows the integrated area under the peak increasing on a timescale of hundredths of t_d. No 'lag' period associated with nucleation-driven phase separation is detected, so I conclude that the phase separation is indeed spinodal in nature, as expected for these parameters [Liu *et al.*, 2005].

Other dynamical observations, presented later (Chapter 6), are in agreement with Ref. [Fasolo and Sollich, 2005] in suggesting that, as in the monodisperse case, the G-L coexistence resulting from the spinodal decomposition is in turn metastable with respect to the growth of a crystalline phase that would coexist with a tenuous gas. On the timescales observed here and in the absence of a template to seed nucleation, no crystals are observed.

5.4 Fractionation

5.4.1 Measuring fractionation

To check for fractionation, a per-particle distinction between the gas and liquid phases is required. For now, I count a particle as liquid if it has more than 12 neighbours within two length units, this threshold corresponding to a local number density in the middle of the G-L coexistence region. Later, in Section 5.4.4 I outline a more complex method of calculating fractionation, used to check for measurement artefacts in the simpler technique. Throughout the simulation the mean diameters in each phase, $\langle d\rangle_{gas}$ and $\langle d\rangle_{liquid}$, are recorded.

5.4.2 Results

In this section I present fractionation data from four simulations with $N = 30000$, where two independent initial configurations have been simulated according to both the scalable and non-scalable definitions of the inter-particle potential, with a Bates (pseudo-Gaussian) particle size distribution of $\sigma = 0.06$. In addition, I ran 20 smaller simulations ($N = 4000$), 10 for each

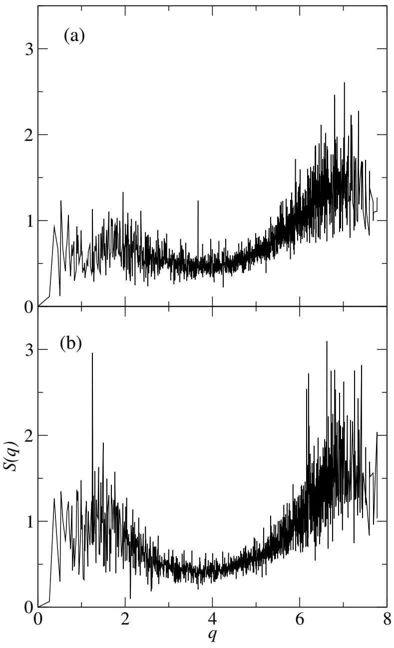

Figure 5.3: Structure factor $S(q)$ for one simulation trajectory, averaged over a) $t = 1-2$ and b) $t = 3-4$. The usual fluid peak is present at $q \approx 7$, while a peak corresponding to the gas-liquid texture appears at $q \approx 2$ and shifts to a lower wavenumber as it grows, indicating coarsening of the two phases.

potential, to confirm the qualitative nature of the fractionation effects, check for sensitivity to initial configuration or system size, and investigate longer timescales. Finally, having confirmed fractionation in the near-monodisperse regime, I ran two simulations at a higher polydispersity $\sigma = 0.2$ and $N = 4000$ in which the strength of fractionation, scaling with σ^2 [Evans, 2001], should increase by an order of magnitude relative to the other simulations.

FIG. 5.5 shows the evolution of $\langle d \rangle_{gas}$ and $\langle d \rangle_{liquid}$ in the large system for the scalable and non-scalable potentials $V_{scal}(r)$ and $V_{non\text{-}scal}(r)$. Fractionation of particle sizes is apparent after the first few hundred time units, although for the non-scalable potential the signal is apparently present at very early times (around $50t_d$). It should be noted, however, that at these times the system has yet to properly separate into distinct gas and liquid regions, as shown in FIG. 5.2, so that the signal may be significantly influenced by the presence of many particles whose neighbour counts lie on the threshold between 'gas' and 'liquid' according to this definition. For such 'threshold' particles, a marginally greater size may exclude neighbours from the local region, leading to the particle being counted as gas, whereas at later times, the system has successfully formed distinct phases with neighbour counts very much less and very much higher than the threshold. This potential measurement artefact and an alternative method used to corroborate the data are discussed in Section 5.4.4 but any effect, if present, seems to be minor for the polydispersities studied here, in that opposite fractionation effects between the two potentials are resolved once the system has formed distinct phases. For this reason, though, I refrain from drawing firm conclusions from the fractionation signal at early simulation times.

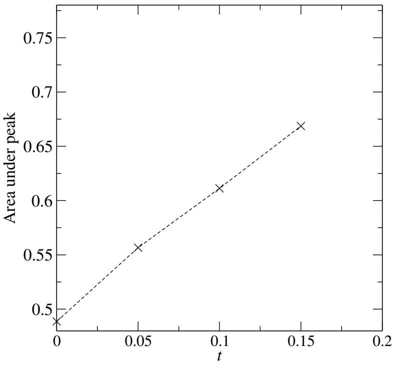

Figure 5.4: The area under the low-q peak, found by integrating from $q = 0.26$ to $q = 2.56$, averaged over 3 independent quenches, for times immediately after the quench. Although background noise, indicated by the value at $t = 0$, is strong, the area under the peak is clearly increasing from these very early times, consistent with spinodal decomposition. The dashed line is a guide to the eye.

As shown in FIG. 5.6, the fractionation in the smaller, longer-time simulations agrees within error with the data from the larger simulations. Thereafter, the fractionation in these simulations shows no detectable change beyond $t \approx 2500$ up to the maximum time simulated. FIG. 5.7 confirms the same qualitative fractionation at higher polydispersity, with the difference in mean size between the phases increased by approximately an order of magnitude as compared to the $\sigma = 0.06$ case.

Since fractionation relies on self-diffusion of individual particles between the phases, whereas bulk phase separation takes place via faster, collective rearrangement of particle density [Warren, 1999], it might be expected that the speed of spinodal decomposition would negate any possibility of fractionation in the early stages of phase separation, if the particles could not rearrange themselves to the required degree faster than the spinodal texture coarsens. Therefore, while the results leave open the question of further ‘ripening’ of the phase composition on much longer timescales, it is significant that fractionation is apparent almost immediately after spinodal decomposition begins.

More strikingly, there is a *qualitative dependence* on the choice of inter-particle potential – the non-scalable potential causes the liquid to accept smaller particles, whereas the scalable potential causes the smaller particles to be incorporated into the gas. Recall that $V_{scal}(r)$ and $V_{non\text{-}scal}(r)$ have the same well depth and mean range – they differ only in the *precise dependence* of their range upon the mean diameter d_{ij} of the particles between which the potential is acting. The two potentials coincide by definition in the monodisperse case (for which they produce exactly identical simulation trajectories), but are ‘split’ by a mild degree of polydispersity so as to produce opposite fractionation effects.

5.4.3 Equilibrium theory

In order to explain this surprising dependence on inter-particle potential, I now make use of a perturbative theory of polydisperse phase equilibria [Evans, 2001]. Although the system is of

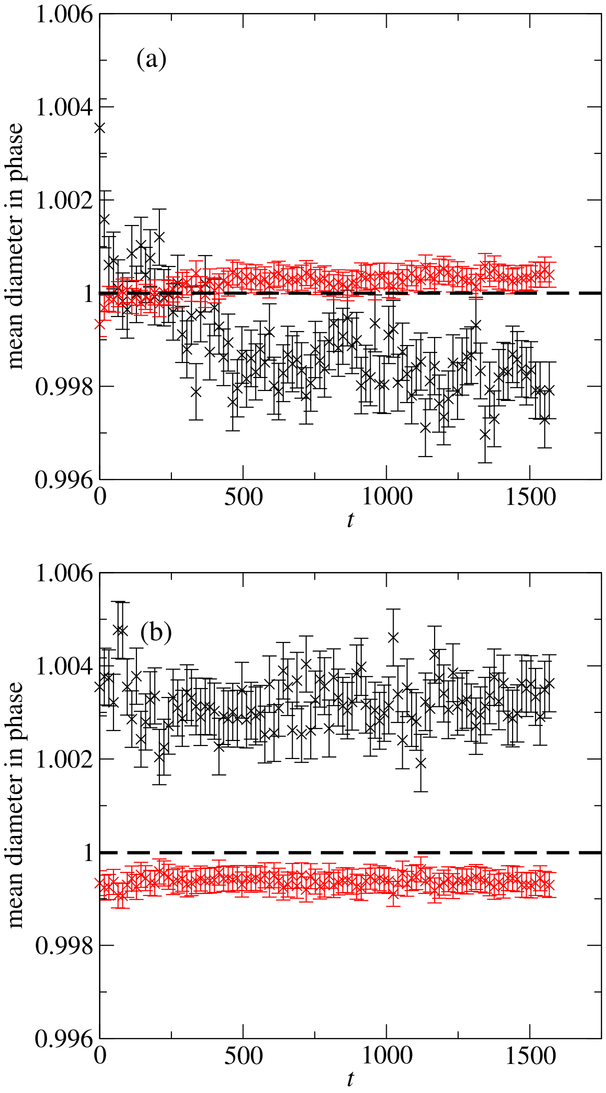

Figure 5.5: $\langle d \rangle_{gas}$ (black, large error bars) and $\langle d \rangle_{liquid}$ (red, small error bars) through time for the two definitions of the inter-particle potential, for four simulations with $N = 30000$. The parent mean particle diameter is $\langle d \rangle \equiv 1$. a) For the scalable potential $V_{scal}(r)$ the liquid prefers on average larger particles than the gas. b) The non-scalable potential $V_{non\text{-}scal}(r)$ causes the liquid to prefer smaller particles. Error bars represent the standard error on the mean.

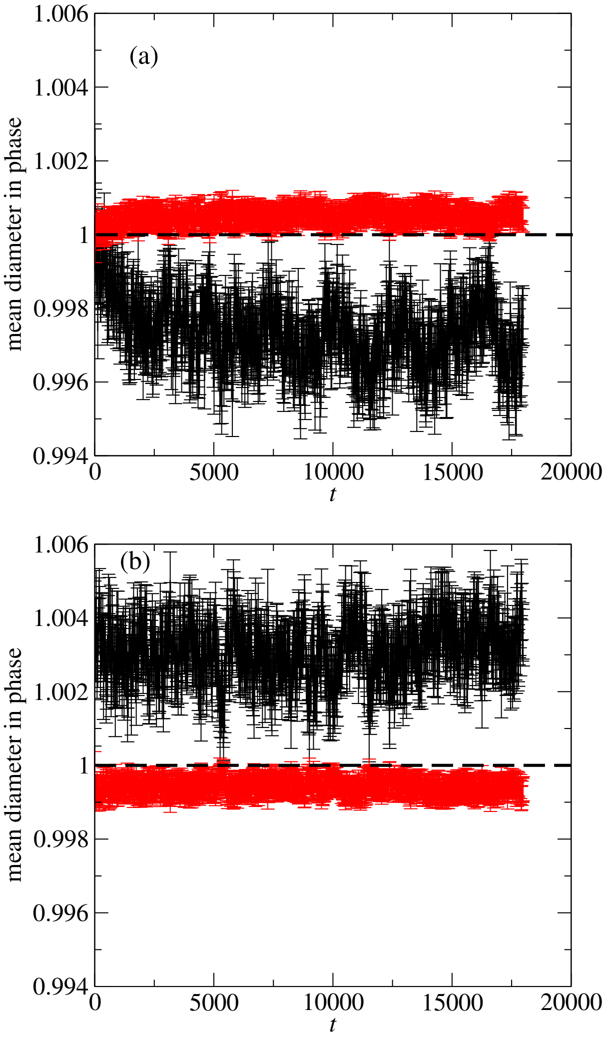

Figure 5.6: As FIG. 5.5, for 20 simulations with $N = 4000$, run for longer times. The fractionation shows no significant change beyond $t \approx 2500$, and is consistent in its magnitude and direction with the data in FIG. 5.5.

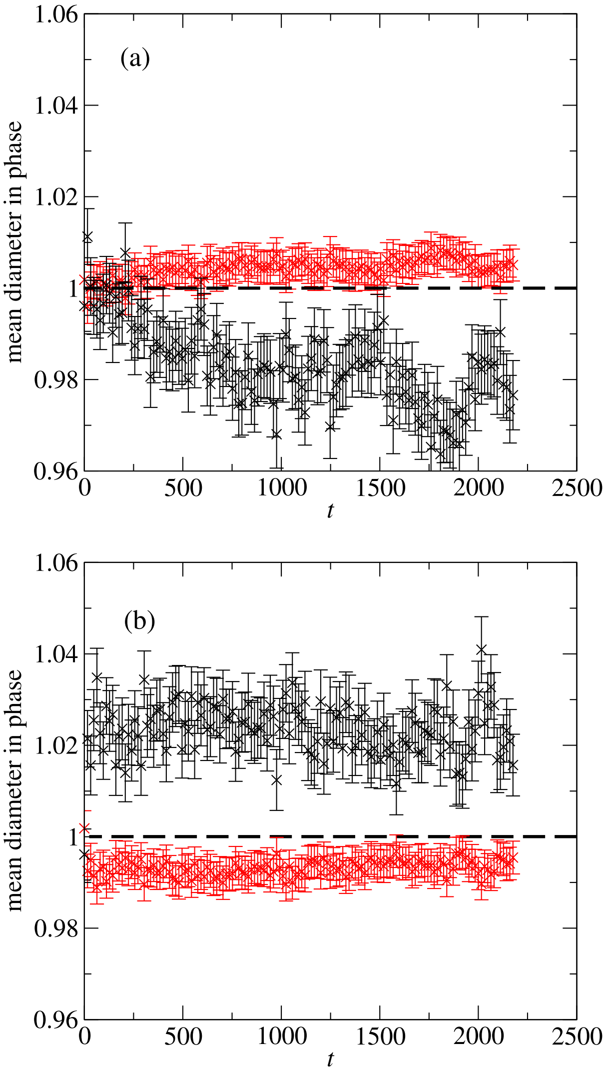

Figure 5.7: As FIG. 5.5, for higher polydispersity ($\sigma = 0.2$) simulations. The direction of fractionation is the same as in the lower-σ simulations, and the difference in mean size between the phases is increased by approximately an order of magnitude.

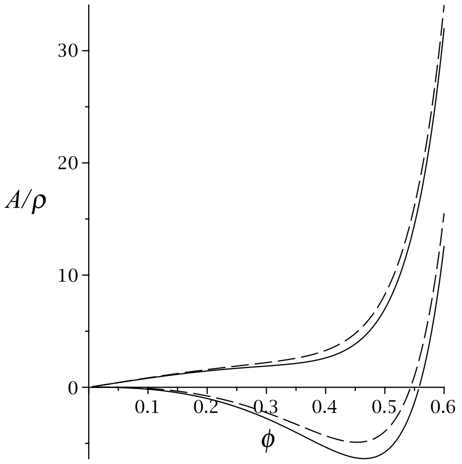

Figure 5.8: A/ρ shown as a function of volume fraction ϕ for a monodisperse reference system with $u = 1.82$ (solid lines), used to predict the fractionation of mean particle size at G-L coexistence [Evans, 2001]. Dashed lines show $u = 1.74$ for comparison. The lower two curves are for the scalable interaction potential $V_{scal}(r)$, while the upper two curves correspond to the non-scalable potential $V_{non\text{-}scal}(r)$.

course not at equilibrium, it is reasonable to invoke such a theory in order to find out what metastable state the system is relaxing towards. As discussed in Section 2.2.3, the theory of Evans becomes exact in the limit of weak polydispersity, i.e. if $\langle\epsilon^2\rangle$ is small where ϵ is a particle's deviation from the mean size $\langle d\rangle$, in units of $\langle d\rangle$. The applicability of that theory is the reason for having primarily focused on the low-σ regime in the simulations. The considerable advantage of Evans' theory is that properties of polydisperse systems may be predicted by inputting properties of a monodisperse reference system to which the polydispersity is treated as a perturbation.

For fractionation, the key thermodynamic potential is $A(\rho) = \rho d\mu^{ex}(\epsilon)/d\epsilon$ a function describing how the (excess) chemical potential μ for a particle in a monodisperse system varies with ϵ. Once $A(\rho)$ and the polydispersity σ of the parent distribution are known, the fractional difference in mean particle size between the phases is given by:

$$[\langle\epsilon\rangle]_g^l = -[A/\rho]_g^l \sigma_p^2 \ . \tag{5.3}$$

The subscripts l, g and p indicate quantities evaluated at the liquid, gas, and parent densities respectively.

To find A, an approximate free energy density f is required for the monodisperse square well fluid. I use the attractive part of the interaction potential $U(r) = -u$ as a perturbation to a hard sphere reference system with the Percus-Yevick radial distribution function $g_{HS}(r)$, giving:

$$f \approx f_{HS} + \frac{\rho^2}{2} \int 4\pi r^2 g_{HS}(r) U(r)\, dr \tag{5.4}$$

where f_{HS} is the Carnahan-Starling hard sphere free energy density [Carnahan and Starling, 1969] and the integral is over the square well range.

For $V_{scal}(r)$, finding $[A/\rho]_g^l$ is made simpler because the definition is such that all lengths, including the well range, scale with the hard core radius. In this case, it has been shown [Evans, 2001] that $A = 3(P - \rho)$ where P is the osmotic pressure calculated from the free energy density. FIG. 5.8 shows that, for the scalable potential, A/ρ decreases across the G-L coexistence region ($\phi_g \approx 0.05$ to $\phi_l \approx 0.40$) so that $[A/\rho]_g^l < 0$, correctly predicting the observed fractionation of larger particles into the liquid phase in FIG. 5.5 a). Conversely, for $V_{non\text{-}scal}(r)$, A/ρ increases monotonically so that $[A/\rho]_g^l > 0$ and the liquid contains on average *smaller* particles, in qualitative agreement with FIG. 5.5 b).

This is intuitively reasonable: from Equations 5.1 and 5.2, it is clear that the range of $V_{scal}(r)$ increases more strongly with d_{ij} than that of $V_{non\text{-}scal}(r)$. In the non-scalable case, increasing the diameter of particles in the liquid results in a greater increase in chemical potential than in the gas – the incorporation of larger particles into the liquid is unfavourable. In the scalable case, the marginally stronger dependence of the well range on d_{ij} is such that an increase in particle diameter *reduces* chemical potential, as more particles can fit into the square wells – this effect is enhanced in the liquid, compared with the gas, due to its smaller inter-particle spacing, so that the liquid now favours larger particles.

For a quantitative comparison, we can use Equation 5.3 to estimate the extent of fractionation when the metastable gas and liquid phases have 'equilibrated', i.e. found their thermodynamically preferred fractionation. This is done by evaluating $[A/\rho]_g^l$ from the curves in FIG. 5.8. Taking the coexistence volume fractions from [Liu *et al.*, 2005], $\phi_g \approx 0.05$ and $\phi_l \approx 0.40$, gives $[A_{scal}/\rho]_g^l = -5.31$ and $[A_{non\text{-}scal}/\rho]_g^l = 2.20$. So, for the $\sigma = 0.06$ system:

$$[\langle d \rangle_{scal}]_g^l = 0.019, [\langle d \rangle_{non\text{-}scal}]_g^l = -0.0079 \; . \tag{5.5}$$

Although the accuracy of these magnitudes is uncertain (given the perturbative free energy used), it appears that, while the fractionation observed on the simulated timescale is in the *direction* of equilibrium, it is substantially weaker than that required to minimise the free energy of the metastably coexisting phases. This is consistent with the hypothesis that the compositions of the phases would typically relax far more slowly than their overall densities [Warren, 1999]. Given this, it is all the more remarkable that the qualitative sensitivity of fractionation to the distinction between $V_{scal}(r)$ and $V_{non\text{-}scal}(r)$ is present even at the very early stages of phase separation.

5.4.4 Neighbour count size distributions

The question of how to define gas and liquid particles was given relatively short shrift so far. I mentioned above a possible measurement artefact whereby particles whose neighbour counts lie on the threshold between gas and liquid (as defined there) might, by virtue of their own size, exclude more or less particles from their immediate surroundings so that e.g. large particles are more likely to be identified as gaseous. This could be particularly important in the early stages where many particles occupy the interfaces between the two phases, so would be likely to have these intermediate neighbour counts.

In order to measure fractionation in more detail, I now measure the particle size distribution as a function of the number of neighbours each particle has, *over the whole system* (i.e. without discriminating yet between 'gas' and 'liquid'), in the $N = 30000$ simulations at $t \approx 1600$, and

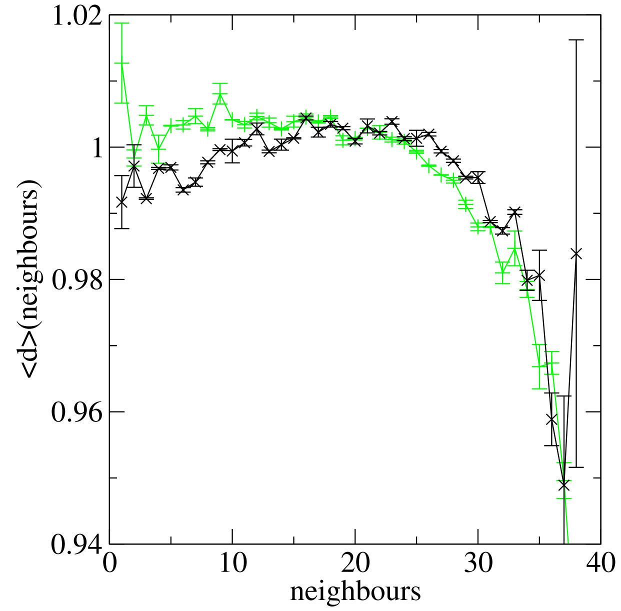

Figure 5.9: For the $N = 30000$ simulations, particle size distributions as a function of neighbour count in the scalable (black crosses) and non-scalable (green plusses) cases.

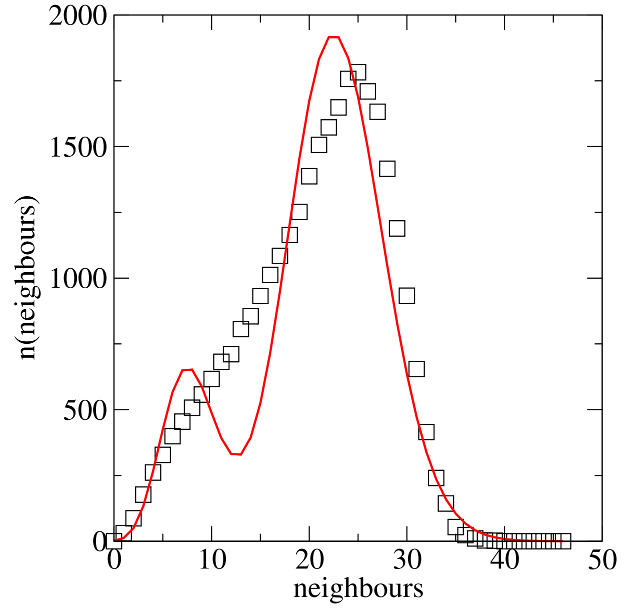

Figure 5.10: Frequency distributions of neighbour counts in the $N = 30000$ simulations (black squares). The red line is a sum of two Poisson distributions of mean values 8 and 23, intended to approximate the shape of the distribution once the phases have ripened and thus reduced the number of interfacial particles.

utilise a more complex method of estimating the fractionation which does not rely on a sharp threshold between the phases. FIG. 5.9 shows the size distributions for the scalable and non-scalable potentials. As expected, the particles with fewer neighbours (predominantly those in the gas) are of smaller size in the scalable case, and larger size in the non-scalable case. These data immediately give some confidence in the results, as the variation of $\langle d \rangle$ with neighbour count is quite smooth – there are no sudden changes at intermediate neighbour counts which might exacerbate the threshold effect described above, and the relative heights (above or below) of the two distributions are relatively constant over a large range of neighbour counts, as opposed to crossing each other in many places.

Plotting the frequency distribution of neighbour counts in FIG. 5.10 allows one to see that, indeed, the particles have not yet split into two bulk phases well-separated in terms of their mean neighbour count. The effect of the large number of interfacial particles is to smear out the peaks corresponding to the mean neighbour counts in the bulk phases. I have attempted a method of reducing the impact of these interfacial particles by 'fitting' two Poisson distributions to the data, to represent the peaks corresponding to particles which are well-defined gas or liquid. ('Fitting' is not quite the right word, as no attempt is made to match the actual, near-unimodal shape of the distribution – instead, I am attempting to 'back out' the distribution that might be expected once the phases have ripened further). This results in mean neighbour counts of 8 in the gas and 23 in the liquid. I then integrate each Poisson distribution in turn over the size distributions in FIG. 5.9 to estimate the values $\langle d \rangle_{gas}$ and $\langle d \rangle_{liquid}$ would take once the phases had ripened and reduced the number of interfacial particles. In the scalable case, this gives:

$$\langle d \rangle_{gas} = 0.9975 \ , \ \langle d \rangle_{liquid} = 1.00035 \tag{5.6}$$

and in the non-scalable case:

$$\langle d \rangle_{gas} = 1.0040 \ , \ \langle d \rangle_{liquid} = 0.9984 \ . \tag{5.7}$$

These values show the correct qualitative relationship to one another and are not far from quantitative agreement with FIG. 5.5. Although necessarily inexact, this method provides at least an estimate of the extent to which the fractionation measurements might be affected by interfacial particles. Given that the results corroborate those of the simpler technique, the fractionation observations appear to be quite robust to measurement artefacts.

5.5 Schulz distribution, skew and high polydispersity

Having studied fractionation in the near-monodisperse regime, I now move onto a more exotic particle size distribution with higher polydipersity. In this section, I make use of the Schulz distribution, for two reasons: i) it naturally truncates at zero particle size, so that no zero or negatively sized particles are created even at high σ; ii) It has a *skew*, i.e. a nonzero third central moment which, according to the perturbative theory [Evans, 2001] should result in fractionation of *variance* (σ^2), as well as $\langle d \rangle$, between the phases. The Schulz distribution occurs widely in real soft matter systems, e.g. in the polydispersity of colloids or polymers.

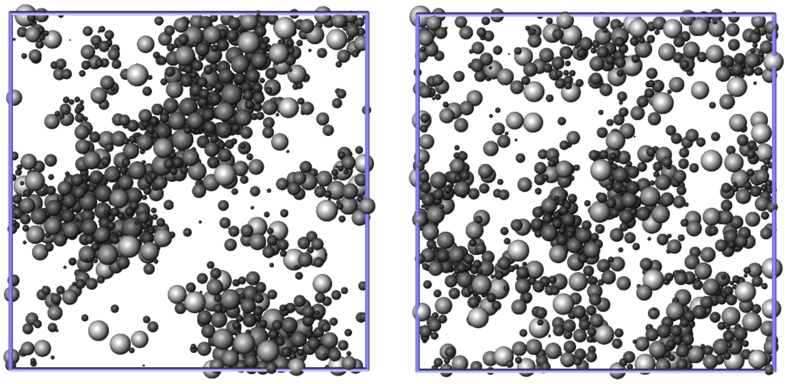

Figure 5.11: Snapshots at $t = 960$ for the $\sigma = 0.44$ Schulz distribution, in the scalable (left) and non-scalable (right) cases, for a slice 3 units thick.

Equation 5.3 can be written more generally in the form:

$$[\langle\epsilon^n\rangle]_g^l = -[A/\rho]_g^l \langle\epsilon^{n+1}\rangle_p \tag{5.8}$$

expressing one of the key results of Evans' perturbative theory: the difference in the value of the n'th central moment between two coexisting phases scales with the $n{+}1$'th central moment of the parent distribution, so that the fractionation of variance may be predicted from the *skewness* $\langle\epsilon^3\rangle$ of the parent. Importantly, this suggests that a parent distribution which is not skewed cannot be purified, at least not in the sense of reducing its polydispersity, via fractionation induced by phase separation. For the simulations in Section 5.4.2, where a skew-free distribution was used, this indeed seems to be the case – no fractionation in variance was detected. We will see in Chapter 6 that when one of the phases is a crystal, variance *can* be fractionated without a skewed parent, in disagreement with the perturbative theory. For now, I extend my study of dynamical G-L separation, using the Schulz distribution to examine the effects of skewness and increased polydispersity on the results.

I now present further simulation results, using the Schulz distribution, at polydispersities $\sigma = 0.10$, $\sigma = 0.20$, $\sigma = 0.29$, $\sigma = 0.37$ and $\sigma = 0.44$ (with $N = 10000$, and all other parameters as in Section 5.4.2). The Schulz distributions are truncated at a maximum diameter $d = 2$ in order to avoid extremely large particles (which slow down the simulation), and have a natural lower cutoff which prevents any zero-sized particles being created. At each state point three independent simulations are used. In addition to checking $\langle d\rangle_{gas}$ and $\langle d\rangle_{liquid}$ through time, I now measure the variance σ^2 of each phase to check for fractionation in this property induced by the skewed parent distribution.

5.5.1 Fractionation of mean diameter

First, I measure $\langle d\rangle$ in the phases through time, as before, to investigate its fractionation at higher polydispersities. The results are shown in FIG. 5.13. The fractionation is qualitatively similar to that observed with the Bates size distribution, with the scalable potential partitioning

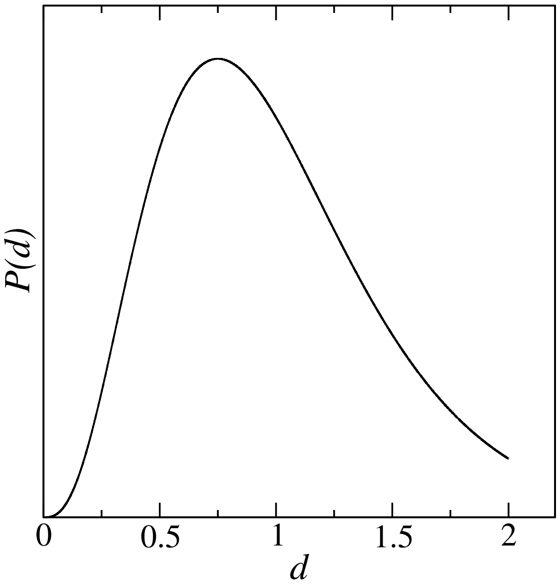

Figure 5.12: A Schulz particle size distribution of positive skew and polydispersity $\sigma = 0.44$, with an upper cutoff at $d = 2$, as used in the simulations.

smaller particles into the gas, while the non-scalable one does the opposite.

The first point to be noted is that this fractionation trend continues even up to $\sigma = 0.44$. Moreover, it can be resolved despite the simple gas-liquid identification method which, as I noted in Section 5.4.4, may introduce some measurement artefact due to the fact that smaller particles are more likely to have enough neighbours within 2 length units to be counted as liquid. As shown in FIG. 5.12, the difference between the minimum and maximum diameters in the $\sigma = 0.44$ system is actually of the same order as the neighbour check range, so it is noteworthy that the difference in fractionation between the potentials is resolvable against what is presumably rather a strong background systematic bias at these high polydispersities. On the other hand, these extremely large and small particles are relatively rare, and one must recall that the strength of the underlying fractionation signal grows with σ_p^2.

Nevertheless, it seems that some bias may be present: for the non-scalable potential, where the fractionation is in the same direction as the suggested measurement artefact (i.e. smaller particles into the gas), the fractionation measured is consistently stronger than the corresponding scalable case, whereas the equilibrium theory (Equation 5.5 and surrounding text) suggests fractionation should be stronger in the scalable case. For this reason, we should refrain from placing much stall in the different magnitudes of fractionation between the two potentials, besides noting that both appear to be weaker than equilibrium.

5.5.2 Fractionation of variance

I have noted already that in the simulations using the (skew-free) Bates distribution (Section 5.4.2), no fractionation in the variance σ^2 of the phases was evident, as expected from

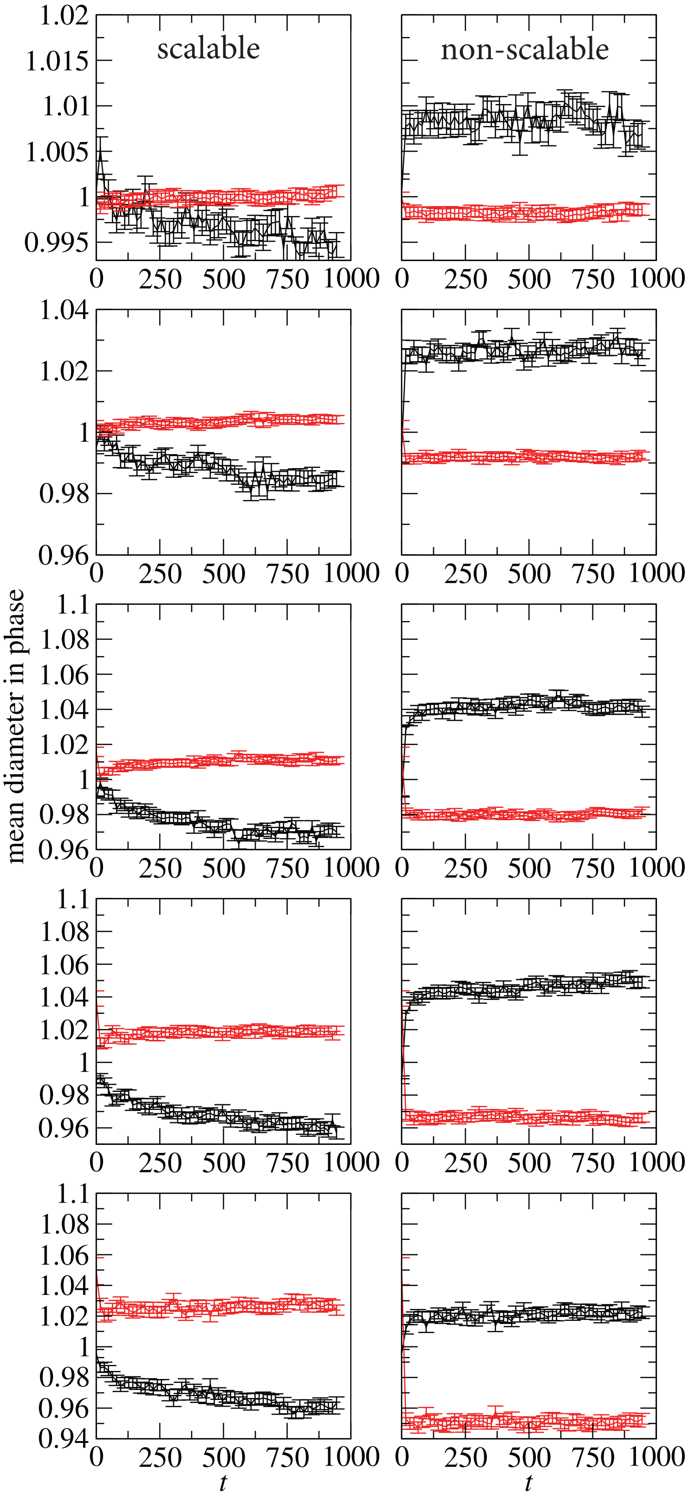

Figure 5.13: As FIG. 5.5, for simulations of $\sigma = 0.10, 0.20, 0.29, 0.37, 0.44$ (increasing down the columns), with scalable potential in the left column, non-scalable the right. Results at each state point are averaged over three independent realisations with $N = 10000$. Note the different y axis on each row.

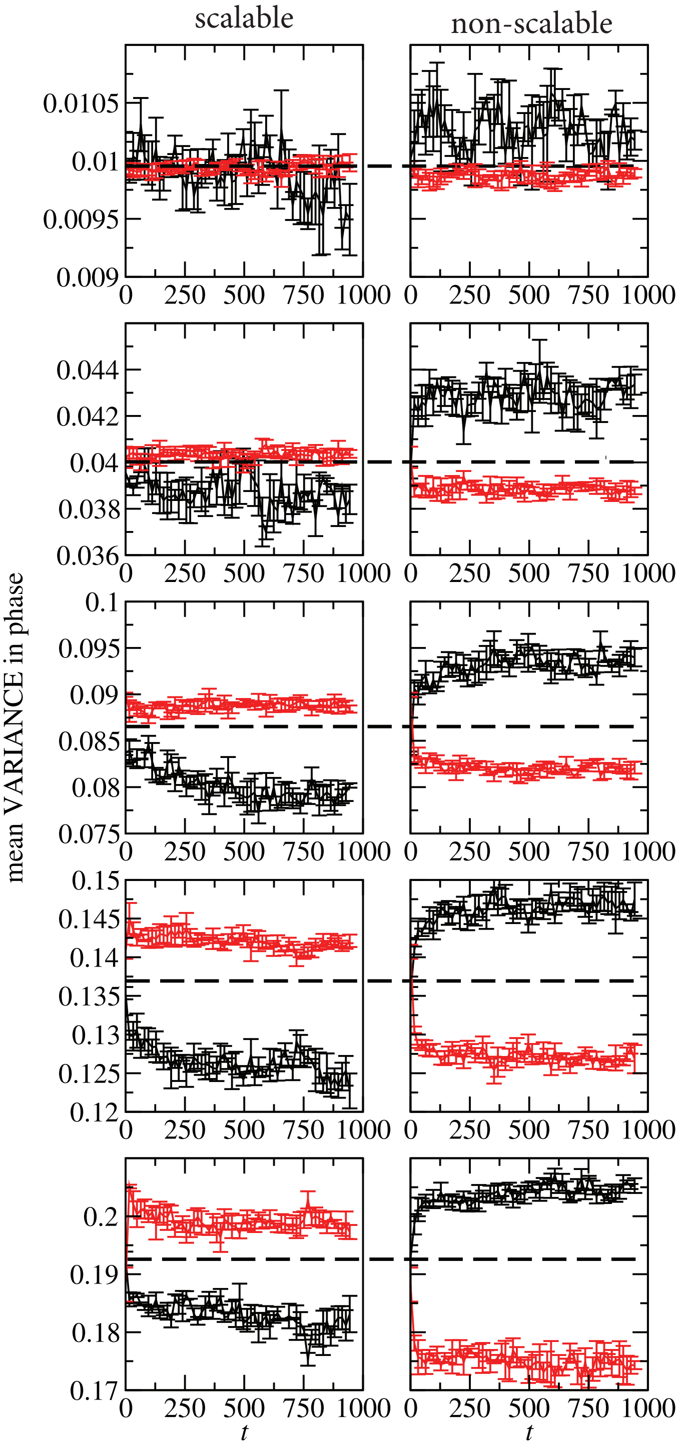

Figure 5.14: As FIG. 5.13, but showing the mean variance σ^2 in each phase over three independent simulations. Again, the left column is for the scalable potential, the right for the non-scalable. In each row, the parental variance is shown as a black dashed line.

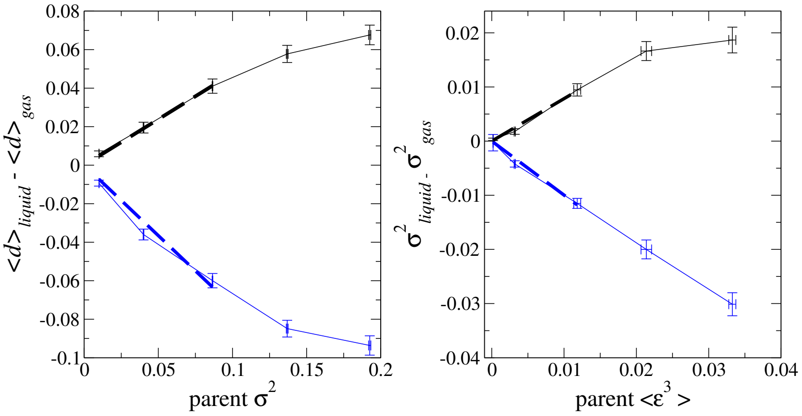

Figure 5.15: Summary of the results in FIGs. 5.13 and 5.14, where fractionation in diameter (left pane) and variance (right) have been averaged at each state point over times beyond $t = 750$. Black shows results for the scalable potential, blue the non-scalable one. Heavy dashed lines are fits to the first three data points, with gradients $0.48, -0.73$ (left) and $0.78, -1.00$ (right).

Equation 5.8. The Schulz distribution exhibits positive skew, i.e. a long tail towards high diameter, which should result in fractionation of σ^2 between the phases. This is observed in FIG. 5.14, and becomes stronger as the polydispersity and skew of the parent increases. The direction is as predicted by the theory, and is intuitively correct: because the high diameter end of the distribution has the longer tail, the phase which contains larger particles will also contain a wider spread of particle sizes.

5.5.3 Collected Schulz results

In FIG. 5.15, the fractionation of diameter mean and variance are plotted as functions of parental variance and skew respectively, so that the gradient corresponds to $-[A/\rho]_g^l$ in Equation 5.8. The dependence is approximately linear, as predicted by the theory, although it seems to 'saturate' at very high polydispersity, suggesting deviation from the perturbative theory and/or increased measurement effects. For this reason, I have used only the first 3 points to extract a gradient. As we saw already, the fractionation of both properties is weaker than equilibrium, so the gradients are shallower than predicted at equilibrium.

Also, for both inter-particle potentials, the measured degree of fractionation in *variance* is stronger, i.e. shows a steeper gradient, than the fractionation of mean diameter on the simulated timescale (note that the constant $-[A/\rho]_g^l$ is a constant, so both gradients should be equal at equilibrium). Although the possibility of measurement artefacts must be borne in mind here as well, this could suggest that the separating phases fractionate σ^2 more quickly than they do $\langle d \rangle$. It is unclear why this should be the case, but equally there is no particular reason, in a nonequilibrium system that is part-way through phase separating, why different moments should not relax at different rates. As our understanding of polydisperse phase transition kinetics becomes more detailed, this kind of comparison between moment relaxation rates may help to elucidate the intricacies of phase separation kinetics with a manageable number of variables, in the same way that consideration of a finite number of moment *densities* clarifies the equilibrium picture. An approach to polydisperse phase separation dynamics based on the relative relaxation rates of different moments has been proposed in Ref. [Warren, 1999]. Dynamical simulations in which the different moments can be tracked through time, as here, may therefore form the basis for tests of theoretical predictions along those lines, although I would suggest that better statistics and perhaps a more sophisticated method of distinguishing one phase from the other would be required in that case.

5.5.4 Summary

In this section, I have extended my study of gas-liquid fractionation to higher polydispersity and to nonzero skew, using the Schulz particle size distribution. I observed fractionation in both mean diameter and, due to the skew, *variance* between the phases, as predicted by equilibrium theory. Given the perturbative nature of the theory it is impressive that its prediction of linear dependence of fractionation strength on the parent variance or skew holds quite well, even up to what are rather wide distributions by colloidal standards. Again, the fractionation on the simulated timescale appears to be substantially weaker than at equilibrium.

5.6 Conclusions

In this chapter I have performed dynamical simulations of gas-liquid separation in a model polydisperse fluid, in order to study how fractionation is enacted during the early stages of phase separation, as opposed to at equilibrium. I first used the evolution of the static structure factor $S(q)$ to confirm that the phase separation was spinodal in nature, i.e. characterised by a sudden, all-over decomposition into small gas and liquid regions which then coarsen through time. Hydrodynamic interactions (HI) are important for spinodal decomposition in real systems, but were neglected in the simulations in order to reduce computational expense. I argued that, especially since HI may be screened to a greater or lesser extent, the question of how fractionation takes place in their absence is a pertinent one. Further work will be required to incorporate HI into simulations still large enough to gather the required statistics for fractionation.

I then measured the mean diameter in the separating phases and confirmed the presence of fractionation. Whilst dynamical simulations of crystallisation in bidisperse mixtures have previously found evidence of fractionation [Williams *et al.*, 2008] these are, to my knowledge, the first such measurements in the early-stage phase separation of a truly polydisperse model colloidal fluid. The equilibrium theory used here has been tested in experimental scenarios [Erné *et al.*, 2005; Evans and Fairhurst, 2004; Evans *et al.*, 1998] in which fractionation measurements are typically taken after a long period of equilibration. For a typical colloidal particle of radius ≈ 100nm in water, the diffusion time is $t_d \sim 10^{-3}$s in which case the simulation time used here represents the first few seconds of phase separation. Experimentally, measurements in these early stages are difficult or impossible in the absence of large single-phase regions or when the system is evolving quickly.

The data presented show that the beginnings of fractionation can take effect from the outset of spinodal decomposition, but with a smaller magnitude than predicted by equilibrium calculations. It must be noted that those calculations used a rather approximate perturbative expression for the free energy, in order to simplify the subsequent differentiation, and that their precise results may be subject to quantitative change if a more sophisticated expression were used. Nevertheless, the theory used was appropriate for our purposes, correctly predicting the dependence of fractionation direction on inter-particle potential that is observed in the simulations. That dependence was switched qualitatively by the choice between two potentials which differ only in the precise form of their dependence on particle size, and are *identical* in the monodisperse case, exemplifying the subtle interplay between polydispersity and other parameters. This finding in particular starkly illustrates the care that must be taken as the treatment of polydisperse systems becomes more sophisticated: under a 'near-monodisperse' approximation, the choice between the two inter-particle potentials studied might have been expected to have a negligible effect on the results and therefore be made simply to maximise theoretical expediency. On the contrary, I have shown that this choice can dramatically affect (and in this case, reverse) the fractionation observed, a fact that must be taken into account if precise knowledge and control of the composition and dependent properties of the two phases is desired.

Having measured and explained the fractionation observed in the near-monodisperse regime (and with a symmetrical, i.e. non-skewed particle size distribution), I moved to the more complex Schulz distribution, which has positive skew, and to higher polydispersities up to $\sigma = 0.44$. Surprisingly, the fractionation trend observed at lower polydispersities remains, even though one would expect measurement artefacts related to the intricacies of actually defining 'gas' and 'liquid' particles to increase also.

Because of the skewed distribution, and in agreement with the perturbative theory used, I also measured fractionation in diameter variance as well as its mean, which was not observed for the symmetrical size distribution – the ability to compare a truly symmetrical distribution with an (arguably more physical) skewed one is a significant advantage of a simulation approach to the problem. As predicted by theory, the dependence of mean diameter fractionation on parental variance, and of variance fractionation on parental skew, is approximately linear, even well away from the monodisperse limit in which the theory becomes exact. I mentioned that the different moments of the size distribution seem to relax at different rates. Such measurements may provide a manageable way to describe the complexities of polydisperse phase transition kinetics, for comparison with theories along the lines suggested by Warren [Warren, 1999], and in Chapter 10.

These results provide insight into how fractionation might take place during the early stages of phase separation, in a system which is approaching equilibrium but is very far from it, and also reveal a striking illustration of the complex way in which fractionation can depend on the details of the inter-particle potential. The fractionation of different moments of the size distribution confirms the qualitative predictions of Evans' theory up to high polydispersity. The fact that dynamical simulation can resolve such phenomena in the early stages of phase separation highlights the role it could play in bridging the gap between equilibrium theories of colloidal polydispersity and the behaviour of real, thus necessarily dynamical, systems.

Chapter 6

Results III – The effects of polydispersity and metastability on crystal growth

In this chapter I move from two phases to three, combining the metastable gas-liquid (G-L) separation introduced in Chapter 5 with crystal growth. I thereby explore how crystal growth kinetics can be influenced by metastability and polydispersity and find a complex interplay between the two, showing that these two factors, ubiquitous in soft matter, must be considered in tandem if their dynamical effects are to be understood. For a certain kind of free energy landscape, a gas layer coats the growing crystal. I find that, in the monodisperse case, this slows growth, but in the polydisperse case enhances it, relative to corresponding state points outside the G-L binodal. I relate the phase transition kinetics to theory, and make novel measurements of fractionation and local size ordering in crystal formation, proposing that these are dynamically facilitated by the gas layer.

6.1 Introduction

As I have already discussed, the physics of soft condensed matter has its own subtleties, and also inherits more general features common to simpler analogous systems. For example, a collection of mesoscopic particles will almost certainly be polydisperse, i.e. will exhibit a distribution in properties such as size, charge, shape or chemical makeup. While a great deal of work has focused on the equilibrium effects of polydispersity, its effects on phase ordering kinetics as actually enacted in real systems are unclear [Liddle *et al.*, 2011; Martin *et al.*, 2003, 2005; Schöpe *et al.*, 2007; Williams *et al.*, 2008; Zaccarelli *et al.*, 2009], and quantitative theoretical work thereon is almost non-existent, exceptions being [Evans and Holmes, 2001; Warren, 1999] and, in this thesis, Chapter 10.

Secondly, and in common with simpler systems, soft matter's path towards equilibrium may be influenced by the presence of *metastable* states. These are the result of local minima in the free energy landscape which, although they do not fully minimise the free energy of the system (and therefore do not appear in the equilibrium phase diagram), may be encountered as an intermediate stage. They may be long-lived, especially if the system must overcome some free energy barrier in order to reach its global minimum. A prime example is the gas-liquid (G-L) separation of square well particles which, as studied in Chapter 5, is metastable for sufficiently short attraction ranges, so may subsequently progress to equilibrium crystal-fluid coexistence. The influence of metastability on phase ordering is a problem of general importance, and within soft matter has particular significance in e.g. colloid-polymer mixtures and protein crystallisation [Anderson and Lekkerkerker, 2002; Dixit and Zukoski, 2000; Haas and Drenth, 1999, 2000].

In this chapter, I dynamically simulate crystal growth in a model colloidal system exhibiting metastable G-L coexistence. By varying the interaction strength and size polydispersity, I study the effects of metastability and polydispersity on the crystal growth process. For the parameters used, I find that the metastable G-L separation results in a gaseous layer coating the growing crystal, an observation which I relate to theoretical free energy curves and existing experimental work. When polydispersity is present, this layer can act to speed up crystal growth, in contrast to the monodisperse case in which crystal growth is *slowed* by the presence of the gas – polydispersity *qualitatively switches* the way in which G-L separation affects the crystal's growth. Interpreted as a dynamical phenomenon, this suggests the existence of a local fractionation process at the crystal interface which is facilitated by the gas layer.

I then test this hypothesis and find that crystal formation is associated with a reduction in polydispersity compared to the parent distribution. Then, I define and measure a local size correlation function, gaining new insight into the structuring of polydisperse phases, and finding a qualitative difference between the crystal and fluid which is indicative of a local rearrangement process in addition to overall fractionation. The results shed light on the influence of polydispersity and metastability on the crystal growth process, demonstrating the complex way in which these factors can interact, and elucidate the details of crystal formation in polydisperse systems generally.

6.2 Simulation

6.2.1 Kinetic Monte Carlo algorithm

The simulation is built on the Kinetic Monte Carlo algorithm and crystal templating procedure introduced in Section 3.2. The polydisperse particle size distributions are taken from a Bates (pseudo-Gaussian) distribution, as described in Chapter 3.

6.2.2 Parameters, crystal template

The system studied consists of $N = 5000$ spherical particles with diameters d_n drawn from a Bates distribution of polydispersity $\sigma = 0$ ('monodisperse'), $\sigma = 0.03$ or $\sigma = 0.06$. The mean hard core diameter $\langle d \rangle$ sets the length unit, and the time unit is the mean time t_d taken for a free particle of diameter $\langle d \rangle$ to diffuse a distance equal to its own diameter. The hard particle cores are surrounded by pairwise additive square wells of mean range $\lambda = 1.15$ and depth u. The interaction range between particles i and j depends on the individual particle diameters in a scalable fashion (see Chapter 5) [Evans and Holmes, 2001; Williamson and Evans, 2012]. The step parameter $\Delta = 0.02$ – I have checked for finite size effects by reducing its value and seeing that the crystallinity measurements through time in Section 6.4.4 are unchanged within error.

The parent volume fraction (hereafter also referred to as 'concentration') is $\phi_p = 0.34$. In the monodisperse limit, the values of ϕ_p and λ are such that the equilibrium state is a coexistence of crystal and gas [Liu *et al.*, 2005]. However, by choosing u (and therefore the effective temperature $T_{\mathit{eff}} = 1/u$) appropriately, the system's state point can be made to lie either inside or outside the metastable G-L binodal, as shown in FIG. 6.1. Given that the critical effective temperature corresponds to $u = 1.72$, I use a well depth of $u = 1.82$ inside the G-L binodal (corresponding to point 1 in FIG. 6.1) and $u = 1.54$ outside (corresponding to point 2), where the energy unit is k_BT. In soft matter, one is often free to vary the interaction strength in this way, for instance by adding more or less polymer into a colloid-polymer mixture.

Since the aim is to study the process of crystal *growth*, not nucleation, I introduce crystallisation artificially by placing a 10×10 template of immobilised particles, arranged in the (100) face of the FCC lattice, at one end of the simulation cell – the procedure is described in Section 3.2. The simulation cell is long in the x axis ($L_x \approx 66$), and relatively short in the (periodic) y and z dimensions ($L_y = L_z \approx 11$), so that the growth of the templated crystal along the x axis can be studied for a long period of time. The template's lattice parameter is set to match the volume fraction of the equilibrium crystal [Liu *et al.*, 2005]. Increasing the y and z dimensions of the lattice had no measurable effect on the results presented. Crystalline particles are identified by constructing bond order vectors $\mathbf{q}_6$ from the spherical harmonics Y_{6m} [Williams *et al.*, 2008; ten Wolde *et al.*, 1996] – see Section 3.2 for details. After the system is initialised as an amorphous hard sphere fluid, the square well attraction is turned on and the template positioned, defining $t = 0$.

6.3 Theory

6.3.1 Free energy landscape

Although this work is concerned with the dynamics of phase transitions, it is possible to gain a surprising level of insight into the expected behaviour from thermodynamic considerations alone [Poon *et al.*, 2000, 1999; Renth *et al.*, 2001], by examining the free energy landscape of the system. To that end, I now calculate the approximate crystal branch of the free energy for the system, in the monodisperse case, by perturbation to a hard sphere reference system. Combined with the fluid branch, which was used in Chapter 5, this specifies the complete free energy landscape.

Again, the free energy density f is given by:

$$f \approx f_{HS} + \frac{\rho^2}{2} \int 4\pi r^2 g_{HS}(r) U(r)\, dr \tag{6.1}$$

where $U(r) = -u$ and ρ is the number density. The integral is over the square well range. For the hard sphere contribution f_{HS}, I use the Carnahan-Starling free energy density [Carnahan and Starling, 1969] for the fluid, and Hall's expression for the crystal [Hall, 1972]. For the hard sphere radial distribution function $g_{HS}(r)$, I use the Percus-Yevick expression in the fluid [Stell, 1963] and that of Choi et. al. for the crystal [Choi *et al.*, 1991].

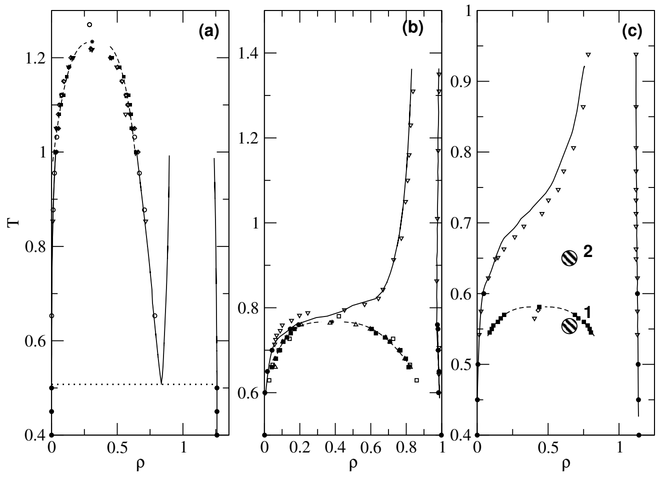

Figure 6.1: Phase diagrams reproduced with permission from Ref. [Liu *et al.*, 2005] for (a) $\lambda = 1.5$, (b) $\lambda = 1.25$ and (c) $\lambda = 1.15$ in the density-temperature plane. With my choice of units, $T = T_{eff} \equiv 1/u$ and $\rho = 6\phi/\pi$. As λ decreases, the gas-liquid coexistence becomes metastable with respect to crystal-gas separation. The relevant diagram for the present work is (c). The points corresponding to the ϕ, u coordinates used in my study are indicated with hashed circles: point 1 is 'inside the G-L binodal' or 'gas-mediated' and point 2 is 'outside the binodal' or 'gas-free'.

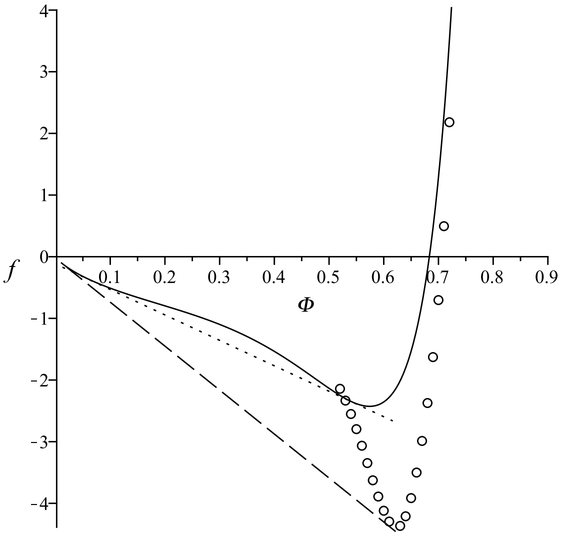

Figure 6.2: Free energy density f as a function of volume fraction ϕ for a monodisperse system with $u = 1.82$ and $\lambda = 1.15$. The solid curve shows the fluid branch of the free energy while the circles show the crystal branch. The dotted line indicates the common tangent for metastable G-L coexistence, and the dashed line shows the tangent for the equilibrium crystal-gas (C-G) coexistence. There is no common tangent between the liquid and crystal phases.

FIG. 6.2 shows the resulting free energy densities plotted as functions of volume fraction for the monodisperse $u = 1.82$ case (inside the G-L binodal for the value of ϕ_p used). Allowed coexistences are given by the common tangent construction introduced in Section 2.1.3. Although quantitative agreement with the equilibrium simulation data of Ref. [Liu *et al.*, 2005] (in terms of the coexistence volume fractions of the gas and liquid phases) is relatively poor, some important qualitative features are present. The common tangent linking the minimum of the crystal branch with the gas is lower than that linking the liquid and gas – the crystal-gas (C-G) coexistence therefore has a lower overall free energy and is hence the equilibrium state, while the G-L coexistence is metastable. Furthermore, there is *no common tangent* between the crystal and liquid, which means that on the basis of free energy considerations, the crystal cannot locally coexist with the metastable liquid.

The lack of a crystal-liquid common tangent implies a growth scenario like that proposed for the experimental colloid-polymer mixture observations in Refs. [Evans *et al.*, 2001; Renth *et al.*, 2001], in which growing crystallites are coated by a layer of colloidal gas because of their inability to coexist with the metastable liquid. (There, a crystalline region was observed to shrink in volume after formation, as the crystallites' gaseous coatings separated off due to their greater buoyancy.) This 'boiled-egg crystal' should deplete the surrounding bulk fluid until the required (equilibrium) C-G coexistence is achieved overall. The simulation setup used here allows detailed observation of the effects of a gas layer on the growth of a single crystal, in a system for which a quantitative calculation of the corresponding free energy landscape is available, allowing new insight into some of the physics described in Refs. [Evans *et al.*, 2001; Renth *et al.*, 2001].

6.3.2 Split interfaces

The expected growth scenario of the 'boiled-egg crystal' described in Section 6.3.1 is that of a 'split interface'. That is, the crystal should form an interface with a gaseous region, which in turn forms an interface with the metastable liquid, since the crystal and metastable liquid cannot coexist on the free energy grounds discussed above. Such interfaces and a method for predicting their growth rates are described in Refs. [Evans and Poon, 1997; Evans *et al.*, 1997], wherein the example used was of a crystal-liquid-gas interface as opposed to the crystal-gas-liquid interface we expect here. I now make use of that theory to predict the evolution of the split interface that should be formed when the system is inside the G-L binodal.

As shown in FIG. 6.3, I assume the growth scenario to be such that the phases on each side of the interfaces are at their 'correct' densities according to the equilibrium phase diagram (FIG. 5.1), the densities $\rho_A, \rho_B, \rho_C, \rho_D$ corresponding respectively to: the metastable liquid, the metastable gas, the stable gas, and the stable crystal. The assumption is therefore one of local equilibrium at the interfaces.

Let $x_1(t)$ be the position of the gas-liquid interface, and $x_2(t)$ that of the crystal-gas interface. The respective fluxes on to and away from the interfaces are j_1, j_2, j_3 and j_4, as shown in FIG. 6.3. I assume $j_1 = 0$, i.e. that the liquid exists uniformly at its metastable density, and $j_4 = 0$ within the crystal region. The remaining fluxes can then be related to the interface speeds and the densities ρ_i:

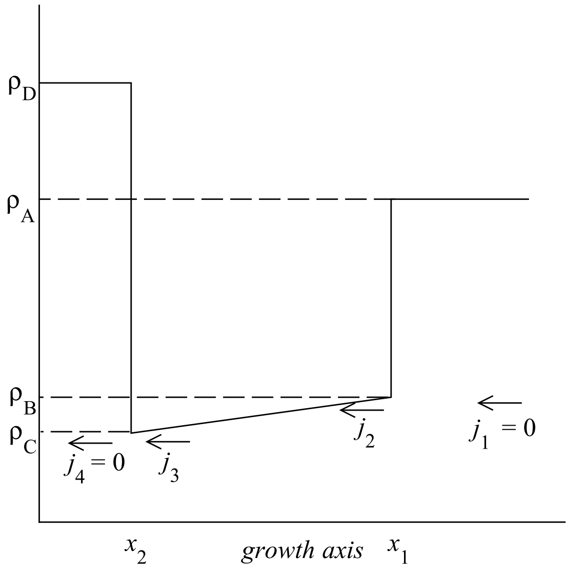

Figure 6.3: Schematic density profile of the expected crystal-gas-liquid split interface inside the G-L binodal, where the crystal cannot coexist with the metastable liquid. The phases from left to right are the crystal, gas, and metastable liquid. The gradient in the gas region is drawn uniform, as is assumed in the calculations.

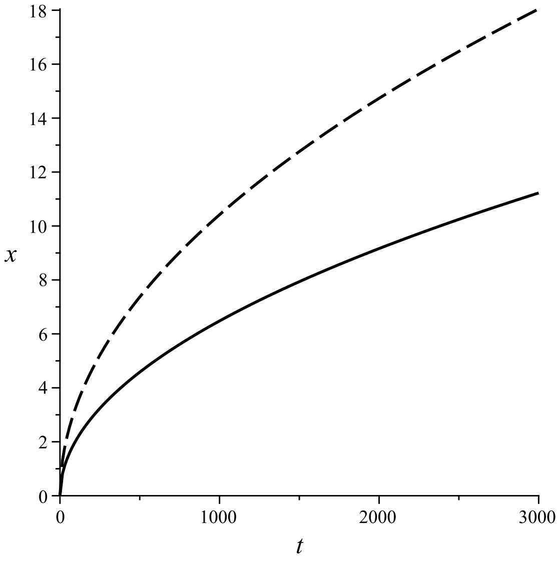

Figure 6.4: The predicted evolution of the crystal-gas (solid line) and gas-liquid (dashed line) interfaces for a split crystal-gas-liquid interface when the system is inside the G-L binodal ($u = 1.82$, $\lambda = 1.15$).

$$j_2 = \dot{x}_1(t)(\rho_A - \rho_B) \tag{6.2}$$

$$j_3 = \dot{x}_2(t)(\rho_D - \rho_C) \tag{6.3}$$

I now make the simplifying approximation that the gas region supports a uniform gradient between the densities ρ_B and ρ_C, as is shown in FIG. 6.3, and further that the diffusion constant therein is equal to the ideal Stokes-Einstein diffusion constant D_0 due to the low density of the gas. Therefore:

$$j_2 \approx j_3 \approx D_0 \frac{\rho_B - \rho_C}{\Delta x} \tag{6.4}$$

where $\Delta x = x_1(t) - x_2(t)$. Substituting into Equations 6.2 and 6.3 and separating variables, with the necessary constants of integration given by the initial conditions $x_1(t) = x_2(t) = 0$, I find expressions for the interface positions:

$$x_1(t) = \sqrt{2D_0 \frac{\beta^2}{\gamma} t} + \sqrt{2D_0 \gamma t} \tag{6.5}$$

$$x_2(t) = \sqrt{2D_0 \frac{\beta^2}{\gamma} t} \tag{6.6}$$

where I have defined $\beta \equiv (\rho_B - \rho_C)/(\rho_D - \rho_C)$ and $\gamma \equiv (\rho_B - \rho_C)/(\rho_A - \rho_B) - (\rho_B - \rho_C)/(\rho_D - \rho_C)$.

To make contact with the simulations, I simply read off the values of ρ_i from FIG. 5.1 at the appropriate effective temperature $T_{\mathit{eff}} = 1/1.82 \approx 0.55$ for the systems inside the G-L binodal, and use the Stokes-Einstein diffusion constant defined by my choice of time unit, $D_0 = 1/6$. The resulting interface positions through time are shown in FIG. 6.4, and in the next section are compared with simulation data. The characteristic $t^{1/2}$ diffusive growth is apparent, with the G-L interface leading the C-G interface as expected.

I note at this stage that the crystal-gas-liquid split interface scenario is very closely analogous to the crystal-liquid-gas interfaces described in Refs. [Evans and Poon, 1997; Evans *et al.*, 1997]. The only difference is that in the present work, the 'pivot' phase which coexists locally with both others (and therefore coats the crystal) is the gas, i.e. *is the equilibrium phase*, whereas in the aforementioned work it is the metastable, i.e. nonequilibrium liquid phase.

6.4 Results and discussion

In this section, I qualitatively and quantitatively examine the simulation results in terms of their time-dependent concentration profiles and in terms of their crystal growth rates, elucidating and discussing the separate and combined effects of metastability and polydispersity on the growth process. Then, in Section 6.5, I outline and test a possible explanation for my findings in terms of fractionation and rearrangement at the crystal interface.

I present results for polydispersities of $\sigma = 0$ (monodisperse), $\sigma = 0.03$ and $\sigma = 0.06$. The ϕ and u coordinates used are marked as points '1' and '2' on the monodisperse phase diagram, FIG. 6.1. I use square well depths of $u = 1.82$ (point 1, inside the G-L binodal, referred to as

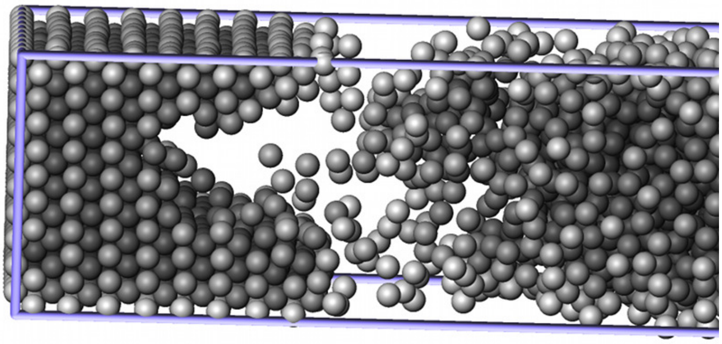

Figure 6.5: Snapshot of the crystal interface for one particular monodisperse gas-mediated trajectory in which the crystal happened to exhibit dendritic growth around an impinging gas bubble. Such dendritic growth was not generally observed.

'gas-mediated') or $u = 1.54$ (point 2, outside the G-L binodal, referred to as 'gas-free'). The combinations of u and σ give 6 state points in total. At each state point, 6 independent initial configurations are used.

6.4.1 Concentration profiles – monodisperse

The clearest way to visualise crystal growth over the whole simulation time is by plotting the time-dependent concentration profile. To achieve this, I plot the 'area fraction' ϕ_{area} of particles intersecting planes at a given x coordinate, corresponding to a local volume fraction in an infinitely narrow interval δx. Since the simulation geometry is such that the crystal interface advances along the x axis, these plots are a powerful way of observing the time evolution of concentration differences along the direction of growth while averaging over the axes perpendicular to it. Further details of the algorithm and a code sample are given in Section A.1.4.

In this section, I compare representative concentration profiles in the monodisperse and $\sigma = 0.06$ cases. Except where noted, the phenomena described occurred similarly in all trajectories of the state point in question.

Let us consider the monodisperse case first. FIGs. 6.6 and 6.7 show example trajectories for simulations at the gas-free and gas-mediated state points respectively. In the gas-free case, the growth scenario is, as expected, relatively simple. The bulk fluid remains homogeneous, since we are above the critical T_{eff} for G-L separation. The crystal template causes the growth of a single crystalline region with an average volume fraction of $\phi \approx 0.6$, which is the expected equilibrium volume fraction for the crystal [Liu *et al.*, 2005]. The crystal, being at a higher volume fraction than the bulk fluid, depletes its surroundings of particles, resulting in a concentration gradient between the local fluid and the bulk; it is this gradient which transports particles toward the crystal so that it can continue to grow.

Inside the G-L binodal (FIG. 6.7), the scenario is quite different. Firstly, there is clearly some G-L separation taking place in the bulk. Focusing next on the crystal template at $x = 0$, we can see that a region of low density gas forms in front of the crystal almost immediately, shielding it

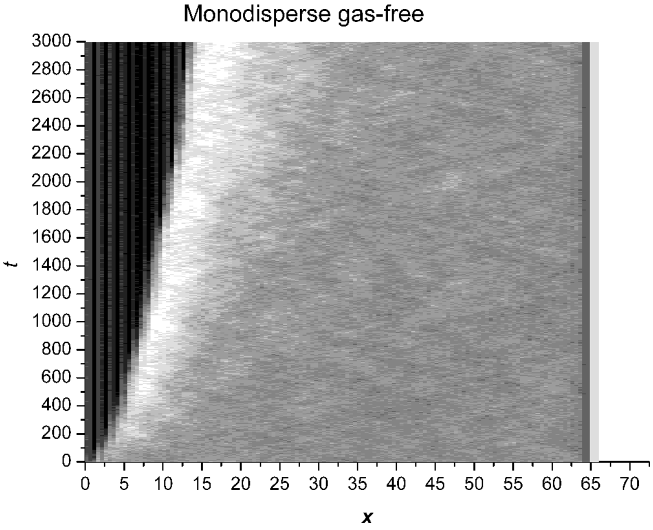

Figure 6.6: Time-dependent concentration profile along the x axis for one of the monodisperse gas-free simulations. The greyscale indicates the local volume fraction, ranging from $\phi < 0.06$ (white) to $\phi > 0.676$ (black). As the simulation progresses, the crystal advances along the x axis, depleting the fluid in front of the interface.

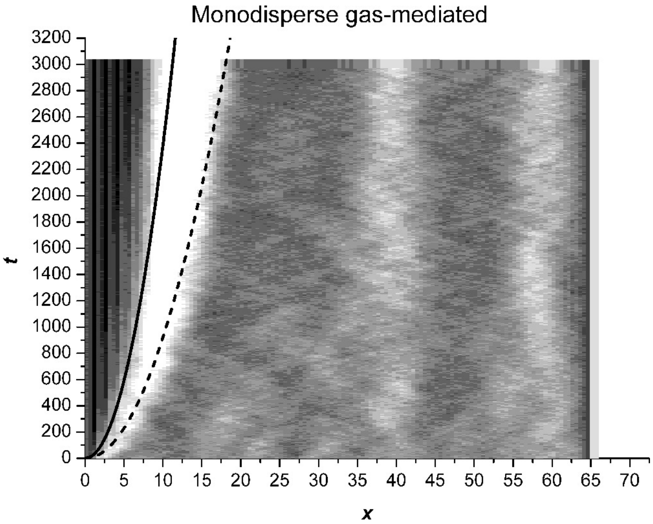

Figure 6.7: As FIG. 6.6, for a monodisperse gas-mediated simulation. Gas-liquid separation takes place in the bulk, and the advancing crystal is coated by a distinct gas layer which advances ahead of it. The theoretical crystal-gas and gas-liquid interface curves from FIG. 6.4 are rotated and superimposed, showing approximate quantitative agreement with the data.

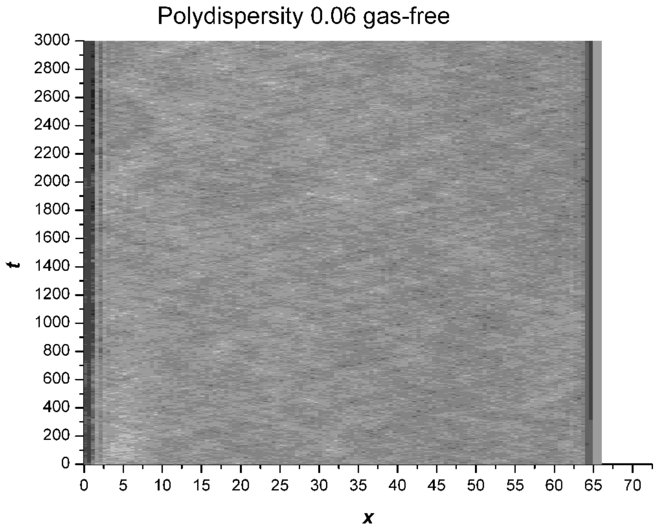

Figure 6.8: As FIG. 6.6, for a gas-free system of polydispersity $\sigma = 0.06$. As in the monodisperse case, the bulk fluid remains homogeneous. However, the crystal template does not induce any growth on the timescale simulated.

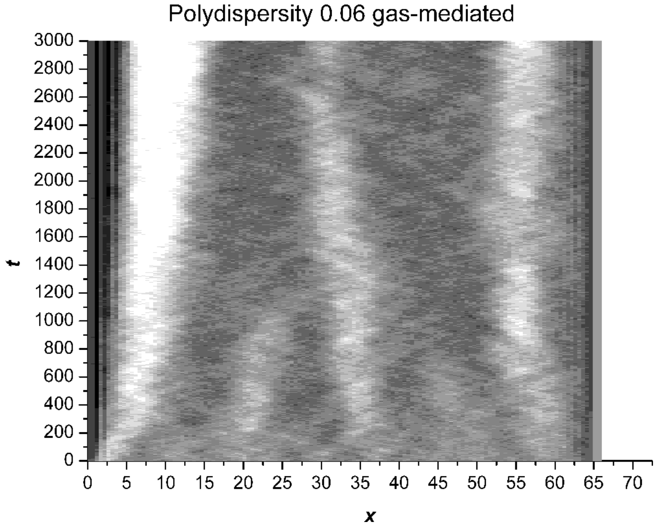

Figure 6.9: As FIG. 6.6, for a gas-mediated system of polydispersity $\sigma = 0.06$. The bulk fluid shows G-L separation, and the gas layer is present in front of the crystal. In contrast to FIG. 6.8, the crystal is able to grow, although more slowly than in the corresponding monodisperse case (FIG. 6.7).

from the liquid with which it cannot locally coexist according to the free energy considerations in Section 6.3.1. This is in contrast to FIG. 6.6, in which the fluid immediately next to the crystal template retains its density for quite some time until the growing crystal depletes it of particles. Note also that whereas the depleted region in FIG. 6.6 fades smoothly into the higher density bulk fluid, the gas next to the crystal in FIG. 6.7 forms a substantially sharper interface with the fluid next to it, indicating a distinct phase boundary between it and the metastable liquid. The formation of this well-defined gas layer shielding the crystal is consistent with the free energy considerations above, and with the experimental observations of [Renth *et al.*, 2001].

The speed of advance of the C-G and G-L interfaces is in approximate quantitative agreement with the theoretical prediction in Section 6.3.2 (FIG. 6.4), which has been rotated and superimposed on FIG. 6.7 for comparison. This agreement is particularly satisfying given that the input densities ρ_i required for the theoretical calculation in Section 6.3.2 were taken from the equilibrium phase diagram, so that there are no free fitting parameters. The accuracy of the prediction thus gives credence to the assumptions of local equilibration of the interfaces and of essentially ideal gas-like Stokes-Einstein diffusion in the gas region.

Dendritic growth

Finally, it is interesting to note that for some trajectories, the crystal can grow incomplete layers, as shown in FIG. 6.5. This is the result of dendritic growth of the crystal around an impinging gas bubble, and is a result of the inherently diffusion-limited growth of the crystal when the gas layer is present. It is possible for a portion of the crystal interface which by chance extends further than its surroundings to become closer to the G-L boundary and thereby increase the concentration gradient supplying it (since the same concentration *difference* is distributed over a shorter distance). That portion will then continue to grow faster than its surroundings and increase the gradient more, i.e. dendritic growth is an instability (known as the Mullins-Sekerka instability [Mullins and Sekerka, 1963]). This phenomenon is relevant particularly to colloids as their growth is often diffusion-limited [Tegze *et al.*, 2009].

Probably due to the relatively small $y - z$ cross-section, such growth was not found to be generally present in the other trajectories at this state point – the crystal typically formed complete layers covered by a layer of gas spanning the whole width of the simulation cell.

6.4.2 Concentration profiles – polydisperse

I now discuss the polydisperse $\sigma = 0.06$ case. FIG. 6.8 shows the gas-free scenario. On the simulated timescale, essentially no crystal growth is seen to take place, except perhaps for a slight 'wetting' of the template by ordered particles. As in the monodisperse case, the bulk fluid remains homogeneous, indicating that the polydispersity has not altered the location of the critical T_{eff} in such a way as might bring this state point inside the G-L binodal.

Moving to the gas-mediated case inside the binodal (FIG. 6.9), we can see that the crystal is able to grow substantially despite the polydispersity. As in the corresponding monodisperse case (FIG. 6.7), the bulk fluid separates into gas and liquid regions, and a gas layer covers the crystal as it grows.

To summarise the qualitative observations in this section: the crystal template successfully induces a physically realistic growth process, in which the crystal depletes its surroundings and

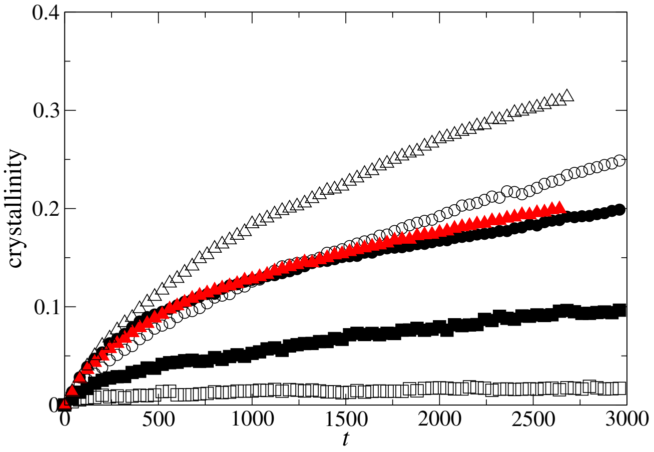

Figure 6.10: Crystallinity through time for the 6 state points studied, in each case averaged over 6 independent realisations. Filled symbols: gas-mediated; open symbols: gas-free. Triangles: monodisperse; circles: $\sigma = 0.03$; squares: $\sigma = 0.06$. The standard errors are approximately the size of the symbols. The filled triangles (monodisperse gas-mediated) are shown in red to distinguish them from the filled circles ($\sigma = 0.03$ gas-mediated).

forms ordered, high density layers. When the system is inside the G-L binodal, a gas layer forms in front of the crystal, shielding the crystal from the metastable liquid. This is exactly the scenario predicted for the system in Section 6.3.1 and proposed in Refs. [Evans *et al.*, 2001; Renth *et al.*, 2001]. The presence of polydispersity $\sigma = 0.06$ substantially slows the crystal growth overall. In the gas-free case the crystal showed barely any growth on the simulated timescale, whereas the gas-mediated case showed substantial crystal growth despite the polydispersity.

6.4.3 Monodisperse crystal growth rates – the effect of metastability

Having summarised the qualitative features of the monodisperse and $\sigma = 0.06$ simulations, I next quantitatively examine the effect of those features on the crystal growth rate. This is done by plotting the crystallinity (the number of particles flagged as crystalline as a proportion of the total N) through time. FIG. 6.10 shows the crystal growth at each of the 6 state points studied, each averaged over the 6 independent realisations of the state point.

In the monodisperse case, the effect of metastability *alone* is apparent. The gas layer in front of the crystal significantly slows crystal growth compared to the gas-free case. The likely explanation for this is that the gas layer inhibits transport of particles into the crystal, as I now discuss.

Since the crystal exists at a higher density than the rest of the system, as soon as the immediate surroundings are depleted of particles, particles must be transported toward the crystal from the bulk in order for it to grow. This transport takes place via collective diffusion down a concentration gradient between the low concentration near the crystal and the relatively higher concentration away from the crystal. The diffusion is described by Fick's law, Equation 2.8. Recall that the relevant collective diffusion coefficient, in contrast to the self diffusion coefficient, increases with ϕ [Tirado-Miranda *et al.*, 2003].

Now, consider the phase diagram for our system, as shown in FIG. 6.1 (c). In the absence of the gas layer, and assuming that the interface is locally equilibrated, a long-range, relatively steep concentration gradient can exist between $\phi \approx 0.05$ at the interface and $\phi_p = 0.34$ in the bulk. However, inside the G-L binodal, when the crystal is covered by a distinct gas layer, the relevant concentration gradient is that which exists *across the gas layer itself* – the G-L interface beyond truncates the concentration gradient. The gradient in the gas layer is between the equilibrium gas ($\phi \approx 0.01$) and the slightly higher concentration *metastable* gas at $\phi \approx 0.05$. The monodisperse results indicate that this small concentration difference, spread across the macroscopic (and constantly growing) gas layer, results in slower particle transport to the crystal than in the gas-free case.

Therefore, the low-concentration gas layer reduces the efficiency of particle transport toward the crystal because: i) the collective diffusion coefficient D increases with concentration; ii) the gas's low concentration (and macroscopic size) mean it can only support a relatively weak concentration gradient in comparison to the gas-free case, in which a strong and long-range concentration gradient can exist between the interface and the homogeneous bulk. Although, on thermodynamic grounds, the driving force for crystallisation is higher inside the G-L binodal, the growth is slowed due to the gas layer's detrimental effect on particle transport to the interface.

The effect of the gas layer in the monodisperse case is comparable to that of the liquid layer in Refs. [Evans and Poon, 1997; Evans *et al.*, 1997], which was also found to inhibit crystal growth. There, as here, the crystal is coated by a phase which advances ahead of the crystal and slows down the growth of the crystal. It is interesting that this comparison holds, given that in the present case the crystal is being coated by the gas (which, once it loses its flux, becomes an equilibrium phase), whereas Refs. [Evans and Poon, 1997; Evans *et al.*, 1997] concern a crystal coated by the *nonequilibrium* liquid. Both cases share the same salient feature, though: a 'coating' phase which results in a relatively small concentration drop being spread over a macroscopic distance (the increasing size of the coating phase), through which particles are therefore transported slowly.

6.4.4 Polydisperse crystal growth rates

I consider first the highest polydispersity studied here, $\sigma = 0.06$. As shown in FIG. 6.10, in agreement with FIG. 6.8, the gas-free state point now shows only very slight growth on the timescale simulated, whereas the gas-mediated state point shows significant crystal growth (although less than in either of the monodisperse systems). This is a *qualitative difference* compared to the monodisperse case, in which the gas layer instead strongly slows the crystal growth. That is, the presence of polydispersity qualitatively alters the effect of the gas layer on crystal growth.

In the $\sigma = 0.03$ case, the gas-mediated system is initially slightly faster, before being overtaken by the gas-free system by around $t = 1000$. This is suggestive of a crossover phenomenon between whatever factors are dominant in the fully monodisperse and $\sigma = 0.06$ cases.

6.4.5 (110) versus (100) face

I will note briefly here the results of some simulations performed to check the effect on these findings of using a different FCC face than the (100) one used so far in the crystal template. With the (110) face, the qualitative relationship between each of the monodisperse and $\sigma = 0.06$ state points was the same and, after correcting for differences in the interfacial area $L_y L_z$, the extent of the crystal through time differed by less than $\sim 10\%$ between different faces at the same state point.

6.4.6 A note on gravity

The simulations presented here implicitly assume perfect density matching between the colloid and solvent, so that there is no gravitational settling of the colloidal particles. In practice, this may not be the case. If the particles are denser than their solvent, the gas phase of G-L coexistence occupies the upper part of a test tube, the liquid the lower. This is the case in the experimental work of Liddle and Poon [Liddle *et al.*, 2011], concerning a colloid-polymer mixture of attraction range ≈ 1.06, i.e. in which G-L coexistence is metastable, as here.

Preliminary calculations suggest that that system may also exhibit a free energy landscape lacking in a crystal-liquid common tangent. One might therefore expect an interplay between the ‘boiled-egg’ gas-coated crystal growth mechanism, and a buoyancy force acting to pull the gas layers away from their crystallites, upwards into the gas region (the crystallites presumably having nucleated in the liquid region).

What effect this might have on the overall equilibration kinetics of the system is unclear at this stage. For now, I simply note that the additional effects of gravity may result in even more complex phase transition kinetics when the free energy landscape is such that crystallites must coat themselves with gas.

6.5 Crystallisation mechanism

6.5.1 Fractionation

The results presented so far show that metastable G-L separation can affect the rate of crystal growth, with a gaseous layer coating the crystal as it grows. However, the qualitative nature of that effect depends strongly upon polydispersity. In the monodisperse case, the resultant gas layer impedes crystal growth by reducing the efficiency of particle transport to the crystal. With polydispersity $\sigma = 0.06$, the metastable separation instead strongly enhances crystal growth. It is this effect of polydispersity on the crystal growth mechanism that I now discuss.

I propose a possible explanation in terms of a local fractionation process at the interface of the polydisperse crystal. In size-polydisperse systems, it is well known that phase separation is typically associated with some degree of thermodynamically-driven fractionation in which one phase e.g. contains on average larger or smaller particles than another, changes its overall polydispersity, etc. Such fractionation has been predicted theoretically [Evans, 2001; Fasolo and Sollich, 2004], and observed experimentally [Evans and Fairhurst, 2004] and in simulation [Williamson and Evans, 2012]. Of course, this kind of fractionation pertains to the *equilibrium* phase composition and may not be fully achieved in real systems on accessible timescales due to kinetic factors. This is a particularly important consideration for the crystal phase, where particles are essentially stationary *once incorporated* [Yiannourakou *et al.*, 2010] and, thereafter, the long-range particle transport required for fractionation is facilitated only by the presence of defects. In general, phase composition is expected to relax slowly in comparison with overall density so that fractionation may lag behind phase separation somewhat [Warren, 1999], leading in the crystal phase to the 'freezing-in' of a nonequilibrium composition [Evans and Holmes, 2001]. (Note, however, that the work in Chapter 5 has shown that *some* fractionation is possible in the very early stages of gas-liquid separation [Williamson and Evans, 2012].)

Nevertheless, one would expect that fractionation *at the interface* of the growing crystal, where particles are still mobile and may be easily exchanged with the fluid, is quite feasible. Indeed, the slowing effect of polydispersity on crystallisation is generally taken as evidence that fractionation is involved to some extent [Martin *et al.*, 2003; Schöpe *et al.*, 2007]. I therefore propose that such interfacial fractionation takes place not just in nucleation but during the subsequent growth of the crystal, so that it is facilitated by the low-density gas layer that forms in front of the crystal when inside the G-L binodal. In contrast to the collective diffusion discussed in Section 6.4.3, self diffusion is hindered by particle density, so that fractionation would be frustrated outside the G-L binodal, where the fluid side of the interface is at a relatively high density compared to that in the gas-mediated case. This effect is essentially similar to that proposed in Ref. [Burke *et al.*, 1988] to explain why crystallisation can be rather fast in a (monodisperse) Lennard-Jones system for which the diffusion coefficient in the liquid is extremely low – the depleted region formed by the crystal's growth reduces the density around the interface and allows faster diffusion than in the bulk, facilitating further crystal growth. The additional claim I am making is that this increased diffusion can also facilitate the *fractionation* that is necessary for crystal growth in a polydisperse system, even more so if the crystal is surrounded by a distinct gas layer rather than just a depleted liquid.

In addition to an *overall* preference for a narrower size distribution, it seems reasonable to suppose that the crystal may be selective in some way as to which particles are incorporated *where*, on a local basis. For instance, it may be frustrated by the presence of regions occupied solely by very large or small particles. Whatever the nature of any local size ordering, one expects that crystal should have some kind of preference as to how particle size is distributed within it, given the fact that the largest and smallest particles in the $\sigma = 0.06$ system differ in size by around 40% of the mean. This local ordering is conceived of as being in addition to any fractionation in terms of the overall balance of particle sizes between the phases.

From the outset, it must be stressed that a full explanation of the results in the previous section requires further work to elucidate and test other possible contributing factors, such as the influence of polydispersity on the equilibrium phase diagram. The proposals here are motivated by the dynamical nature of my simulation approach and by the intuition that the increased diffusivity of particles in the gaseous layer would help any dynamical sorting process at the crystal interface, a consideration which I expect to be a significant aspect of any full explanation of the results. In the following, I perform further analysis on the simulation data in an attempt to detect the presence of fractionation processes.

6.5.2 Polydispersity of the crystal

The results now presented concern the gas-mediated and gas-free $\sigma = 0.06$ simulations. To be sure of sufficient sampling in the crystal, I allowed these to run up to $t \approx 9000$. By this time, the crystal interface in the gas-mediated simulations reached $x \sim 8$, while the gas-free simulations had only grown to $x \sim 3$. I check for fractionation by measuring the mean diameter $\langle d \rangle_{cryst}$ and polydispersity σ_{cryst} in the crystal. To within statistical error, $\langle d \rangle_{cryst}$ was equal to the parental mean, $\langle d \rangle$.

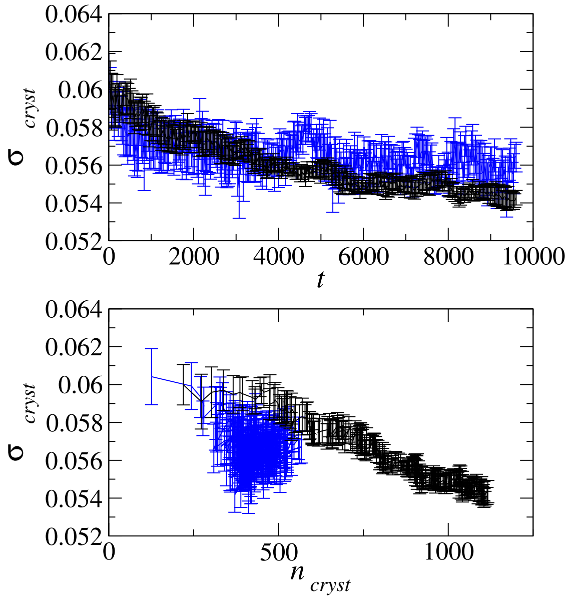

Figure 6.11: Crystal polydispersity as a function of t (upper pane) and n_{cryst} (lower pane) for the gas-mediated (black) and gas-free (blue, larger error bars) crystals. The error on the n_{cryst} axis is approximately 30.

However, changes in σ_{cryst} were detected. In FIG. 6.11, the evolution of σ_{cryst} in the gas-mediated and gas-free crystals is shown, as a function of time and of n_{cryst}, the mean number of particles in the crystal. In the gas-mediated case, there is a clear reduction in the polydispersity of the crystal as compared to the parental $\sigma = 0.06$, demonstrating that the crystal is selecting a narrower subset of the parent size distribution, eventually attaining a value of $\sigma_{cryst} \approx 0.054$ at the end of the simulated time. This is qualitatively consistent with Fasolo and Sollich's equilibrium calculations on hard spheres, in which even a small parental polydispersity was found to lead to reduced polydispersity in a crystal coexisting with a fluid [Fasolo and Sollich, 2004]. This measurement in itself is significant: to my knowledge, such fractionation has not previously been measured in the dynamical crystal growth of polydisperse spherical colloids[1],

---

[1] Fractionated crystal growth in colloidal *platelets* was observed in Ref. [Byelov *et al.*, 2010]. During preparation of this thesis, experimental results on fractionated crystal growth in spherical colloids appeared which show, in a system whose size distribution has a 'fat tail' at low particle size, rejection of the smallest particles from the growing crystal [Leocmach and Tanaka, 2013].

although it is proposed as an explanation of slow crystallisation in such systems [Martin *et al.*, 2003; Schöpe *et al.*, 2007].

The data for the gas-free crystal are subject to significantly larger error, because very little crystal has actually formed, but appear to be consistent with a small reduction in polydispersity corresponding to the frustrated crystal growth. In any case, it is clear that the successful growth of a crystal involves a measurable reduction in polydispersity from the bulk fluid, indicating a fractionation process which would be enhanced by the gas layer.

6.5.3 Local size correlations

To check for local size ordering, in addition to the overall preference for a narrower particle distribution, I define a local diameter-diameter correlation function $\xi(r)$:

$$\xi(r) = \langle d_i d_j \rangle_r - \langle d_i \rangle^2 \tag{6.7}$$

where the averaging in the first term is over neighbours separated by a distance r. The function $\xi(r)$ is therefore analogous to a 'radial distribution function,' but for size correlations, as opposed to density correlations, and tends to zero in the limit of an ideal system (i.e. one in which particles do not interact and 'size' just becomes a label which does not affect the behaviour of the particles).

Measurements are taken from the simulations by binning in r, and can be made within either the crystalline or amorphous regions, to observe any differences induced by crystallisation. Then, the averaging over the second term in Equation 6.7 is over all particles *in that phase.* I compare $\xi(r)$ in the gas-mediated $\sigma = 0.06$ crystal with that outside the crystal in the same system (remember that in the gas-mediated case, 'outside the crystal' refers to a metastable coexistence of gas and liquid in the bulk). I also measure $\xi(r)$ outside the crystal in the gas-free case, i.e. a homogeneous fluid of $\sigma = 0.06$.

The results of this analysis are shown in FIG. 6.12, in which are also shown the measured radial distribution functions $g(r)$ corresponding to the $\xi(r)$ data. The data have been averaged within each simulation over $t = 8,000 - 9,000$, during which time both $\xi(r)$ and $g(r)$ appeared essentially static due to the very slow growth of the crystal (whose size scales as $t^{1/2}$), and then averaged over the independent simulations.

The comparison reveals a fascinating '$g(r)$-like' appearance to the local size correlations both inside and outside the crystal. However, closer inspection reveals that whereas the oscillations in $\xi(r)$ for the crystal appear in phase with those in its $g(r)$, those in the fluid – both the homogeneous and G-L separating fluid cases – seem to be approximately in phase quadrature with $g(r)$.

The proposed explanation for these findings is as follows. We will consider first the gas-mediated crystal, then the two fluid datasets. Before starting, I note the known fact that, in a polydisperse system, larger particles will tend to contribute a stronger signal to structural correlation functions. See, for instance, the measurements of partial static structure factors in Ref. [Weysser *et al.*, 2010] and Section 4.2.3. One should therefore be guided by this in interpreting $\xi(r)$: e.g. a positive value indicates a correlation between either 'big-big' or 'small-small' particle pairs, but will contain a stronger contribution from 'big-big' pairs due to the stronger overall structuring of large particles. The terms ‘big’ and ‘small’ are here shorthand for particles greater or lesser in size than the mean $\langle d \rangle \equiv 1$.

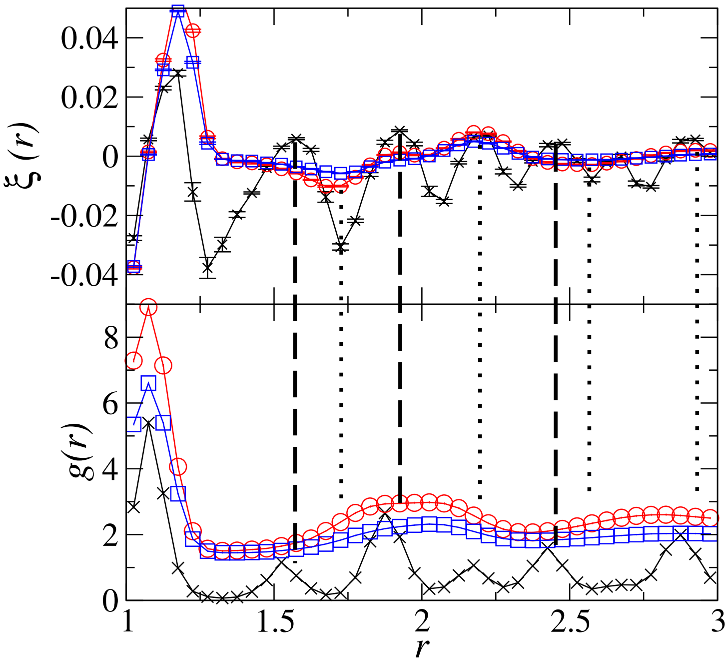

Figure 6.12: Upper pane: the local size correlation function $\xi(r)$ as defined in the text, measured at late times for systems of polydispersity $\sigma = 0.06$. Black crosses indicate the gas-mediated crystal, red circles the gas-mediated outside-crystal region (comprising gas and liquid coexisting), and blue squares the gas-free outside-crystal region (a homogeneous fluid). Lower pane: the corresponding radial distribution functions measured in the same systems. Colours and symbols as for the upper pane. Both the outside-crystal $g(r)$ datasets have been multiplied by a factor of 2 for ease of visualisation. Lines between the data points are guides to the eye. Vertical lines linking the panes are provided to show the relationship between the functions – dashed lines for the crystal, and dotted lines for the two fluid datasets, whose dependence on r is essentially identical. The data points are placed in the centre of each bin in the r axis.

Crystal

In the crystal, the features in $\xi(r)$ seem to more or less mimic, and are in phase with, those of $g(r)$. Where there is a peak in $g(r)$, i.e. a pair of neighbours is likely to be found separated by this distance, it is likely that those particles will show a slight positive correlation in terms of their size, indicating that the crystal structure has a weak preference for particles in the same region to be of similar size, presumably to minimise distortion of the lattice. Conversely, particle pairs separated by the unusual distances corresponding to minima in $g(r)$ have a rather strong *negative* correlation (see the first minimum in $\xi(r)$ at $r \approx 1.3$). This is consistent with a scenario in which these neighbour separations are associated with the presence of ‘wrong’, i.e. unusually-sized, particles, for a given region of the crystal. They may, for instance, result from an interstitial defect, with a particularly small particle squeezing itself in amongst a region of average or slightly large on-lattice particles. The relationship between $\xi(r)$ and $g(r)$ in the crystal is therefore due to the need to minimise lattice distortions over small regions of the crystal, and the fact that unusual neighbour separations will tend to be associated with particles of unusual size for a given region distorting the crystal structure. Since the crystals grown in my dynamical simulations are not necessarily at equilibrium, it would be interesting to know whether the effects described here are present to a similar or greater extent in an equilibrated polydisperse crystal. I intend to pursue this question in future work.

Fluid

Both the homogeneous (gas-free, point 2 in FIG. 5.1) and G-L separating (gas-mediated, point 1 in FIG. 5.1) fluid regions show near-identical behaviour with respect to the r-dependence of $\xi(r)$ and $g(r)$, so I will not distinguish between them for the purposes of this analysis. In the fluid, $\xi(r)$ appears to be approximately in phase quadrature with $g(r)$, so that minima or maxima in $\xi(r)$ appear, respectively, halfway up or halfway down the slope around a peak in $g(r)$.

This can be interpreted in terms of shells around a nominal test particle i, in analogy to the explanation of similar oscillations in terms of shells near a hard wall in Refs. [Buzzacchi *et al.*, 2004; Pagonabarraga *et al.*, 2000; Yu *et al.*, 2004]. Consider that a single shell of surrounding particles corresponding to a peak in $g(r)$ may be roughly divided into those which are nearer than average and those which are further away than average. In order to most efficiently fill space, particles which are closer should be small, and those which are further away should be large. Put another way, particles which are further away should be further away *because they are big*; placing small particles further away than is required by their size would waste space, as it were, reducing the packing efficiency and increasing the free energy of the fluid. Given the stronger structural signal from larger particles [Weysser *et al.*, 2010], the data are dominated by the case where particle i is big. Therefore, the smaller, closer, particles in the shell will on average contribute negatively to $\xi(r)$, resulting in a minimum therein at a value of r slightly less than the peak in $g(r)$. Conversely, the larger, further away particles contribute positively, giving a peak in $\xi(r)$ just outside the peak in $g(r)$.

In the Density Functional Theory study of Ref. [Pagonabarraga *et al.*, 2000], an observed phase quadrature between the size distribution *near a wall* and the mean density profile (roughly corresponding to our $\xi(r)$ and $g(r)$ functions) was explained in these terms. These findings are therefore significant in showing that similar 'local size segregation' appears to be present in the bulk fluid, rather than being solely the result of spatial inhomogeneity. I note for completeness that measuring $\xi(r)$ in a cubic system with no crystal template (so that the system is truly homogeneous and isotropic) had no significant effect on the behaviour of $\xi(r)$; the measurements do seem to capture the behaviour of a bulk fluid.

Thus, the observed relationship between $\xi(r)$ and $g(r)$ in the fluid is a manifestation of the need for efficient packing in the fluid, coupled with the tendency for larger particles to contribute more strongly to the structuring signal. The fact that local size segregation previously observed near a hard wall appears to be manifest also in a bulk fluid is very interesting, and merits further investigation. In particular, I aim to investigate in future work whether this same effect is present in an equilibrium polydisperse fluid.

Comparison with G-L fractionation

It is interesting to compare this fractionation with that discussed in Chapter 5 for G-L separation. There, no fractionation in terms of σ between the phases was detected in the absence of a skewed parent distribution, even with parental $\sigma_p = 0.2$ – this was consistent with the findings of Evans' perturbative theory [Evans, 2001] in which the fractionation of the n'th moment (e.g. variance, σ^2) is found to depend linearly on the parent value of the $n+1$'th moment (e.g. skew). In contrast, when a crystal phase is involved, it seems that some amount of fractionation in σ is desirable even for a parent of zero skew (i.e. the pseudo-Gaussian used here). This finding corroborates the calculations of Fasolo and Sollich [Fasolo and Sollich, 2004] for hard spheres.

In itself, it is not surprising that, at finite polydispersity, the more exact theory would capture this effect that the perturbative theory does not – in any case, the theories embody different approaches to the problem and each have their own advantages. However, it is interesting to note that, while Evans' theory held up to the non-skewed $\sigma_p = 0.2$ in the G-L case, here the *crystal* phase introduces some fractionation of σ even at much lower σ_p, which is not captured by the theory. This to be expected: heuristically, the scale for what constitutes 'significant' polydispersity is set now by the very low polydispersity tolerable by a single crystal, so that the perturbative approximation breaks down at correspondingly smaller values of σ.

6.5.4 Summary

The results above show that the crystal phase is able to reduce its polydispersity relative to the fluid even as it is growing. In addition, I have measured a local radial size correlation function $\xi(r)$ to quantify the nature of local size ordering in each phase. Both the overall form of $\xi(r)$ *and* its relationship to $g(r)$ must change qualitatively between the fluid and crystal phases. This indicates a further, local ordering process *in addition to* the overall preference of the crystal to reduce its polydispersity. To my knowledge, this phenomenon has not previously been observed, and gives new insight into the detailed structural features of polydisperse phases. I have also found that local size segregation, previously found in a polydisperse fluid near a wall, seems to have a close analogue in the structuring of a *bulk* polydisperse fluid.

I have proposed that the necessary fractionation must take place at the interface and therefore, relying on self diffusion near the interface, is enhanced by the presence of the gas layer. If true, this mechanism would explain the very slow crystallisation in the $\sigma = 0.06$ gas-free case, where no such layer exists.

The $\sigma = 0.03$ simulations provide an interesting intermediate case, in which the crystal growth in the gas-mediated case is initially faster before being overtaken by the gas-free case. In light of the hypothesis presented here, this may be because the gas-free growth is relatively slower until the interfacial fluid is sufficiently depleted to allow adequate diffusion near the interface. In the gas-mediated growth, a gas layer is formed *immediately*, so that growth is initially faster. However, once the gas-free system has sufficiently depleted the interfacial fluid, the interfacial diffusion is enhanced and it now grows the fastest, the gas-mediated system lagging behind for the reasons discussed in Section 6.4.3. In contrast, for $\sigma = 0.06$ the gas-free growth is so slow (presumably due to the greater fractionation required) that the interfacial fluid is not depleted enough for such a crossover to happen on the simulated timescale.

I stress that the ideas outlined in this section constitute only one potential explanation of the results observed in Section 6.4. Further work will be required to determine other factors (e.g. the influence of polydispersity-induced changes to the equilibrium phase diagram), and to more rigorously test the dynamical explanation outlined here, which I expect will remain an important part of any more comprehensive explanation. Nonetheless it is clear that, whatever the mechanisms, the influence of metastability on crystal growth kinetics is far from trivial, and that its role can be qualitatively switched by the presence of a very mild degree of polydispersity.

6.6 Conclusions

In this chapter I have studied the kinetics of crystal growth using a model system in which a single crystal grows from an amorphous fluid. By varying the interaction strength and the polydispersity, I was able to investigate the separate and combined effects of i) metastable gas-liquid separation and ii) polydispersity on the crystal growth process. Taking advantage of the predetermined growth direction of the crystal, I used 1-dimensional concentration profiles to clearly display the phase ordering scenarios when the system is inside or outside the metastable G-L binodal. Inside the binodal, I observed the formation of a gaseous layer coating the crystal, in agreement with approximate free energy curves and with previous experimental observations of three-phase ordering. The diffusive growth of the crystal-gas-liquid 'split interface' was modelled theoretically, and I found approximate quantitative agreement between theoretical and simulation interface growth profiles. This agreement validated the use of two approximations in the theory, namely that the interfaces are locally equilibrated to the densities given by the phase diagram (FIG. 5.1) and that particle transport through the gas layer is approximately ideal gas-like due to the gas's low density.

The dynamic influence of this gaseous layer, in the monodisperse case, was to impede growth by slowing particle transport to the crystal interface, suggesting that, while metastable G-L separation can enhance crystal *nucleation* [Anderson and Lekkerkerker, 2002; Fortini *et al.*, 2008], it may (in the ideally monodisperse case at least) have the opposite effect on the subsequent *growth* of the crystal if the free energy landscape is such that the metastable liquid cannot coexist with the crystal, as was the case here. However, introducing a small amount of polydispersity, corresponding roughly to that found in a 'near-monodisperse' colloidal system, I found that the gaseous layer instead strongly *enhanced* growth.

The crystals grown show an overall reduction in polydispersity versus the parent, consistent with existing equilibrium theory and equilibrium simulations. I have compared this fractionation with that discussed in Chapter 5, noting that, in contrast to the G-L fractionation there, the crystal phase reduces its polydispersity even in the absence of skew in the parent distribution; an effect in qualitative agreement with the hard sphere results of Fasolo and Sollich, showing that the perturbative approximation in Evans' theory breaks down more easily (i.e. at lower polydispersity) when a crystal is involved.

I have postulated that the fractionation process, requiring self diffusion, is facilitated by the gas layer, perhaps explaining why this layer enhances crystal growth in the polydisperse case, rather than impeding it as in the monodisperse case.

Additionally, I observed a process of *local* size ordering in crystal formation. A radial size correlation function $\xi(r)$ resembled the radial distribution function $g(r)$ in each phase. The findings for the fluid phase seem to indicate that local size segregation previously observed near a hard wall also appears in the structuring of a bulk fluid, a finding which merits further investigation. I described and explained the relationship between these two functions in both the crystal and fluid, showing that both the form of $\xi(r)$ and its relationship to $g(r)$ must be qualitatively altered as the crystal is grown. Like the overall fractionation driving reduced crystal polydispersity, this local process would also be enhanced by the presence of the gas layer.

It must be stressed that the proposed explanation requires further investigation and probably does not constitute the full story. Other factors, such as the effect of polydispersity on the equilibrium phase diagram of the system, must be considered – further work will be required to establish what effect those factors have. Also, as with other comparable simulations, hydrodynamic interactions (HI) have been neglected; more complex and expensive simulations would be required to quantify possible effects of HI on the kinetics observed here.

I have presented a detailed dynamical study of crystal growth scenarios in the presence of metastability and polydispersity, and shown how the findings relate to the free energy landscape of the system. Whatever emerges from future work, the results demonstrate the importance of both metastability and polydispersity for soft matter phase transition kinetics and, moreover, that these two factors interact in a complex and previously unknown manner, the causes of which remain to be investigated further.

Chapter 7

Results IV – Short attraction range regime in colloid-polymer mixtures

In this chapter I describe work aimed at understanding the phase behaviour of polydisperse colloid-polymer mixtures in which the attraction range is short compared to the colloid size. This direction was prompted by experimental work on such a system by S. Liddle and W. C. K. Poon in Edinburgh, but serves as a general test of the effects of smaller attraction range on the commonly used Mean-Field Asakura-Oosawa (MFAO) model of colloid-polymer (CP) mixtures. Both the perturbative and moment free energy theories of polydispersity described in Section 2.2.3 can be applied to the MFAO model of C-P mixtures, which motivates an evaluation of its behaviour at low size ratio and a comparison of the resultant fractionation predictions, in this regime, with the new experimental data of Liddle and Poon. Before beginning, I acknowledge their invaluable contribution to this chapter, in the form of extensive discussions and the provision of hard-won experimental data.

7.1 Introduction

Colloid-polymer (C-P) mixtures are a paradigmatic example of a highly tuneable soft matter system. They consist of a colloidal species suspended in a solvent in which are also suspended coils of a nonadsorbing polymer. The polymer induces an effective attraction – dependent on the properties of polymer – between the colloid particles, the physical origin of which was first described theoretically in the Asakura-Oosawa-Vrij (A-O) model [Asakura and Oosawa, 1958; Vrij, 1976]. In this model, the polymers do not interact with one another, but interact sterically with (i.e. cannot overlap with) the colloid particles. Therefore, polymers of radius of gyration δ are excluded from a spherical shell of thickness δ around the colloid particles, whose radius a we will consider for now to be constant across all particles, as is the case in a monodisperse system. Since these shells are *depleted* of polymer no matter what the arrangement of the colloids, there is a free volume benefit to the polymer in having the shells overlapping with one another: there is no free volume cost to the overlap since polymers could not enter these regions anyway, and the overlap means that the *available* volume for polymers in the rest of the system is increased. Therefore, configurations in which colloid particles are close enough for their depletion regions to overlap increase the free volume and thus decrease the free energy of the polymer. This is manifested as an effective attraction between the colloid particles, which is referred to as the *depletion attraction.* Its tuneability arises from the fact that the range and strength of the attraction depend respectively on the size and concentration of the dispersed polymer.

The physics described above is broadly applicable to any system in which a 'depletant' species, viz. the polymer in this case, is dispersed along with another. Colloid-polymer mixtures are important in industry and also, by virtue of their customisable interactions, as a system in which the interactions and therefore phase behaviour of many other systems (such as molecular fluids) can be reproduced for the purposes of testing and formulating theories.

The work in this chapter is aimed towards modelling fractionation in C-P mixtures, particularly in the case where the attraction range, given by the size ratio δ/a, is small, as is the case in an experimental system under investigation by S. Liddle and W. C. K. Poon in Edinburgh [Liddle *et al.*, 2011]. I will begin by outlining a well-known theory used to derive phase behaviour from the A-O model. I then introduce concepts by reproducing results for the gas-liquid (G-L) phase behaviour predicted at moderate size ratio, and the fractionation results of Evans' perturbative theory in this regime.

Moving to smaller size ratio, I show that the MFAO theory, in its original form, becomes unphysical in terms of its predictions of the (now metastable) G-L binodals. This poor behaviour also affects the derived fractionation predictions, and is particularly important given that both main theories of fractionation can be applied to the MFAO free energy when used in C-P mixtures. I find that it can be improved by using an alternative hard sphere compressibility, and use the improved MFAO free energy to make fractionation comparisons with new experimental data from Liddle and Poon.

7.1.1 The Mean-Field Asakura-Oosawa model of C-P phase behaviour

Proceeding from the A-O model, earlier descriptions of the phase behaviour of C-P mixtures treated the effective attraction as a perturbation to the hard sphere system (in much the same way as I have derived free energies for square well systems in this thesis), without considering the degrees of freedom of the polymer, i.e. assuming that polymer concentration is equal between all phases [Gast *et al.*, 1983]. However, there is no *a priori* reason why the polymer should behave in this way – after all, the colloid does not. By definition, it is partitioned between phases, so there is reason to suppose that the polymer may be too.

The first theory to deal properly with the polymer's translational degrees of freedom, and one which consequently has become extremely widely used, is the Mean-Field A-O (MFAO) model derived by Lekkerkerker et al. [Lekkerkerker *et al.*, 1992]. The mean-field element of the theory consists in the approximation that the volume V_{free} available to the polymers does not depend on the colloid particle positions except through the colloidal volume fraction ϕ, and can therefore be written as $V_{free} = \alpha V$, where α denotes the free volume fraction and is a function only of ϕ and the size ratio δ/a. The subsequent derivation is concisely outlined in Ref. [Lekkerkerker *et al.*, 1992]. As usual, phase coexistences are found by equating the colloidal chemical potential

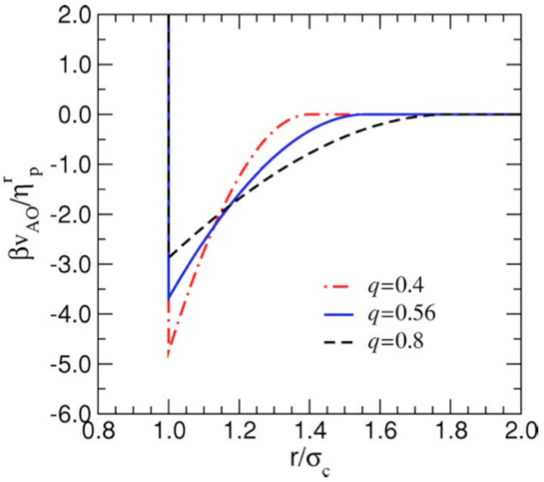

Figure 7.1: The shape of the A-O depletion potential as a function of r/σ_c ($\sigma_c = 2a$ in this figure only) for different values of $q = \delta/a$ as indicated (q represents wavenumber elsewhere in this thesis). The potential depth is normalised by multiplication by $\beta \equiv 1/k_BT$ and division by the polymer reservoir packing fraction $\eta_p^r \equiv (4/3)\pi\delta^3 n_R$. Reproduced (permission obtained) from Ref. [Lo Verso *et al.*, 2006].

and total osmotic pressure in each phase. Both these quantities are found to depend on n_R, which is the number density of polymer *in the volume available to the polymer*. This, being a function of the polymeric chemical potential, is also equal between phases, meaning that the polymer number density $n = \alpha n_R$ differs between phases of different ϕ. That is, the polymer is partitioned between phases, just as is the colloid. Experimentally, n_R corresponds also to the polymer density in a (polymer-only) reservoir with which the system is at equilibrium, via some membrane through which only the polymer and solvent, not the colloid, can pass.

Working in terms of (Helmholtz) free energy density f and setting $k_BT \equiv 1$ as usual, the free energy is a sum of colloidal and polymeric terms respectively:

$$f = \frac{3}{4\pi a^3}\phi \int \frac{Z}{\phi}\, d\phi + n \log n_R \tag{7.1}$$

where $Z(\phi)$ is the hard sphere compressibility. Recall from Chapter 2 that the addition of constant terms, or those linear in ϕ or n, to f does not affect the phase behaviour. These terms have therefore been dropped in the expression above. The constant of integration for the colloidal term is only required when finding crystal-fluid coexistences because these phases, lying on different free energy 'branches', have different Z terms. In that case, the relative values of the integration constants can be set by requiring that, in the absence of polymer, the hard sphere coexistences $\phi_{fluid} = 0.494$ and $\phi_{crystal} = 0.545$ are reproduced.

7.1.2 The hard sphere compressibility Z

To use Equation 7.1, an expression for the hard sphere compressibility Z is needed. (Z specifies uniquely the hard sphere free energy, up to constants – I therefore use 'compressibility' and 'free energy' interchangeably.)

Since the hard sphere system forms the starting point for so many theoretical approaches to related systems, finding an accurate but useable expression for Z has been the focus of much effort – this has resulted in the availability of many expressions of varying accuracy and complexity [Mulero *et al.*, 2001], and an exact result is not available. So far, I have used the Carnahan-Starling form for the fluid [Carnahan and Starling, 1969], and Hall's expression for the crystal [Hall, 1972]. These expressions are those used in Ref. [Lekkerkerker *et al.*, 1992] and are the default choices in the literature.

However, the Carnahan-Starling form, $Z_{CS} = (1+\phi+\phi^2-\phi^3)/(1-\phi)^3$, while convenient and accurate for moderate fluid densities, has the curious property that it diverges at $\phi = 1$. That is to say, the divergence which should correspond to the limiting volume fraction achievable in an amorphous fluid occurs at a value for which *all* the available volume is occupied by spherical, hard particles. This is clearly unphysical: the maximum packing achievable even in a fully ordered hard sphere system is $\pi\sqrt{2}/6 \approx 0.74$. Empirically, the random close packing limit is actually $\phi_{RCP} \approx 0.64$. This feature of Z_{CS} will turn out to be important in the application of MFAO theory to systems with small size ratio δ/a, later in this chapter.

7.2 MFAO phase behaviour – moderate size ratio

In this section I reproduce existing theoretical results on the position of the G-L binodals and the degree of fractionation for moderate size ratio, in order to confirm correct implementation of the theoretical apparatus and introduce the relevant concepts.

7.2.1 The semi-grand potential

In Chapter 2 I described how phase coexistence may be visualised using the common tangent construction, looking for cotangential points on a plot of the free energy density f. In the case of a C-P mixture, however, f is a function of both ϕ and n, so a common tangent plane, rather than a line, would be required, since we are allowing that n may vary between phases as the polymer is partitioned. In order to retain the convenience of a common tangent line, we can move to the semi-grand potential density $h(\phi, n_R)$ [Lekkerkerker *et al.*, 1992]. This is given in terms of the (Helmholtz) free energy density f by:

$$h = f - \mu_p n \tag{7.2}$$

with $\mu_p = 1 + \log n_R$ the polymer chemical potential. Since n_R, unlike n, is equal in both phases, h can be plotted as a function of, and minimised with respect to, colloidal volume fraction ϕ alone. The polymer reservoir density n_R acts, like the square well depth u in Chapter 5, as an inverse temperature which determines the shape of h as a function of ϕ. It can be converted into a dimensionless polymer reservoir volume fraction by multiplying by $(4/3)\pi\delta^3$.

An example plot of the fluid branch of h with size ratio $\delta/a = 0.4$ and $(4/3)\pi\delta^3 n_R = 0.52$ is shown in FIG. 7.2. The double minimum structure familiar from the Helmholtz free energy density f in a pure system is seen, and enables coexistences to be found in a C-P mixture whilst allowing for polymer partitioning between phases as required by the MFAO theory. The

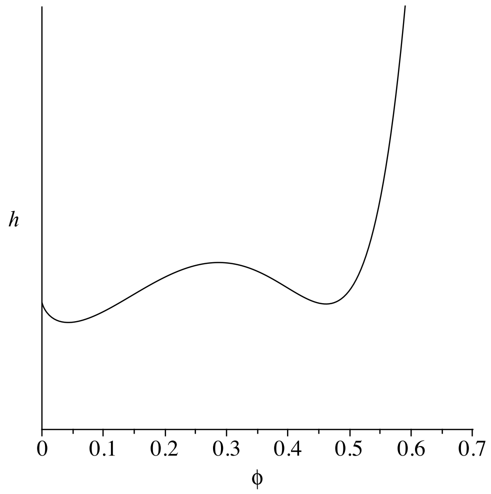

Figure 7.2: Behaviour of semi-grand potential h in a C-P mixture as a function of colloidal volume fraction, for size ratio $\delta/a = 0.4$ and polymer reservoir volume fraction $(4/3)\pi\delta^3 n_R = 0.52$. The Carnhan-Starling expression for the hard sphere compressibility Z has been used. The potential exhibits double minimum structure, enabling the G-L coexistence to be found in the same way as is done using the Helmholtz free energy density f in a colloid-only system.

Carnahan-Starling expression for Z has been used, and as a result there is no divergence in h at the maximum packing fraction for the fluid $\phi_{RCP} \approx 0.64$. Instead, the divergence occurs at $\phi = 1$, which is unphysical but does not severely affect the accuracy of Z_{CS} where only low to moderate fluid densities are involved [Mulero *et al.*, 2001].

7.2.2 Binodals

I now use equality of chemical potential and osmotic pressure, embodied in the common tangent construction outlined above, to map out the G-L binodals (positions of the cotangential points in the ϕ axis) at $\delta/a = 0.4$ as a function of $(4/3)\pi\delta^3 n_R$. This is precisely the calculation used to find binodals at this same size ratio in Ref. [Lekkerkerker *et al.*, 1992], so serves only to lay the ground for the subsequent calculations and ensure correct implementation of the theory. Until explicitly stated, I will be using the Carnahan-Starling compressibility Z_{CS}.

FIG. 7.3 shows the calculated G-L binodals, which are in agreement with the same plot in Ref. [Lekkerkerker *et al.*, 1992], confirming that the MFAO theory has been implemented correctly.

7.2.3 Fractionation

To calculate fractionation, I again apply Evans' perturbative theory of polydispersity, described in Chapter 5:

$$[\langle\epsilon\rangle]_g^l = -[A/\rho]_g^l \sigma_p^2 \tag{7.3}$$

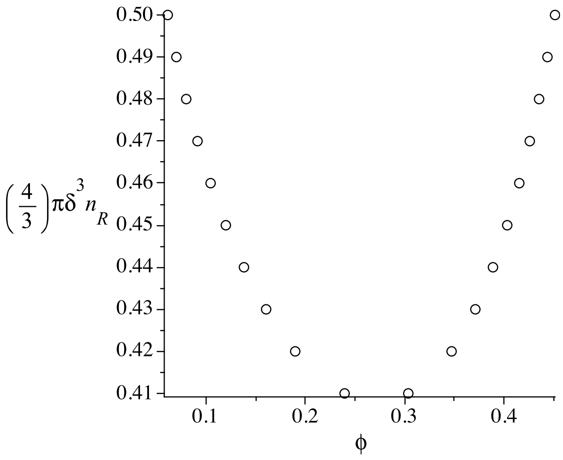

Figure 7.3: G-L binodals at $\delta/a = 0.4$ as a function of polymer reservoir volume fraction (vertical axis), determined using the Carnahan-Starling expression for Z. Since the vertical axis indicates the volume fraction of polymer in the volume *available* to the polymer, the tie lines are horizontal. The results are in agreement with the same calculation in the original MFAO paper [Lekkerkerker *et al.*, 1992].

where, as before, subscripts l and g indicate quantities in the liquid and gas phases, and $A(\rho) = \rho d\mu^{ex}(\epsilon)/d\epsilon$ describes the variation in the excess chemical potential μ^{ex} as a function of dimensionless size deviation ϵ, where ϵ is the deviation (positive or negative) from the overall mean diameter in units of the mean. The variance of the parent distribution is σ_p^2. We can divide through by σ_p^2 in Equation 7.3 to give that the difference in mean diameter between phases, normalised by the overall mean and by the parental variance, is equal to $-[A/\rho]_g^l$. Plotting the latter as a function of polymer reservoir volume fraction allows one to see how the amount of polymer in the system affects the degree of fractionation for any given parent distribution width. FIG. 7.4 shows the results, which reproduce those of Evans' original implementation [Evans, 2001]. It can be seen that increasing the amount of polymer in the system results in stronger fractionation, which results from an interplay between two factors. Firstly, the shape of the $A(\rho)$ curve is changed by the stronger attraction, as was seen for the square well system in FIG. 5.8. Secondly, the points along the curve between which $-[A/\rho]_g^l$ is evaluated are shifted according to the movement of the binodals in FIG. 7.3.

7.3 Decreasing size ratio

Having reproduced known results for the binodals and degree of size fractionation at moderate size ratio by way of introducing the method, I now examine the effects of decreasing the polymer-colloid size ratio δ/a. I will continue to focus on G-L phase separation, i.e. the fluid branch of the free energy. Below $\delta/a \sim 0.24 - 0.3$, the G-L coexistence becomes metastable with respect to crystal-fluid separation [Poon, 2002], as was the case for the G-L separation in Chapter 5.

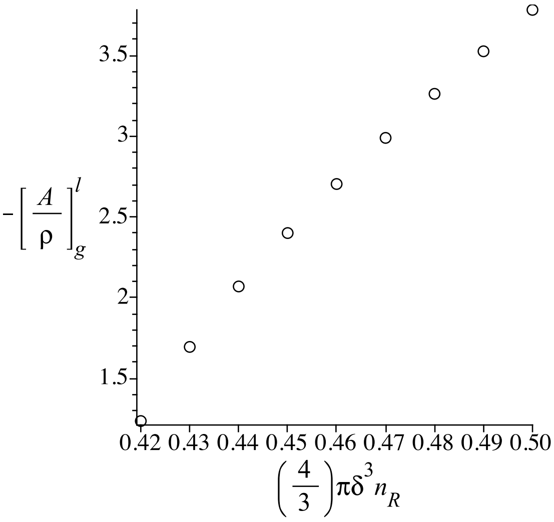

Figure 7.4: Fractional difference in mean diameter between phases per unit parental variance, for the same range of polymer reservoir volume fractions as in FIG. 7.3. The results are in agreement with the same calculation in Evans' paper on the perturbative theory of polydispersity [Evans, 2001].

Replacing Z_{CS}

As noted already, the Carnahan-Starling hard sphere fluid compressibility Z_{CS} results in a divergence in the free energy at $\phi = 1$, whereas it is known that a disordered fluid of monodisperse hard spheres should not be able to pack more densely than $\phi_{RCP} \approx 0.64$. I therefore test the effect of using the alternative 'WKV3' expression [Wang *et al.*, 1996], which diverges at $\phi \approx 0.74$ (still not ϕ_{RCP}), and has been recommended in Ref. [Mulero *et al.*, 2001] as the most accurate hard sphere fluid expression for high densities.

7.3.1 Binodals

I first track the movement of the G-L binodals with decreasing size ratio down to $\delta/a = 0.1$. The results are shown in FIG. 7.5. The overall trend with decreasing size ratio, for equivalent polymer reservoir volume fraction $(4/3)\pi\delta^3 n_R$, is for the coexistence gap between the gas and liquid binodals to become wider. This is because reducing the size ratio in the A-O model results in increased attraction strength and shorter range, as seen in FIG. 7.1, encouraging the liquid to densify to reduce its energy.

For Z_{CS}, even a cursory comparison of the theoretical results at $\delta/a = 0.15$ with the simulation data of e.g. Ref. [Fortini *et al.*, 2008] at the same size ratio shows that the theory vastly overestimates the concentration of the liquid, but such a comparison is not necessary to show that the theoretical results are unphysical – the known packing limit ϕ_{RCP} is already exceeded. Indeed, for the lowest size ratio, even the maximum *crystalline* packing fraction $\phi = 0.74$ is exceeded!

In discussion with W. C. K. Poon, one of the authors of the original MFAO paper [Lekkerkerker *et al.*, 1992], it was suggested [Poon, 2012] that the cause of this is the use of the Carnahan-Starling Z_{CS}, the shortcomings of which were described in Section 7.1.2. The divergence of Z_{CS} at $\phi = 1$ rather than $\phi = \phi_{RCP}$ means that there is a spuriously low free energy cost to the MFAO theory compressing the liquid beyond ϕ_{RCP} – indeed, the actual free energy cost should be *infinite*, corresponding to a divergence at ϕ_{RCP}. Using the alternative compressibility Z_{WKV3}, a much more realistic scenario is recovered. At moderate size ratio, the two expressions give similar results, but the more realistic behaviour of Z_{WKV3} at high concentration (i.e. its divergence lower in ϕ) means that the the liquid does not become more concentrated than is physically possible when the size ratio decreases.

The binodals calculated here show that, although the MFAO theory in its original form has been found to predict reasonably well the overall positions of experimentally observed phase boundaries over size ratios $\delta/a \approx 0.08 - 0.57$ [Ilett *et al.*, 1995; Poon, 2002], it becomes dramatically unphysical with respect to G-L separation (i.e. considering the fluid branch of the free energy only) below $\delta/a \sim 0.25$. It seems that, as suggested by Poon, this is mainly due to there being no 'wall' in the Carnahan-Starling fluid compressibility at ϕ_{RCP}, which allows the predicted liquid volume fraction to become higher than is physically possible. Replacing the Carnahan-Starling expression with one that has a divergence at lower volume fraction gives more realistic behaviour at low size ratio, while giving more or less identical results at moderate size ratio.

Coincidentally, the size ratio $\delta/a \sim 0.25$ at which the liquid-side binodal begins to go beyond ϕ_{RCP} in the original MFAO theory is also around the point where the G-L transition becomes metastable and thus no longer appears in the equilibrium phase diagram; it becomes 'hidden' within the equilibrium fluid-crystal transition. The behaviour of the high concentration side of the equilibrium coexistence is then given, in the MFAO theory, by Hall's expression for Z in the crystal, which *does* have a divergence at $\phi = 0.74$, as expected for the crystal.

The fact that the MFAO theory's poor behaviour at lower size ratio seems largely to be due to the C-S free energy is very important, given that such a marked change results from replacing it with an alternative expression. This *must* be taken into account when assessing the performance of the MFAO theory (or indeed any other which uses a hard sphere free energy as an input). For instance, Ref. [Brader, 2001] (see also Ref. [Dijkstra *et al.*, 1999]) contains a detailed study of both perturbative and free volume (i.e MFAO) theories based on the A-O model, and it is observed that in both cases the gas-liquid coexistence predicted is unrealistic at low δ/a. However, no mention is made of the C-S free energy's unrealistic behaviour at high volume fraction which, as I have shown here, *causes* theories based on it to predict unphysical coexistences, and can be easily remedied. In Ref. [Rotenberg *et al.*, 2004], perturbation theory (rather than the MFAO model) is used. Here again, the C-S free energy is applied, at low polymer size ratio, with no mention of the unphysicality of its divergence or of the manifestly impossible predictions for the metastable liquid binodal which result. Anecdotally, my experience in discussions with colleagues suggests that the C-S free energy is prone to being used without question, or is taken as being an intrinsic part of the MFAO theory, when in fact it is a part which is known to be incorrect when higher fluid densities are concerned, and which is easily replaceable within MFAO or other theories.

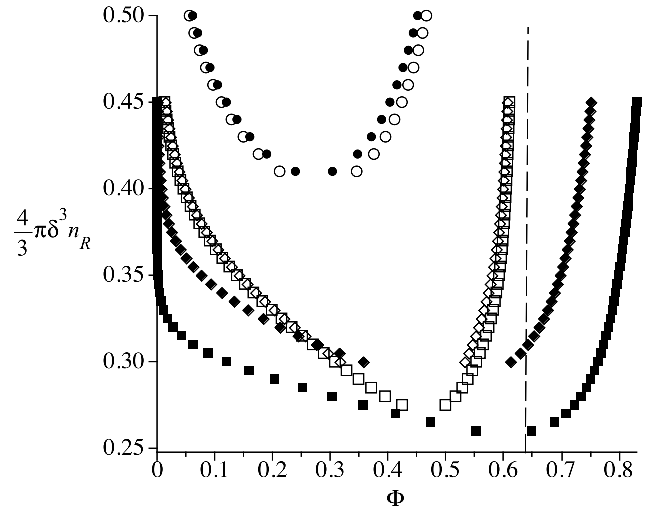

Figure 7.5: Gas-liquid binodals calculated from MFAO, with two choices of the HS compressibility. The full symbols show Z_{CS}, the open ones Z_{WKV3}. Circles, diamonds and squares indicate size ratios $\delta/a = 0.4, 0.15, 0.1$ respectively. The dashed line shows the location of $\phi_{\mathrm{RCP}} \approx 0.64$.

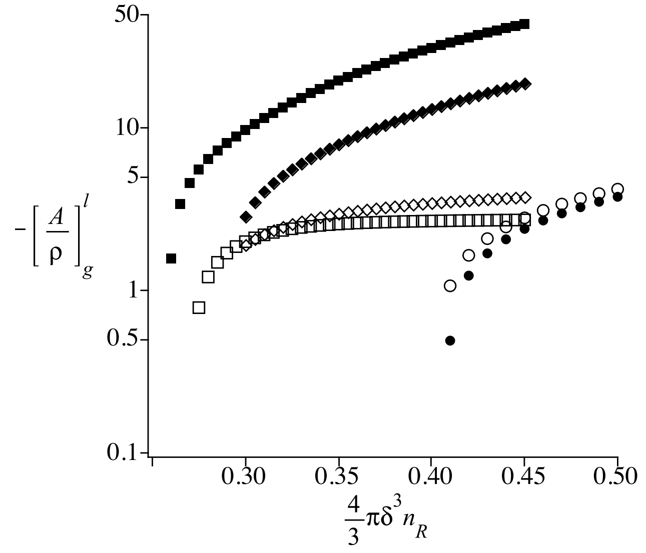

Figure 7.6: Predicted fractionation strength as a function of polymer reservoir volume fraction, within Evans's perturbative theory applied to the MFAO model, with two choices of the HS compressibility. The full symbols show Z_{CS}, the open ones Z_{WKV3}. Circles, diamonds and squares indicate size ratios $\delta/a = 0.4, 0.15, 0.1$ respectively. A log scale has been used on the vertical axis.

7.3.2 Fractionation

I now calculate the corresponding fractionation strengths for smaller size ratios, using the theory mentioned in Section 7.2.3. In light of the unphysical behaviour of the predicted binodals with decreasing size ratio, it seems reasonable to suppose that the predicted fractionation, which takes as input the positions of the binodals, would also become inaccurate. However, the difficulty in characterising polydisperse systems, particularly phase separating ones, means that there is very little experimental data of this type available. In preparation for comparison with some new experimental data in Section 7.5, the fractionation predictions corresponding to the binodals in FIG. 7.5 are shown in FIG. 7.6. It is clear that, as size ratio decreases, the two expressions for the hard sphere compressibility result in dramatically different fractionation predictions.

7.4 Existing experimental tests

In the moderate size ratio regime discussed in Section 7.2.3, experiments in Ref. [Evans and Fairhurst, 2004] have tested and confirmed the scaling of fractionation strength with parental variance, where the constant $-[A/\rho]_g^l$ is fixed but does not need to actually be known. This kind of test does not say anything about the accuracy of any theoretical calculation for $-[A/\rho]_g^l$, i.e. the absolute degree of fractionation at given parental variance. Later, predictions made using the moment free energy method [Fasolo and Sollich, 2005] showed reasonable quantitative agreement with the aforementioned experiments in terms of *both* the absolute degree of fractionation as well as its scaling with the parent distribution width. The smaller size ratio regime was not investigated there either, however. Refs. [Fasolo and Sollich, 2005] and [Evans, 2001] constitute the two most widely cited and tested theoretical approaches to fluid-fluid size fractionation in a polydisperse C-P mixture, and both can take the MFAO free energy as their starting point (in Ref. [Fasolo and Sollich, 2005] the polydisperse 'BMCSL' generalisation of Carnahan-Starling is used, but it inherits the unphysical divergence at $\phi = 1$). Neither has been tested or compared with experiment in the region of smaller size ratios for which G-L coexistence is metastable and in which, as we have seen in Section 7.3.1, the MFAO predictions for the metastable binodals become unphysical.

Therefore, there is a need to determine the accuracy of equilibrium fractionation predictions in this regime, particularly given that the MFAO free energy with Z_{CS} behaves so poorly as size ratio decreases. This question is particularly pressing because i) both the moment free energy and perturbative methods, which have been tested in the moderate size ratio regime may use the MFAO free energy as their input and ii) there is little systematic experimental data on the extent of fractionation, particularly in the small size ratio regime where the G-L transition is metastable and therefore perhaps more difficult to access without the intrusion of the crystal phase. Therefore, assessing how well the perturbative approach to fractionation (applied to the MFAO theory) performs in this regime is of relevance not just to the perturbative theory, but also to the moment free energy method – indeed, *any* approach to predicting the extent of fractionation must presumably involve some expression for the system's free energy.

7.5 New experimental data for small size ratio

With the above discussion in mind, I now compare the fractionation predictions embodied in FIG. 7.6 with some new fractionation data from Liddle and Poon. The experimental system is that described already – in terms of its overall phase behaviour – in Ref. [Liddle *et al.*, 2011]. Work is ongoing on detailed compositional characterisation. This is a rather difficult task, when one is concerned with precisely measuring particle size distributions not just in the system overall but in the separate resultant phases, there is a limited stock of the particular batch under study, and there is no option to simply 'run more simulations' to determine or reduce uncertainties. Therefore, the present comparison should be taken as somewhat preliminary – collaborative work is underway towards a comprehensive and detailed comparison of fractionation and other aspects of the system's behaviour with theoretical and simulation insights.

7.5.1 State points

We will deal presently with four state points within the metastable gas-liquid binodal of the system, characterised by polymer reservoir volume fractions: $(4/3)\pi\delta^3 n_R = 0.388$, 0.396, 0.405, 0.435; colloid volume fractions $\phi = 0.37, 0.22, 0.22, 0.35$ respectively; average size ratio $\delta/a = 0.062$; mean colloid radius $a = 152$nm; parent colloid polydispersity $\sigma \approx 0.1$. The shape and polydispersity of the parent size distribution is currently rather unclear. Small-angle X-ray scattering (SAXS) data assuming a Schulz form gives a parent polydispersity $\sigma = 0.0612$ but the resultant phase all show higher polydispersities than the parent, which seems unlikely. Transmission electron microscopy (TEM) data gives $\sigma = 0.0744$. However, the initial assumption of a Schulz form may well be incorrect, with a better fit being given by a Weibull distribution (with a tail at low size, rather than high size as in a Schulz), of parent polydispersity $\sigma \approx 0.099$. Given the still-evolving situation with respect to the particle size distribution, I simply note the existence of these multiple possible values of the polydispersity at this stage.

I use the SAXS phase characterisation data for fractionation comparisons (TEM data is not presently available for all four state points), and assume a parent polydispersity $\sigma = 0.1$, with the caveat that the predicted degree of fractionation of mean diameter depends on σ^2, i.e. is very sensitive to uncertainty in the parent polydispersity.

7.5.2 Fractionation

The experimental fractionation data available so far from the system in Ref. [Liddle *et al.*, 2011] are shown in FIG. 7.7, along with theoretical predictions using the WKV3 expression in replacement of the C-S one. Fractionation is predicted using Equation 5.3, with the appropriate value of the parameter $-[A/\rho]^l_g$, whose behaviour is shown in FIG. 7.6. Note that the experimental fractionation data are taken after hours of equilibration so that, unlike for the early-stage fractionation in Chapter 5, we can reasonably expect the phases to have reached the equilibrium fractionation magnitude in terms of their compositions.

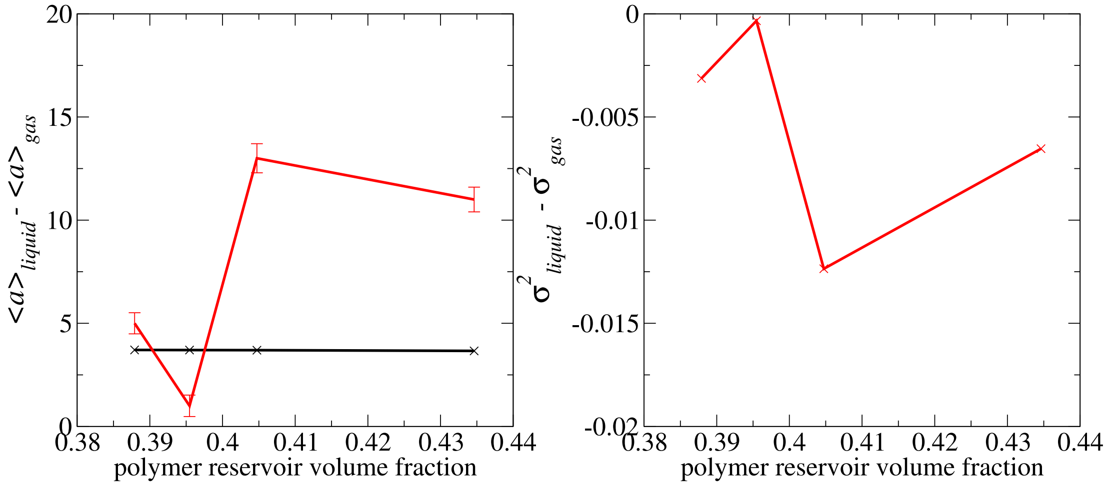

Figure 7.7: Experimental fractionation measurements from Liddle and Poon. The red series show the fractionation in mean radius (left) and variance (right) as a function of estimated polymer reservoir volume fraction. Lines guide the eye. The black line shows the predicted fractionation made using Evans' perturbative theory applied to the MFAO free energy, with the C-S free energy replaced by the WKV3 expression.

Mean radius

The direction of fractionation observed experimentally is qualitatively consistent with theory, in that the liquid contains larger particles on average. There is an apparent overall trend in the experimental data for increasing magnitude of fractionation with polymer reservoir volume fraction (this is the same trend as in theory at higher size ratio, see FIG. 7.4), but this trend appears not to be monotonic. This could be physical but, perhaps more likely, could suggest inaccuracy in the calculated values of the reservoir volume fraction: a scaled particle approximation is used in calculating this value from the experimental polymer *concentration* [Lekkerkerker *et al.*, 1992]. Quantitatively speaking, the theoretical prediction of mean radius fractionation is in order of magnitude agreement. However, I note also that the theory predicts very slightly *decreasing* fractionation magnitude with reservoir volume fraction, in opposition to the experimental data and to theory at higher size ratios. There are a number of possible contributors to this, including the fact that the WKV3 free energy, while vastly better than the C-S one in terms of the position of its divergence, is still only an approximate empirical expression. Further, irrespective of the WKV3 expression, the MFAO theory is known to be subject to other sources of inaccuracy at low size ratio [Brader, 2001; Dijkstra *et al.*, 1999]. In summary, Evans' perturbative theory, combined with the MFAO theory, provides the correct sign and order of magnitude of gas-liquid fractionation at low size ratio, as long as the MFAO free energy is modified by replacement of the C-S expression. A brief look at FIG. 7.6 suffices to show that the predicted fractionation magnitude using the C-S expression would be much too high, since it scales linearly with $-[A/\rho]_g^l$.

Variance

Lack of information regarding the overall skew of the experimental particle size distribution prevents a quantitative theoretical comparison in terms of variance fractionation. However, the relationship between variance and mean radius fractionation is important, so is included here. It can be seen that the liquid is always less polydisperse than the gas, and that the qualitative relationship between magnitude of fractionation at the different state points is the same as when considering mean radius. If, as seems likely, the parent size distribution is of Weibull-like form (i.e. has negative skew, unlike a Schulz), both these features are predicted by Evans' perturbative theory. Equation 7.3 says that the fractionation of the n'th central moment scales with the $n+1$'th central moment of the parent distribution, with the same constant $(-[A/\rho]_g^l)$ in both cases. Since the variance (the second moment) and the skew (the third moment) have opposite signs in the present case, the fractionations of the first and second moments have opposite sign. Even where quantitative experimental characterisation of polydispersity is difficult or is not required, insight into these 'gross' features of fractionation – it's direction, rough magnitude, and the relationship between moments – is valuable.

7.6 Conclusions

In this chapter I have studied the details of deriving fractionation predictions for a colloid-polymer mixture. Prompted by experimental data, I focused on the regime of low polymer-colloid size ratio. This is an important regime, since it is one in which the widely-used MFAO free energy becomes unphysical in terms of its predictions for gas-liquid coexistence. Also, no theory/experiment comparison of fractionation at low size ratio is yet available in the literature.

I successfully reproduced existing findings at moderate size ratio, before reducing the size ratio and showing that the MFAO theory predicts, unphysically, a liquid binodal above the RCP limit of $\phi_{RCP} \approx 0.64$. However, replacing the C-S free energy (which has a divergence at $\phi = 1$) with the WKV3 expression (whose divergence is closer to the right place) provided a drastic improvement at low size ratio, while not strongly affecting the behaviour at moderate size ratio. Irrespective of polydispersity, this is an important, if simple, finding. It should prompt a reevaluation of instances where the MFAO theory (or any theory built on the C-S expression which deals with high fluid densities) has been found to behave poorly, because the MFAO theory is commonly used in a way that suggests the C-S expression is somehow an indispensable part of it. This is not the case and, what's more, the C-S expression is essentially wrong. Coincidentally, however, the regime in which C-S causes the most problems is also the one in which G-L coexistence is metastable. This could explain why such a modification to the MFAO theory has not, to my knowledge, been used before: as the G-L coexistence becomes metastable, the Hall crystal free energy 'takes over' from the C-S one on the high density side of the equilibrium coexistence. It should be pointed out that, although the 'MFAO+WKV3' free energy is clearly better for short size ratio than the original form, a direct comparison of its predicted binodals with experiment was not possible here, so it may still be quantitatively inaccurate. Further work, particularly simulations, would help to clarify this. Importantly, it can also be reasonably expected that theories of other related systems (e.g. square well or Yukawa potentials) would be similarly affected at short attraction range by the C-S free energy – C-S may well constitute a common 'point of failure' in such theories.

Having improved the MFAO theory by use of the WKV3 expression, I calculated fractionation predictions with Evans' perturbative theory for comparison with some early experimental data from Liddle and Poon. The direction and order of magnitude of mean radius fractionation were predicted correctly, although the theory did not capture an apparent trend for increasing fractionation with polymer concentration (which *is* observed in theory at larger size ratio). The WKV3 expression is not a panacea: it too is inexact, and also there are other sources of inaccuracy at low size ratio in the MFAO theory [Brader, 2001; Dijkstra *et al.*, 1999], so further work (e.g. the application of Evans' theory to a completely different colloid-polymer free energy) would be required to determine whether this discrepancy arises from Evans' theory, or from remaining inaccuracies in the MFAO+WKV3 theory. Moving to fractionation of polydispersity between the phases, I found a relationship consistent with the best current experimental suggestion that the system's parent particle size distribution is of Weibull-like form, i.e. has negative skew, in contrast to the Schulz studied in Chapter 5.

Note that the predicted and observed fractionation direction, for all parameter choices in this chapter, was for the liquid to contain larger particles. However, the non-scalable square well potential in Chapter 5, whose range has a 'depletion-like' additive dependence on hard core size, caused the gas to contain the larger particles. The discrepancy between the two arises because, for the A-O potential, the *depth* (as well as the range) of the attraction depends upon particle size (via the different effective size ratio δ/a for a given particle), whereas in the square well case only the range depended on size. This is another example of the sensitivity of fractionation to interaction details, highlighted in Chapter 5.

In summary, I have shown how an existing and very widely-used theory of colloid-polymer phase behaviour may be easily modified to cope better with small polymer-colloid size ratio. The modification to the original MFAO theory performed here is of relevance to any other theory of polydispersity, or phase behaviour generally, which requires a colloid-polymer free energy (e.g. the moment free energy approach to polydispersity). Fractionation predictions derived using the MFAO+WKV3 free energy along with Evans' perturbative theory were in qualitative agreement with emerging experimental data. I have shown that, with an approximate free energy and a perturbative treatment of polydispersity, some of the most important qualitative features of experimental fractionation may be reliably predicted in the low size ratio regime. The ground is paved for further work involving detailed experimental characterisation and the testing of other free energies, towards the goal of precisely predicting and measuring phase compositions in real systems.

Chapter 8

Results V – Testing for diffusion-induced fractionation in a hard sphere crystal

We have seen already how fractionation of particles in a size-polydisperse colloidal system can take place during phase separation (Chapter 5), and how dynamical frustration or enhancement of fractionation processes can affect the rate of crystal growth (Chapter 6). In this chapter, I try instead to test a prediction that the central dynamical process in diffusion-limited crystal growth – namely, the diffusive transport of particles towards the crystal – may *itself* induce fractionation.

8.1 Introduction

The fractionation observed so far in this thesis has been thermodynamic in origin. That is, the driving force for fractionation is the reduction of free energy. Even though the degree of fractionation may be weaker than that required to actually *minimise* the free energy with respect to phase composition, it is in the direction of that minimisation for whichever (equilibrium or metastable) phase coexistence the system is approaching. This is why, in Chapter 5, it was appropriate to use free energy considerations to predict the direction of fractionation between separating gas and liquid phases, although the system was a long time away from reaching equilibrium either in terms of the spatial distribution of the phases or, as it turned out, in terms of the degree of fractionation achieved.

This chapter is concerned with testing a prediction that the dynamics of a phase transition can also *cause* fractionation, which is not thermodynamic in origin and may therefore drive the formation of phase compositions which are not necessarily at equilibrium. Specifically, the prediction of Evans and Holmes [Evans and Holmes, 2001] concerns the diffusion-limited growth of a crystal in a size-polydisperse hard sphere system. In this regime, the limiting timescale for crystal growth is the diffusive transport of particles towards the crystal interface. The prediction is that, because small particles are able to diffuse faster (as discussed in Chapter 4, this is due to their intrinsically higher short-time self diffusion coefficient as well as their ability to more easily avoid geometric frustration in dense systems), they will arrive at the crystal interface in greater abundance than larger ones, and therefore be preferentially incorporated into the crystal. This dynamical fractionation would be in direct opposition to the *thermodynamic* preference for the crystal to incorporate larger particles (in order to maximise positional entropy for those remaining in the fluid) and, because particles become almost stationary once incorporated, could lead to the 'freezing-in' of a nonequilibrium phase composition.

This is a stronger claim than simply to note that equilibrium phase composition may fail to be achieved on experimental timescales. It is a claim that the diffusive transport necessary for the incorporation of particles into the crystal may *actively oppose* the fractionation expected on thermodynamic grounds. It should not be taken to mean that the crystal will *necessarily* contain, on average, smaller particles than the overall (parental) mean – a meaning mistakenly used in Ref. [Martin *et al.*, 2003] – rather, there should be an effect *in this direction* and the possibility remains that the thermodynamic preference for large particles may still 'win' overall.

The work in this chapter is an attempt to detect nonequilibrium diffusion-induced fractionation in simulation. I am able to access a growth regime in which diffusion becomes important, but not a true diffusion-limited regime, due to long-lived transients. I find that the predicted effect of diffusion-induced fractionation is extremely weak, but argue that there are good reasons to expect the theory to underestimate this effect. I achieve very small error bars in the crystal composition measurements, and find that the crystal at no stage actually achieves its equilibrium fractionation, even when diffusion is not a significant factor in growth. The predicted diffusion-induced fractionation proves difficult to test for in this case, in part because the prediction assumes local composition equilibration between crystal and fluid, which is not apparent. The results show that, for the low polydispersity studied here, crystallisation can take place without any appreciable equilibrium fractionation, leading to a nonequilibrium composition being frozen into the crystal. I also take the opportunity to measure local size correlations (previously studied in Chapter 6 for crystals of higher polydispersity than here and with square well rather than purely hard sphere interactions).

The structure of the chapter is as follows. I first introduce the theoretical apparatus for the equilibrium and nonequilibrium fractionation predictions. In Section 8.3 I outline some methodological considerations and the simulation parameters used. In Section 8.4 I then state the quantitative fractionation predictions, before examining the growth profile of the crystal and checking for fractionation and local size correlations.

8.2 Theory

I now briefly go over the relevant equilibrium theory, and introduce the extension given in Ref. [Evans and Holmes, 2001] to incorporate the effect of diffusion-induced fractionation.

8.2.1 Equilibrium

In contrast to much of the other work in this thesis, I here study the hard sphere model – arguably the simplest possible model of interacting particles – with size-polydispersity the only complicating feature. An outline of the relevant thermodynamics has been given already in Section 2.2.1.

The thermodynamic preference at crystal-fluid coexistence in polydisperse hard spheres is for the crystal to incorporate larger particles. For a given volume of crystal, populating it with larger particles results in smaller particles remaining in the fluid, and the resultant increase in the positional entropy of the fluid particles is enough to make this thermodynamically preferable. This heuristic explanation can be stated quantitatively using theory familiar from Chapter 5. Once again, $A(\rho) = \rho d\mu^{ex}(\epsilon)/d\epsilon$, describes how the (excess) chemical potential of a monodisperse system changes as the size of the particles in the system is varied. The parameter ϵ is a dimensionless quantity describing a particle's deviation from mean size in units of that mean, and is by definition zero for a mean sized particle, negative for smaller-than-average particles, etc. Once $A(\rho)$ and the parent size distribution $p_p(\epsilon)$ are known, the equilibrium crystal size distribution is given:

$$p_c^{equil}(\epsilon) = p_p(\epsilon)\left(1 + \left\{(\chi_0 - 1)\frac{\rho_f}{\rho_p}[A/\rho]_f^c\right\}\epsilon\right) \tag{8.1}$$

where the subscripts c, f and p indicate quantities evaluated at the equilibrium, monodisperse, coexistence densities of the crystal and fluid phases and the overall system density, respectively. $\chi_0 = (\rho_p - \rho_f)/(\rho_c - \rho_f)$ is the supersaturation of the system. The important value is the *difference* in A/ρ between the two phases, contained in the term $[A/\rho]_f^c$. For hard spheres, $[A/\rho]_f^c = -3.55$ in units such that $k_B T = 1$ [Evans, 2001]. Because this term is negative, the equilibrium mean size in the crystal phase will be greater than that of the parent distribution.

8.2.2 Nonequilibrium

I now outline how diffusion-induced fractionation is incorporated to yield the nonequilibrium or 'kinetic' crystal particle size distribution. Evans and Holmes' prediction of nonequilibrium size fractionation [Evans and Holmes, 2001] is achieved by requiring that the crystal is only in *local* equilibrium with the fluid immediately surrounding it which, due to the diffusive transport of particles into that region, is expected to contain an extra abundance of smaller particles. The modified expression for the particle size distribution in the crystal is then:

$$p_c^{diff}(\epsilon) = p_p(\epsilon)\left(1 + \left\{(\chi_0 - 1)\frac{\rho_f}{\rho_p}[A/\rho]_f^c - \chi_1\frac{(\rho_c - \rho_f)}{\rho_p}\right\}\epsilon\right) \ . \tag{8.2}$$

Here, χ_1 accounts for the variation of diffusion coefficient with particle size and its effect on the relative densities of particles of different sizes in the fluid at the interface. For details of its derivation, the reader is referred to [Evans and Holmes, 2001]. Finally, I note that, in the diffusion-limited regime, all length scales (including the size of the crystal) should scale as $t^{0.5}$ – this criterion will be used to try and check whether the required regime is reached.

8.3 Method

I now discuss some practical considerations that crop up in attempting to detect the diffusion-induced fractionation predicted by Evans and Holmes, before describing the particular setup used to generate results for analysis.

8.3.1 Considerations

Small polydispersity

Firstly, there is, even more so than in the other work in this thesis, an imperative to get very good statistics for fractionation measurements. In Chapter 5, I dealt initially with a rather small polydispersity of $\sigma = 0.06$ (polydispersity is, as usual, the standard deviation of a size distribution in units of its mean), but was able to go to higher polydispersities, which caused the fractionation signal to grow much stronger and to be detectable with simulations of relatively few particles. In Chapter 6, the maximum polydispersity was $\sigma = 0.06$ but, as noted there, this is actually quite large relative to the maximum polydispersity tolerable by a single crystal phase, and reduction in the crystal polydispersity, though not a change in its mean particle size, was detectable with relatively few particles.

In this chapter, I am attempting to reach a growth regime where diffusion is the most important factor, i.e. is the slowest timescale, but in which the bulk fluid is homogeneous (in contrast to the diffusion-limited gas-mediated simulations in Chapter 6). This requires the crystal to grow many layers, in order to deplete a long enough region of the fluid that diffusion through that region becomes the limiting timescale in crystal growth. There is the intuitive consideration that large polydispersity, which slows down crystal growth, could result in growth being reaction- (including fractionation-) limited for much longer than in the monodisperse case. Therefore, the polydispersity used must be small enough that the ‘reaction’ is not too slow, and diffusion has a chance to become the slowest process. In turn, this necessitates the measurement of many crystalline particles – we have seen in Chapter 5 that fractionation signals are weaker if the system’s polydispersity is lower.

Bulk depletion

The prediction of Evans and Holmes pertains to a regime where growth is diffusion-limited, but the ‘bulk’, i.e. the region far from the crystal, from which particles are being transported, still contains the same size distribution as the parental fluid. Otherwise, the depletion of some species of particle from the bulk would add another complicating factor into the determination of the crystal composition. Experimentally, this represents, for instance, the case where individual crystal nuclei are initially separated from one another by many times the particle diameter. In simulation, it is harder to achieve, as will become apparent later in this chapter. The simulation cell must be extremely long in the growth direction, so as to allow the crystal to grow many layers without depleting the ‘bulk’ fluid at the other end of the cell.

In FIG. 8.1, I show an example plot of crystallinity through time, alongside the local volume fraction in a region at the far end of the simulation cell. There is an apparent ‘turnover’ in the crystallinity growth curve to a gradient of approximately 0.5 on a log-log plot, which might have been taken to indicate the change to diffusion-limited growth, but it is clear that this is associated with a sharp drop in the bulk volume fraction. That is, the crystal has begun to deplete the concentration of the far end of the simulation cell, representing the bulk fluid from which the crystal is growing. Therefore, the system has not reached the required regime of diffusion-limited growth *without* bulk depletion.

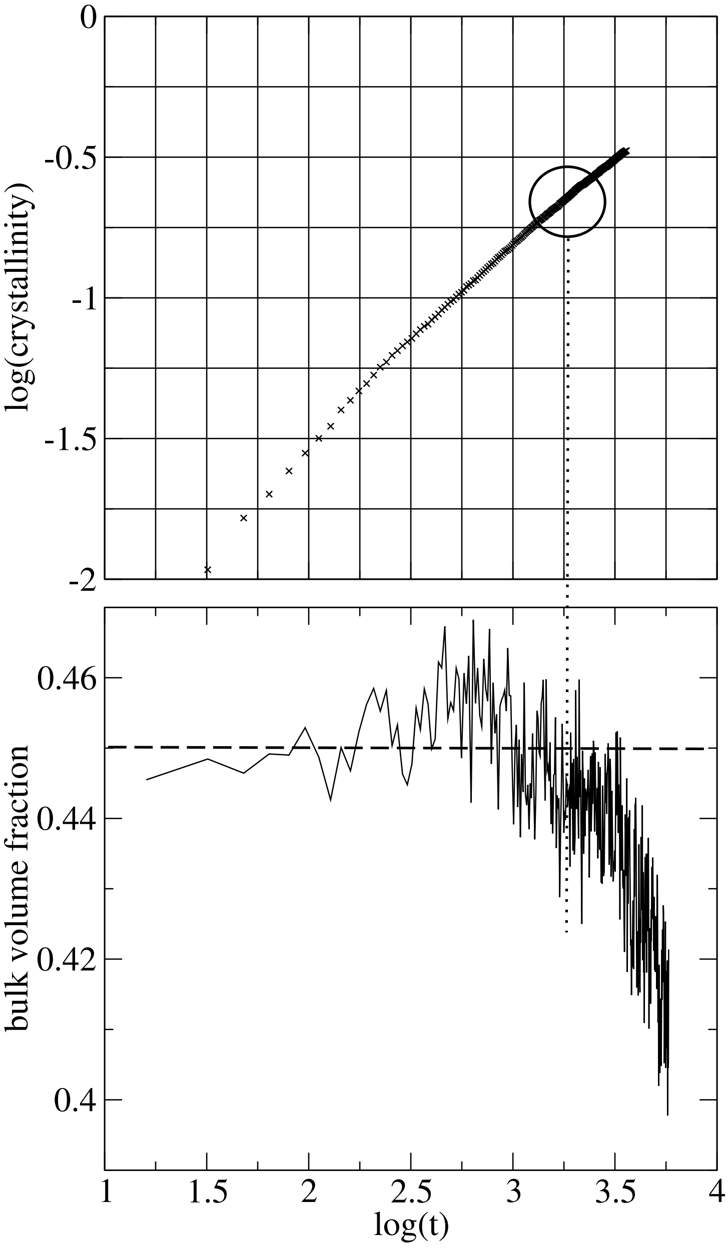

Figure 8.1: An illustrative plot of an apparent turnover in the gradient of the crystallinity curve, which might be taken to indicate a transition to diffusion-limited growth (upper panel). However, as shown in the lower panel, this is in fact the result of the depletion region in front of the crystal reaching far end of the simulation cell. The ‘bulk volume fraction’, averaged over the last 5 units of length of the system, drops sharply from the parent value, 0.45 in this case, indicated by the dashed line. Unlike the others in this chapter, the simulation here used square well interactions – however, the important physics, viz. bulk depletion, is the same as in the hard sphere simulations. Logarithms are taken to base 10.

8.3.2 Parameters used

Having explained and illustrated the considerations impacting upon the choice of parameters, I now detail the parameters used for the simulations presented in the rest of this chapter. The algorithm, crystal templating procedure and basic simulation geometry are the same as used in Chapter 6, and explained in Chapter 3. The particles interact via a purely hard sphere potential, which is infinite at particle contact and zero everywhere else. In the monodisperse limit, therefore, the phase diagram has a single axis, the volume fraction ϕ, and exhibits crystal-fluid coexistence between volume fractions $\phi_f = 0.494$ and $\phi_c = 0.545$. I choose a parent volume fraction $\phi_p = 0.52$ for the present work. The crystal is templated with the (110) face of the FCC lattice, 10 particles wide in the periodic y and z dimensions. The lattice parameter is chosen to match the equilibrium crystal volume fraction, $\phi_c = 0.545$. The number of particles ($N = 100000$) and therefore the length of the simulation cell ($L_x \approx 580$) are chosen very large, so that the crystal can grow many layers without depleting the bulk fluid, as discussed in Section 8.3.1. There is in principle nothing to stop other crystals nucleating in the bulk fluid, which would spoil the attempt to reach and precisely characterise the diffusion-limited regime of the single templated crystal. However, probably thanks to the polydispersity, no independent nucleations occured in the simulations presented here.

Particle diameters d_n are taken from a Bates (pseudo-Gaussian) distribution of polydispersity $\sigma = 0.03$ – this is intentionally rather small, in order to allow the crystal to grow quickly and reach a regime where diffusion is the slowest timescale (again, as discussed in Section 8.3.1). The mean diameter $\langle d \rangle$ sets the length unit.

Since the central piece of physics in this chapter is the difference in diffusivity between small and large particles, it is worth reiterating that the mean step size for particle i is $\Delta\sqrt{\langle d \rangle / d_i}$ (where $\Delta = 0.02$ as usual), so as to give the correct dependence of the short time Stokes-Einstein diffusion coefficient upon particle diameter. It is important to remember, though, that this ‘programmed-in’ dependence on particle size is vastly enhanced, in dense systems like the one under investigation here, by the geometric effects of smaller particles being better able to fit through gaps in their environment – see FIG. 4.11.

9 independent simulations with the same parameters are used to improve statistics, and they are allowed to run long enough for approximately 50 crystal layers to grow. This means that up to 80000 crystalline particles are available for measurement. For curiosity’s sake, I note that these simulations take approximately a month to run per 1500 time units, where the time unit is t_d, the mean time taken for a free particle of mean size to diffuse its own diameter.

8.4 Results and discussion

I now quickly present the theoretical predictions for fractionation of mean particle diameter, for the parameters specified above, before analysing the results of the simulations.

8.4.1 Theoretical predictions

The quantity to be found for comparison is the mean size deviation in the crystal, $\langle \epsilon \rangle_{cryst}$ – due to the definition of ϵ, the mean *diameter* in the crystal is given by $\langle d \rangle_{cryst} = 1 + \langle \epsilon \rangle_{cryst}$. For the equilibrium prediction, using Equation 8.1 to find the expectation value of ϵ, the parameter $[A/\rho]_f^c = -3.55$ and the parent volume fraction $\phi_p = 0.52$ give:

$$\langle d \rangle_c^{equil} = 1.001488 \ . \tag{8.3}$$

Using the expression incorporating diffusion-induced fractionation, Equation 8.2, the parameter $\chi_1 = 0.081156$ results in:

$$\langle d \rangle_c^{diff} = 1.001481 \ . \tag{8.4}$$

Clearly, in theory at least, the correction due to diffusion is quite small. One reason is that the correction is reduced in magnitude relative to that in Ref. [Evans and Holmes, 2001] because I am dealing with diffusion down a 1-dimensional concentration gradient, rather than a 3-dimensional one as would be the case for a spherical, not flat, nucleus. More importantly, the theory in Ref. [Evans and Holmes, 2001] is made tractable by the use of a severe approximation, namely that each particle size species diffuses only according to concentration gradients in *its own* species, controlled by its ideal Stokes-Einstein coefficient $D_0(d)$. (The key parameter χ_1 depends linearly on $dD_0/d\epsilon$, which in the present case is equal to $D_0(\langle d \rangle) = -1/6$). An exact treatment would require i) exact knowledge of how each species' flux depends on other species' concentration gradients (i.e. a non-diagonal, not to mention infinite, diffusion matrix) and ii) incorporation of the scaling of diffusion coefficient with concentration. I do not attempt such a task here. (The work in Chapter 10 suggests a route to a more complete, but still tractable, treatment of polydisperse diffusion which could be applied in future work to the theory in Ref. [Evans and Holmes, 2001].)

However, a possible phenomenologically-motivated improvement might be to replace the ideal gas value of $dD_0/d\epsilon$ with one that takes into account the enhancement of diffusion's species-dependence brought about by finite concentration, described by FIG. 4.11. From that figure, a good approximation seems to be that the dependence of the *self* diffusion coefficient on particle size increases by a factor of 2 for a volume fraction $\phi \approx 0.5$, relative to the dilute limit. It seems reasonable to assume that this provides an order of magnitude estimate as to the scaling of the *collective* diffusion coefficient with particle size, so I make a further prediction by doubling χ_1 in Equation 8.2, which gives:

$$\langle d \rangle_c^{diff\text{-}modified} = 1.001474 \ . \tag{8.5}$$

Of course, it must be admitted that this modification is likely to be quite unimportant compared to the effect of approximating the diffusion matrix as diagonal.

8.4.2 Concentration profiles

I first use concentration profiles along the x axis, as detailed in Section A.1.4, to visualise crystal growth at particular points during the simulation. In FIG. 8.2, I show such profiles at two times during crystal growth. It can be seen that the crystal does indeed achieve a volume fraction approximately equal to the equilibrium value $\phi_c = 0.545$. The fluid immediately next to the interface approaches its equilibrium value of $\phi_f = 0.494$ as the crystal grows. The behaviour of

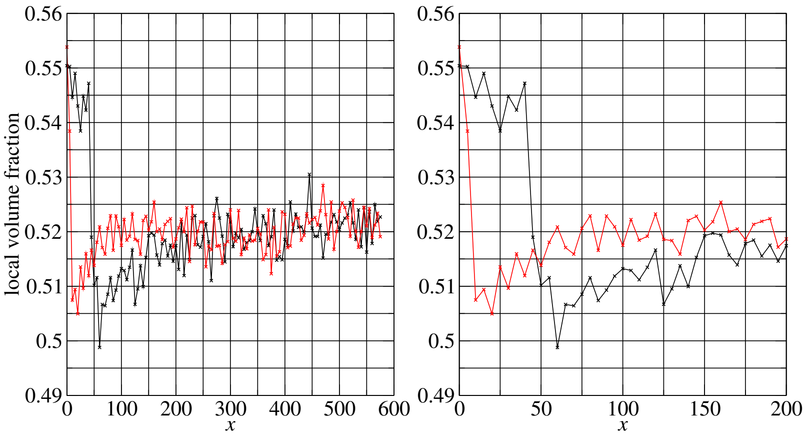

Figure 8.2: Concentration profiles along the x axis from one trajectory at $t \approx 200$ (red) and $t \approx 1400$ (black). The left panel shows the whole x axis, while the right panel is slightly zoomed in. Averaging is over 5 units in x.

these profiles at high-x suggests that the crystal has not yet begun depleting the bulk fluid from its initial $\phi_p = 0.52$, but this will now be explicitly checked for.

8.4.3 Crystal growth rate

I now examine the growth of crystallinity through time, averaged over all 9 simulations, along with the bulk (far-x) volume fraction, to check for depletion of the bulk. The results are shown in FIG. 8.3. As in FIG. 8.1, there is a clear change in the gradient of the crystallinity curve, but it is not now associated with a drop in bulk concentration. However, strangely, the gradient of the curve switches to 0.7, not 0.5 as would be expected for the $t^{0.5}$ growth characteristic of the diffusion-limited regime.

A possible contributing factor towards this is that the crystal is, while growing, also annealing so that the number of particles flagged as crystalline increases slightly more quickly than the rate at which particles are being added to the crystal. However, this explanation seems to fail given that a similar examination of the interface position (which should not be subject to changes due to annealing) as a function of time yields the same 0.7 exponent. Instead, the crystal growth may simply not yet have reached fully diffusion-limited growth, and rather be in some temporary intervening regime.

Ref. [Ackerson and Schätzel, 1995] provides a detailed numerical study of hard sphere growth crystal kinetics. There it is found that while the growth exponent must approach the diffusion-limited value 0.5 eventually, very long-lived transients can be observed where the exponent sits between 0.5 and 1 before finally relaxing. The time taken for this relaxation to happen becomes longer as the dimensionless growth velocity δ – a parameter describing how quickly matter can be added to the crystal – is decreased. In the words of Ref. [Ackerson and Schätzel, 1995], this is because diffusion is then better able to “keep up” with the reaction process for longer. This is in line with the discussion in Section 8.3.1. The idea that the system remains under the influence of such a transient effect seems sensible, given that polydispersity would act to reduce δ and that, as seen in FIG. 8.2, the fluid side of the interface may not yet actually have reached the condition of local equilibration in terms of its volume fraction.

Nevertheless, it is clear that *some* change in the timescales controlling growth happens at the time identified (approximately $t = 800$) so, for the purposes of performing a fractionation analysis, I identify this as the point in time where crystal composition may become subject to the effect predicted in Ref. [Evans and Holmes, 2001]. Although I cannot claim to be accessing the fully diffusion-limited regime in which the prediction was derived, the simulations have reached a regime in which diffusion is becoming more important, as evidenced by the reduction in the growth exponent, and the bulk fluid has not depleted. With this in mind, I proceed to take measurements of the crystal composition.

8.4.4 Fractionation measurements

I now make measurements of the crystal composition, in order to determine whether diffusion-induced fractionation is taking place. This is performed by designating particles behind the crystal interface as crystalline, and monitoring the average properties of those particles through time. The results are averaged over all 9 independent simulations.

Polydispersity

I first measure the polydispersity of the crystal, σ_{cryst}, as it grows. According to the perturbative theory of fractionation, there should be no fractionation in this property between phases, since the skew of the (symmetrical) parent size distribution is zero. This was confirmed for gas-liquid fractionation in Chapter 5, but for crystal growth at $\sigma = 0.06$ in Chapter 6, a reduction in crystal polydispersity was observed even though the parent size distribution was symmetrical. I attributed this to the breakdown of the perturbative approximation of a narrow distribution, given that $\sigma = 0.06$ is actually quite significant if considered relative to the polydispersity that a single crystal phase can tolerate, and noted that the more exact calculations in Ref. [Fasolo and Sollich, 2004] *do* predict a reduction in crystal polydispersity for a symmetrical parent.

However, as shown in FIG. 8.4, no reduction in crystal polydispersity is observed for the present simulations. This indicates that the parent polydispersity $\sigma = 0.03$ does not sufficiently distort the crystal to drive fractionation of polydispersity as the crystal grows: the crystal is happy to simply 'swallow up' the local particle size distribution presented to it as it grows, without needing to reduce the polydispersity before incorporating it. The very fast crystal growth observed here, compared to the $\sigma = 0.06$ simulations in Chapter 6, is consistent with this picture.

Mean diameter

I now move on to measurement of the crystal mean particle diameter $\langle d \rangle_{cryst}$, the property which the equilibrium and nonequilibrium fractionation predictions pertain to. Based on the discussion in Section 8.4.3, I nominate $t \approx 800$ as the point where diffusion of particles towards the crystal becomes, if not the limiting process, at least more important. FIG. 8.5 shows the results, where I have taken measurements of $\langle d \rangle_{cryst}$ both in the whole of the deposited crystal, and in only the region deposited after $t = 800$. Although growth is becoming slower as the simulation

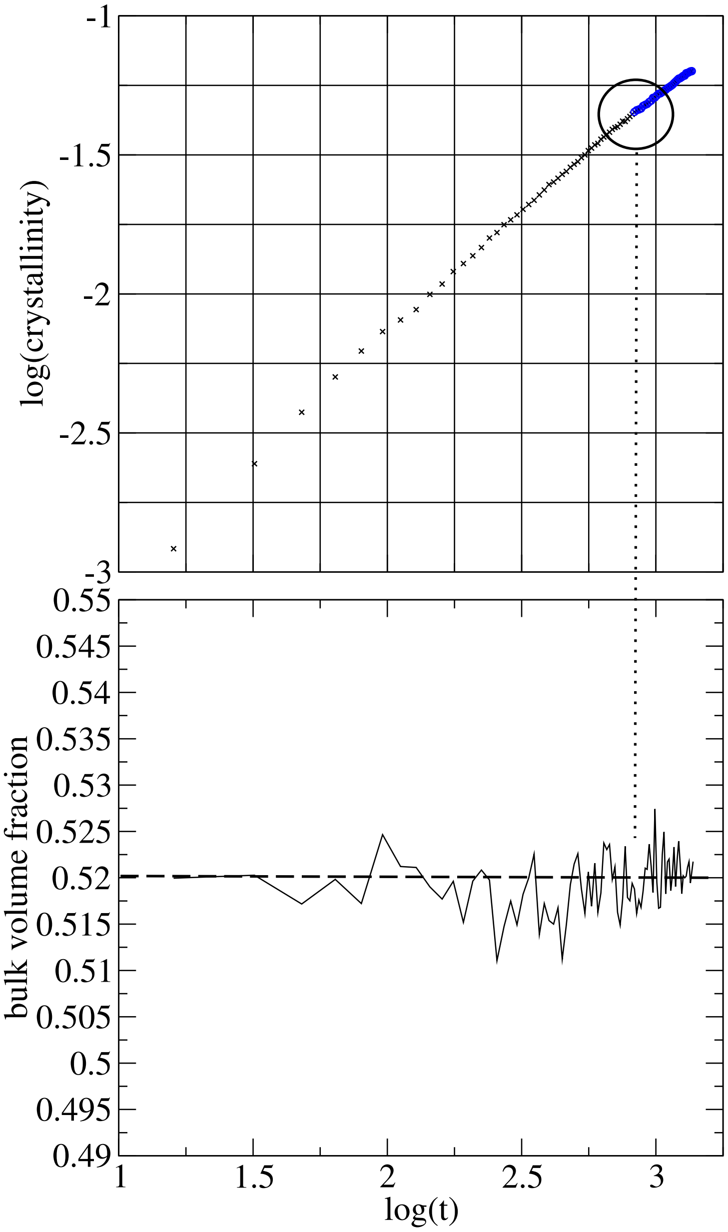

Figure 8.3: Crystallinity through time accompanied by the local volume fraction averaged over $x = 565 - 570$, i.e. the far end of the simulation cell. There is a clear turnover in the gradient of the crystallinity curve (although to a value of 0.7, not 0.5 – discussed in the text), but this turnover, unlike in FIG. 8.1, is not associated with a drop in the bulk concentration from the parental $\phi_p = 0.52$. Blue indicates the region from which the gradient is extracted.

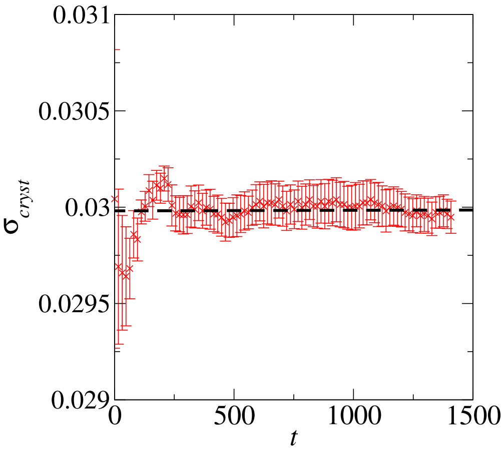

Figure 8.4: Polydispersity of the crystalline region through time. The dashed line indicates the parental polydispersity $\sigma = 0.03$, and error bars show the standard error over 9 independent simulations.

progresses, there are still around 10 layers of crystal deposited in the interval $t = 800 - 1400$ in each simulation, giving up to 24000 particles for analysis. The error bars are extremely small, allowing one to easily see that the equilibrium crystal composition is never achieved, even in the early stages when any diffusion-induced effect should not be important.

Given that we have also seen that the crystal is not selective in terms of its polydispersity, this could perhaps be expected. Notwithstanding the perturbative prediction of Ref. [Evans, 2001] that a non-skewed parent should result in *no* fractionation of polydispersity (which has been shown to break down by $\sigma = 0.06$ if the crystal phase is considered), one can put together an argument that, for the crystal phase specifically, fractionation in mean diameter is *much less important* than fractionation in polydispersity. A nonequilibrium value of $\langle d \rangle_{cryst}$ raises the free energy of the system as a whole, since the fluid suffers a reduction in positional entropy due to comprising larger particles than it would prefer to at equilibrium. However, a too-high value of the crystal polydispersity σ_{cryst} internally destabilises and distorts the crystal *itself*, irrespective of the properties of the fluid phase – this is seen dramatically in Chapter 9. Put another way, the crystal, at all stages of growth, will 'notice' a nonequilibrium value of σ_{cryst} more than it will one of $\langle d \rangle_{cryst}$.

The lack of an equilibrium signal makes it difficult to check for diffusion-induced fractionation: the nonequilibrium prediction of Evans and Holmes assumes *local* equilibrium, *in terms of composition, not just volume fraction,* at the crystal interface. Based on FIG. 8.5, the crystal does not fractionate even at early times. Two broad possibilities arise: i) diffusion-induced fractionation has occurred immediately, i.e. local compositional equilibrium is achieved, but the diffusive effect cancels out the global equilibrium preference for larger crystalline particles; ii) local equilibrium is not being achieved, so that, even in the early stages where growth is not strongly dependent on diffusion towards the crystal (so the crystal interface is confronted with a 'bulk-like' fluid composition), equilibrium fractionation does not occur. In my view, ii) is likely to be a stronger factor; given that, even at $t = 800$, the growth is not *fully* diffusion-limited, it

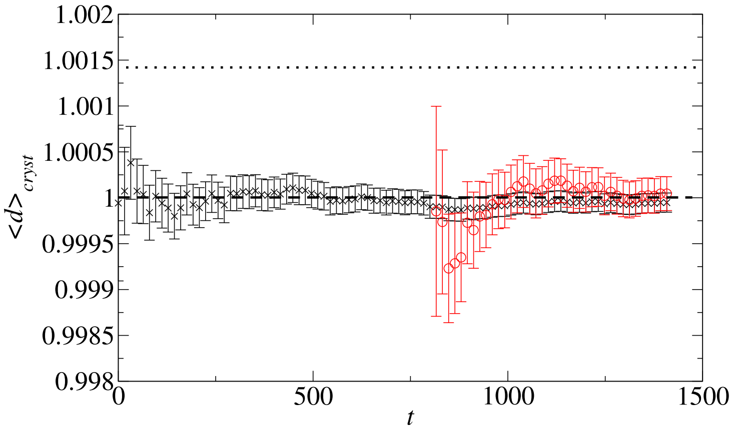

Figure 8.5: Mean particle diameter in the crystalline region through time. The dashed line indicates the parental value. The black crosses are for the whole crystal, whereas the red circles are for only particles deposited after $t = 800$. The dotted line indicates the equilibrium prediction from Equation 8.3. On the scale of this plot, the nonequilibrium predictions are indistinguishable from the equilibrium one. The error bars are the standard error on the average over 9 simulations.

seems unlikely that diffusion-induced fractionation should begin cancelling out thermodynamic fractionation at the very early stages, when growth is reaction-limited. Thus, one could interpret the lack of fractionation at early times as an indication that local compositional equilibrium is not being established, rather than as a signal that the diffusion-induced fractionation predicted (derived in the diffusion-limited regime) has somehow asserted itself when growth is reaction-limited. However, it is not possible on these data to be sure on this point. Conceivably, through time, the interface becomes more and more equilibrated locally (as the dynamics of growth slows down), *but* the diffusive effect continuously becomes more important, leading to an overall cancelling out of two signals.

However, measuring the composition of the crystal deposited after $t = 800$ is a way of checking if, when diffusion becomes more important, the mean crystal particle size drops from its earlier value, which would indicate some impact from diffusion-induced fractionation. Within the error bars, no such drop is observed here. Probably the most important reasons are i) the true diffusion-limited regime has not been reached, due to long-lived transients (see Section 8.4.3) and ii) the magnitude of any diffusion effect that has asserted itself on the simulated timescale may be too small to detect.

Given that the y and z dimensions are periodic, there is a possibility that the exact value of $L_y = L_z$ could *impose* a preferred size of particle on the crystal. To check for this effect, I have also run simulations in which $L_y = L_z$ is increased or decreased by an amount proportional to the predicted equilibrium deviation from mean particle size. The results are unaffected, again showing no fractionation within error, so it appears that the lack of fractionation is not an artefact of the simulation geometry.

Fractionation summary

The results above constitute a very precise characterisation of the composition of a single hard sphere crystal, grown over a long time. I was unable to reach a truly diffusion-limited regime, but did reach one where diffusion becomes more important, evidenced by a drop in the growth exponent. No signal was detected from diffusion-induced fractionation on the simulated timescale, although this could be partly due to progressively better equilibration of the interface masking an increasingly stronger diffusive effect. However, the error bars were sufficiently small to show conclusively that the equilibrium composition is not achieved. For this low a polydispersity, the crystal seems to grow 'blindly', failing to reduce its polydispersity or to increase its mean particle size, with error bars on the latter less than 10% the magnitude of the expected equilibrium signal. This is in stark contrast to the results for $\sigma = 0.06$ in Chapter 6, where significant thermodynamically-driven reduction in polydispersity was observed.

8.4.5 Local size correlations

In light of the measurements of local size ordering in Chapter 6, I take the opportunity to measure local size correlations here, now in a hard sphere as opposed to square well crystal, of lower polydispersity ($\sigma = 0.03$ versus $\sigma = 0.06$). Again, I measure a local size-size correlation function $\xi(r)$ (for neighbours separated by a distance r), defined by Equation 6.7. As before, I also measure the radial distribution function (RDF) $g(r)$ for comparison.

The results are shown in FIG. 8.6. As for the other data (FIG. 6.12), $\xi(r)$ follows the form of $g(r)$, so that nearby lattice sites are occupied predominantly with similar-sized particles, and unusual separations are associated with particles of the 'wrong' size for a given region of the crystal. The magnitude of the oscillations in $\xi(r)$ is reduced here, compared to FIG. 6.12, likely due to the reduced polydispersity and/or reduced volume fraction, but the large crystals allow excellent statistics despite this. The two functions are completely in phase, and track each other closely in terms of the relative heights of each peak as r is increased.

The observation of similar crystal $\xi(r)$ behaviour between the hard spheres of FIG. 8.6 and the square well particles of FIG. 6.12 confirms the intuitively sensible idea that such size correlations are mostly or completely a result of the polydisperse hard particle cores, rather than being caused by the square well attraction. This is as expected, since the structuring of e.g. dense fluids is known to be influenced strongly by the repulsive part of its inter-particle potential. However, it may be interesting to compare with crystals formed by particles with strong, short range attractions, as in Ref. [Foffi *et al.*, 2002]. In these systems, strong enough attractions give a very dense crystal which sacrifices vibrational entropy in order to bring particles as close together as possible, reducing interaction energy. In this way, the crystal 'chooses' energy (the attractions) over entropy (the hard particle cores) so that, in the polydisperse case, the form of polydispersity taken by the attraction could become important.

It is fascinating to observe that, even when the crystal overall shows, as here, *no* signs of fractionation, significant *local* size ordering can be resolved. The measurements show directly that even very weak polydispersity distorts the crystal structure measurably – oppositely-sized particles in a crystal region are associated with unusual centre to centre particle distances. This would not be the case if the crystal lattice 'ignored' polydispersity, as might have been expected

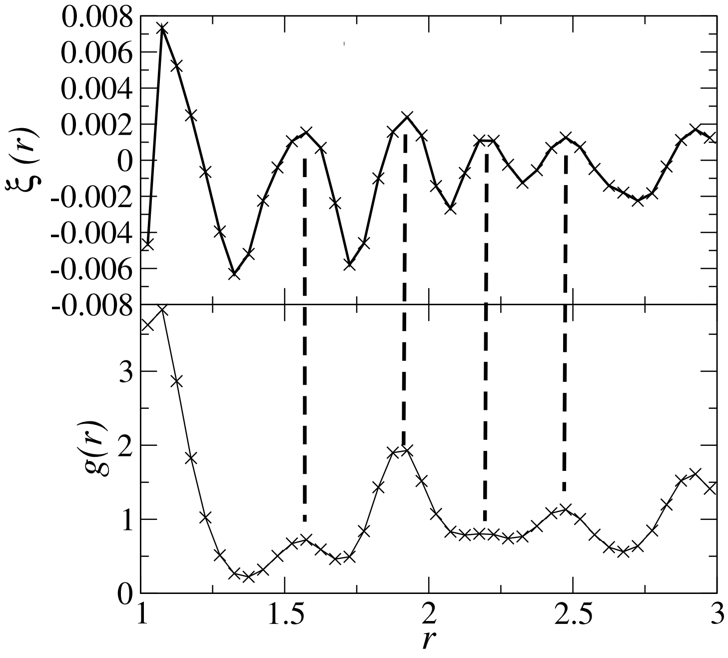

Figure 8.6: Neighbour size-size correlation $\xi(r)$ (upper), as FIG. 6.12, for the present simulations of hard sphere crystals. Results are averaged over $t = 800 - 1200$ then over the 9 simulations, and the error bars are invisible on the scale of the plot. The lower panel shows the RDF, $g(r)$, for comparison. Dashed lines illustrate the relative position of some peaks in the two functions. Lines joining the data points are guides to the eye.

when, as here, the crystal is much less dense than the close-packed limit $\phi \approx 0.74$. This could allow new insight into *how* polydispersity destabilises/suppresses crystals, or reduces their quality (an important consideration in photonics [Allard and Sargent, 2004]). A quantitative, detailed study of local size ordering, including its dependence on the shape of the particle size distribution, the crystal density, etc. would be a very interesting piece of future work, and may also be a good test for detailed theories of crystalline structuring (e.g. density functional theory, perhaps).

8.5 Conclusions

In this chapter I have presented very large simulations of crystal growth in a system of slightly polydisperse hard spheres. This was aimed at providing a test of the prediction of Evans and Holmes of nonequilibrium diffusion-induced fractionation in such crystals [Evans and Holmes, 2001], where thermodynamics favours the incorporation of large particles into the crystal, but small particles should reach the crystal more quickly. After discussing the theoretical outline of this prediction, I introduced the particular considerations relevant to constructing a simulation, motivating my choice of a very large number of particles ($N = 100000$) in order to achieve a long enough x axis that the bulk fluid concentration is not depleted on a long timescale.

I used Evans and Holmes' theory to make quantitative predictions of the nonequilibrium correction to the crystal mean particle diameter due to diffusion. The predicted correction is extremely small but I argued that, due to some severe approximations in the theory, it may underpredict the correction's magnitude. I then analysed the growth profile of the crystal, while monitoring bulk depletion, and found a regime where the growth exponent suddenly drops, although to a value 0.7, rather than 0.5 as would be expected for diffusion-limited growth. I cited literature which suggests this is due to long-lived transients, which would be enhanced by the polydispersity present in my system, as it slows down the 'reaction' of particle incorporation thus allowing diffusion to keep pace for longer before becoming the limiting process. Nevertheless, I argued that this at least indicates a regime where diffusion becomes 'more important'. Future work to access the true diffusion-limited regime in which Evans and Holmes' theory was derived would require much longer times, and therefore larger N (and therefore slower simulations ...) to avoid bulk depletion.

On analysing the crystal composition in terms of its mean particle diameter I found that, even before the point when the growth exponent changes, the equilibrium crystal preference for larger particles is not achieved. These measurements were made with error bars less than 10% of the equilibrium fractionation magnitude. This, and the apparent lack of fractionation in polydispersity also, suggest a growth scenario in which the crystal almost blindly incorporates whatever particle distribution is presented to its interface, due to the low polydispersity here ($\sigma = 0.03$). This is in contrast to the $\sigma = 0.06$ crystals in Chapter 6, where significant reduction in crystal polydispersity was found to be required for growth.

For this reason, a quantitative test of Evans and Holmes' theory, which assumes *local* compositional equilibration of the crystal-fluid interface, was hampered: the background equilibrium fractionation against which any diffusion effect should take place was not manifest at early times when growth should be reaction-, not diffusion-limited. Nevertheless, a qualitative confirmation could still show itself in a reduction of the crystalline particle size through time. However, either because the fully diffusion-limited regime has not yet been reached, or because the effect is smaller than the error bars, this was not detected.

The mean diameter measurement suggests that Evans and Holmes' theory could be improved in part by modifying the assumption of local compositional equilibrium. It seems in some sense that if the crystal 'can' form without fractionation, in practice, it will, leading to the freezing-in of a nonequilibrium composition. In this case, 'can' means: low enough parent polydispersity that the crystal itself is not significantly distorted if it chooses not to reduce its polydispersity. I have argued that the fractionation of mean diameter, being driven by an entropic benefit to the *fluid*, would be (kinetically speaking) less important than fractionation in polydispersity, which has benefit to the crystal *itself*, irrespective of the coexisting fluid's properties. At the low parent polydispersity studied here, though, neither of these kinds of fractionation seems to occur.

I briefly measured local size-size correlations, for comparison with those in the square well, higher polydispersity crystals in Chapter 6. I found similar behaviour in the present case, confirming the intuition that such correlations are predominantly due to the excluded volume interactions of the hard particle cores, as opposed to the square well attraction. Further study of these correlations and their dependence on system parameters would be very interesting, and could allow detailed insight into polydispersity's role in crystal structuring.

In summary, the enterprise of quantitatively or qualitatively testing Evans and Holmes' prediction has been frustrated somewhat by difficulty in reaching a true diffusion-limited regime, and the apparent absence of any equilibrium fractionation even in the very early reaction-limited regime (suggesting that even *local* compositional equilibration might not be taking place). The main result of this chapter is therefore that, with low polydispersity ($\sigma = 0.03$), a crystal can apparently grow without any thermodynamically-driven fractionation taking place to within the small error bars achieved here. The question of diffusion-induced fractionation remains unresolved, and would require significantly larger simulations even than those used here in order to reach the correct regime for testing for it. Finally, the local size-size correlations observed in Chapter 6 would seem to be general to any size-polydisperse crystal composed of particles with hard cores.

Chapter 9

Results VI – Epitaxial growth from dual templates in a polydisperse system

In this chapter, I study the effects on crystal growth of a particular epitaxial setup comprising two crystal templates with different and finely tuned properties.

9.1 Introduction

When studying the growth of crystals in this thesis, I have induced growth by placing a single-layer template of crystalline particles at one end of the simulation cell. Further details are given in Chapters 3 and 6. This situation can be interpreted as representing the limit of a large nucleus which, on the scale of a few particle diameters, appears flat. However, another possible interpretation is that the simulation represents *epitaxy* [van Blaaderen *et al.*, 1997; Hoogenboom, 2002; Hoogenboom *et al.*, 2004; Lin *et al.*, 2000], i.e. the template-controlled growth of a crystal, in which the aim is to exert control over the properties (orientation, lattice parameter etc.) of the resultant crystal. The need, in applications such as photonics, for fine control over the properties of colloidal crystals makes epitaxy a subject of clear technological importance.

As briefly mentioned in Chapter 1, one of the key advantages of soft matter, and colloids specifically, is that their large size may enable direct visualisation using confocal microscopy. As well as facilitating their visualisation, this size also makes colloids more easy to manipulate, on a per-particle basis, than atoms or molecules. An excellent example is the experimental work in Ref. [van Blaaderen *et al.*, 1997]. There, epitaxy was performed by patterning a polymer substrate with holes of characteristic size equal to that of colloids suspended in a solution (radius $a \approx 500\text{nm}$). This is almost an exact analogue of the crystal templating procedure in my simulations (described in Section 3.2.4): the holes are used to arrange the first layer of particles that drops onto the substrate into an essentially stationary template, which induces crystal growth by deposition of particles from the bulk. Note however that, in contrast to my simulations, slow sedimentation by gravity, as opposed to supersaturation of the bulk fluid, is used to encourage deposition. Ref. [van Blaaderen *et al.*, 1997] contains a number of important demonstrations. For instance, it is found that arranging the holes to match the (100) face of the FCC lattice causes a pure FCC crystal to form – this is because, in that growth direction, there is no possibility of stacking faults between sequential layers. Also, adjusting the distribution of templating holes on the substrate allows side-by-side regions of differing crystal orientation to grow. Such complex arrangements of crystal structures may be used in the construction of optical waveguides.

However, no mention is made in that work of polydispersity, which was presumably minimised as far as possible to $\sigma \sim 0.04$. A benefit of a simulation approach, particularly evident in Chapter 6, is that a 'truly' monodisperse system can be compared with one that would, experimentally, often be *approximated* as monodisperse. As I have shown that even low polydispersity is associated with measurable fractionation and local size ordering in the crystal, and given e.g. the strong dependence of a photonic crystal's quality on polydispersities of only a few percent [Allard and Sargent, 2004], it seems sensible to perform epitaxial simulations explicitly focusing on the effects of colloidal polydispersity in the presence of epitaxial templating strategies.

Indeed, the findings in Chapter 6, although they were discussed in the context of crystal growth generically, could be taken as providing some design principles for epitaxial growth in polydisperse systems: a template, combined with a free energy landscape such that a gas layer coats the crystal and facilitates fractionation can provide significant growth, even if the polydispersity is high enough that crystallisation in the bulk, or in the absence of gas, is extremely unlikely.

In this chapter, I move to a templating method involving side-by-side templates of differing lattice parameter and particle size, in a system in which the bulk particle size distribution is composed of two overlapping pseudo-Gaussian distributions, to which the template properties are matched. The aim is to see to what extent manipulation of the epitaxial substrate, as opposed to the bulk dispersion, can be used as a route to controlling crystal properties, in the presence of polydispersity, for two templates which are in the same local environment (viz. their x coordinate, local particle size distribution, etc.).

9.2 Method – dual templates

The templating procedure used so far has been explained in Section 3.2.4, so here I will take that as given and discuss only the modifications made to introduce dual templates. The central idea is that the $y - z$ plane at $x = 0$, where the single template has been positioned up to now, be instead occupied by two templates side-by-side in the y direction, each with different particle size and lattice parameter. Since the aim now is to test the effect of precisely-controlled templates on growth in a less precisely controlled (i.e. polydisperse) dispersion, I assume a scenario in which the templates are created perfectly monodisperse (and, as before, perfectly immobile). This is contrast to the other chapters involving templates, in which the particles in the template were taken from the bulk without being resized, and therefore took on the same size distribution.

9.2.1 System properties

Given the results in Chapter 6, it is to be expected that any selectivity induced by the dual templates will be greater, and easier to detect, if the crystal is shielded by a gas layer which allows greater particle mobility on the fluid side of the interface. For that reason, I use the same inter-particle potential as in the inside-binodal (or ‘gas-mediated’) systems in that chapter: a mean square well range $\lambda = 1.15$, and well depth $u = 1.82$. The overall volume fraction is again set to $\phi_p = 0.34$, and the system comprises $N = 5000$ particles in total.

9.2.2 Bulk particle size distribution

The starting point is to create a bulk particle size distribution composed of two narrower, but still polydisperse, distributions. As in Chapter 6, I use the pseudo-Gaussian Bates distribution. The distributions, labelled A and B, have mean diameters $\langle d \rangle_A$ and $\langle d \rangle_B$, smaller and larger respectively than the overall mean, $\langle d \rangle \equiv 1$. For simplicity, I set the standard deviation of both distributions to be the same, $\sigma_A = \sigma_B$.

9.2.3 Lattice geometry

The two component distributions A and B of the bulk size distribution, described in the previous section, are next used to define the corresponding A and B templates, which are to be placed side-by-side. The monodisperse particle diameter in each template is set to the mean, $\langle d \rangle_A$ or $\langle d \rangle_B$, on which its corresponding distribution is centred. The templates are arranged in the (100) face of the FCC lattice, and the lattice parameter for each is set to give the correct crystalline volume fraction according to the equilibrium phase diagram in Ref. [Liu *et al.*, 2005], *if* the crystal on each template is populated only with particles from its own corresponding component distribution.

An important consideration is that the lattice parameters must be such that an integer number of lattice points from each fit into the z dimension of the simulation box, because the simulation is periodic in the y and z directions. This is enforced by choosing an *exact* integer ratio of lattice parameters and ‘working backwards’ from that ratio to set the other properties of the component distributions A and B. The lattice spacing along the boundary lines *between* the two templates is set to be the average of the spacing of each template individually which, along with the individual spacings, specifies the y dimension of the simulation box.

The A template, composed of smaller particles, is moved forwards slightly so that the right-most (highest-x) surfaces of the particles in each template have the same x coordinate. Finally, the relative proportion of the A and B component distributions in the bulk is set so that, if all the bulk particles crystallised, the crystals would have equal volumes.

Crystal lattice mismatch along the template boundary

At the value of y where the dual templates meet, the geometry of the templates imposes a mismatch on the resultant crystals. It is clear that, since each template fits a different number of lattice points into the z dimension, any resultant crystals which mirror the geometry of the templates must have a lattice mismatch along the template boundary. This would presumably result in a free energy penalty, and would tend to distort crystals near the boundary between them, negatively impacting growth. However, wider templates may cause particles to have to diffuse further in order to be selected for the ‘right’ template, negatively impacting selectivity. Because of the possible complexity of effects along these lines, a range of values of the number of y lattice points will be used for the quantitative observations in Section 9.4, to check for any measurable dependence. Until explicitly mentioned, though, the results pertain to templates which are each 5 lattice points wide in their y dimension, as shown in FIG. 9.1.

9.2.4 Summary, parameters used

The aim of the setup procedure outlined above is to create a bulk particle size distribution which can be decomposed into two closely spaced but individually slightly narrower distributions, and to introduce templates which are precisely matched to the properties of those component distributions. The proportion of each component distribution is set so that there is no strong overall preference in terms of simple probabilities as to which dominates the crystals formed – each template should be free to choose, to whatever extent is preferable, particles from 'its own' component distribution, without either of them running out, as it were. Note that this simplified description should not be taken to imply that the two distributions are actually separate in any sense – as seen below, they may overlap significantly, so that the system cannot know *a priori* which distribution a particle is from.

For the results in this chapter, I have used two combinations of the A and B distributions (and therefore the associated template properties). In all cases, the individual standard deviations are set to $\sigma_A = \sigma_B = 0.04$. That is, the component distributions have equal width. Note that they may not be said to have the same *polydispersity*, since polydispersity refers to their standard deviation in units of their own mean values.

Setup 1

In the first setup, the *overall* particle size distribution has polydispersity $\sigma \approx 0.062$ (and, as always, mean diameter $\langle d \rangle \equiv 1$). The component mean diameters are set to $\langle d \rangle_A \approx 0.959$, $\langle d \rangle_B \approx 1.055$. The particles that form the templates are therefore made monodisperse at exactly these diameters. These values are the result of a choice of a ratio of lattice parameters 10 : 11 – i.e. the A template has 11 lattice points in the z dimension, and the B template has 10.

The resulting template pattern is shown in FIG. 9.1. As can be seen in FIG. 9.2, the overall particle size distribution is weakly bimodal.

Setup 2

The second setup uses a wider spacing of the component distributions, corresponding to a lattice parameter ratio 10 : 12, giving an overall polydispersity $\sigma \approx 0.098$. The component mean diameters are now $\langle d \rangle_A \approx 0.932$, $\langle d \rangle_B \approx 1.118$.

Again, the resultant template pattern is shown in FIG. 9.1. In contrast to Setup 1, the particle size distribution in this case is strongly bimodal, with almost no overlap between the component A and B distributions (FIG. 9.2). One could choose to think of this distribution (and, in fact, that of Setup 1 also) as representing a binary mixture of particles where the polydispersity of *each* of the binary components has been taken into account.

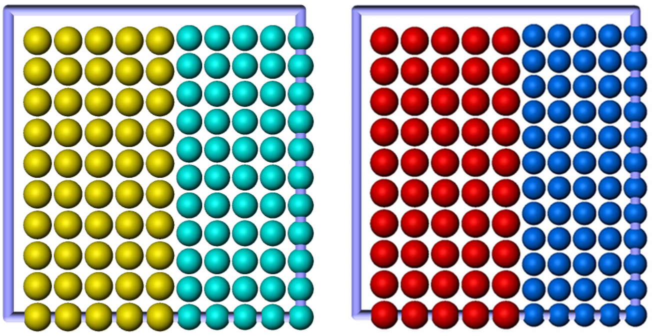

Figure 9.1: Simulation snapshots from Setup 1 (left) and Setup 2 (right) showing the dual templates, in the $y-z$ plane. The horizontal axis is y, the vertical z. Hotter colours indicate larger particles, as described in the text. The ratio of lattice parameters is 10 : 11 in Setup 1 and 10 : 12 in Setup 2, as can be seen by counting the number of particles in the z axis of the snapshots.

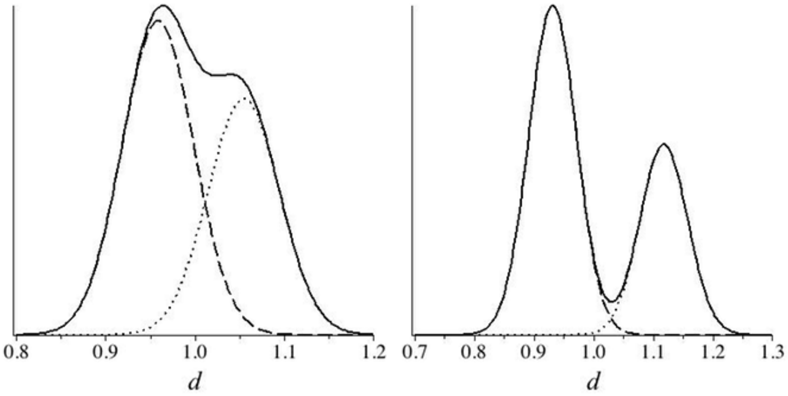

Figure 9.2: Idealised Gaussian forms of the overall particle size distributions (solid lines), composed of the A (dashed) and B (dotted) component distributions. Note that, in simulation, to avoid zero-sized particles and for efficiency, the size distributions are realised as pseudo-Gaussian Bates distributions. The left pane shows Setup 1, the right pane Setup 2.

9.3 Qualitative crystal observations

I now qualitatively study crystal growth in the Setup 1 system, using two representative trajectories. For clarity, I will refer to these two independent trajectories as Setup 1a and Setup 1b.

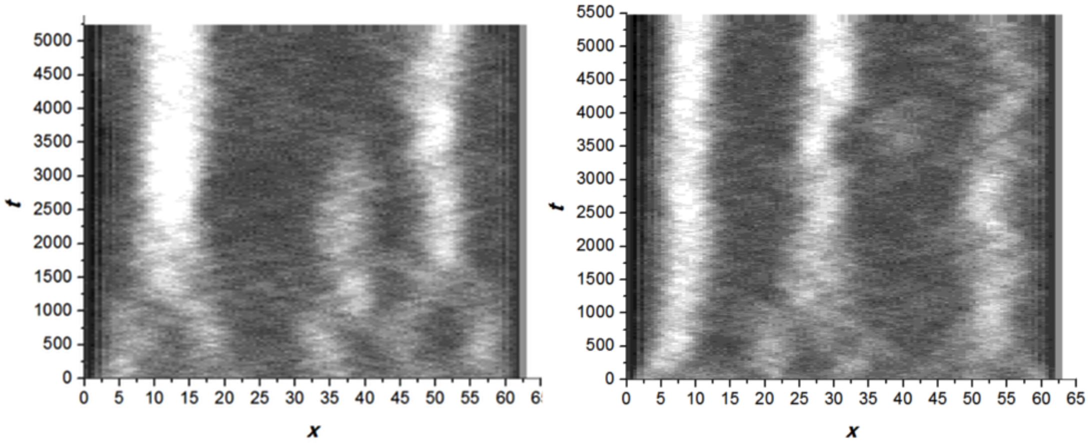

Figure 9.3: Concentration profiles through time, as e.g. FIG. 6.9, for Setup 1a (left) and Setup 1b (right). As with the single template simulations of Chapter 6, a gas layer is present in front of the crystal, and gas-liquid separation happens in the bulk. Note however that, especially in Setup 1a, the crystal interface seems extremely spread-out and unclear. This is associated with poor ordering and incomplete layers in this particular crystal.

9.3.1 Concentration profiles

As in Chapter 6, the best way to illustrate the phase behaviour of the system over the whole simulation time is by plotting concentration profiles, such as FIG. 6.7. FIG. 9.3 shows such plots for trajectories Setup 1a and Setup 1b. It can be seen that only a few layers have grown and, as we will see in the next section, those that have are quite likely to be disordered and incomplete.

The gas layer observed in Chapter 6 (see FIG. 6.9) appears to be present on a cursory viewing of FIG. 9.3. However, closer inspection reveals that the crystal-gas interface region is now much more spread-out and ill-defined, if it can really be said to be present at all. This is illustrated in FIG. 9.4, in which instantaneous concentration profiles for the Setup 1b trajectory are shown at two times during the simulation. At the earlier time, there is quite a sharp concentration drop in front of the crystal, indicating the presence of a gas layer. However, at the later time, there seems to be a *liquid* region in front of the crystal, although the free energy considerations in Chapter 6 should prohibit this local coexistence (the coating of the crystal with a liquid region can be seen for this trajectory also in the snapshot FIG. 9.12). Given the similar polydispersity between Setup 1 and the simulations in Chapter 6 ($\sigma = 0.062$ versus $\sigma = 0.060$), and the identical square well parameters, it seems that this must be related to either i) the slightly different particle size distribution shape or ii) the dual crystal templates.

Perhaps the most likely cause is the use of dual templates in the present case, with the

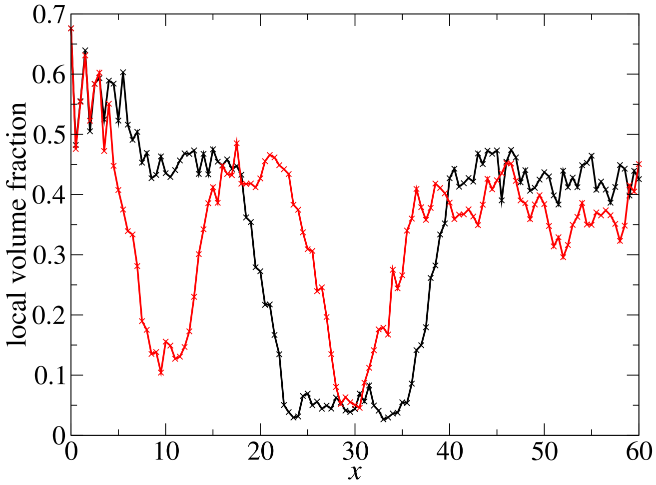

Figure 9.4: Instantaneous concentration profiles from trajectory Setup 1b, at $t = 8000$ (red) and $t = 18000$ (black). At the earlier time, the crystal is coated by a gas region, but at the later time, there appears to be a liquid region coating it. This did not take place for the single-template simulations in Chapter 6, and the present simulations have the same parameters as there, *except* for the presence of side-by-side mismatching crystal templates. A possible explanation therefore is that the dual templates distort the crystal so much that its free energy minimum is effectively 'raised', allowing a common tangent to be drawn with the metastable liquid minimum. Further discussion is in the text.

resultant mismatch along the lattice boundary which I have already discussed. Although the dual templates may induce size-selectivity of particles, it seems unlikely that side-by-side coexistence of two crystals with mismatched lattices is favourable on free energy grounds. The resultant 'dual crystal', with the mismatch along the lattice boundaries, would not be able to reach the true bulk minimum of the crystal branch of the free energy. This could effectively raise the free energy of the crystal minimum so that, comparing FIG. 6.2, a common tangent *could* then be drawn and local crystal-liquid coexistence became possible. Put another way, the two side-by-side crystals may be so distorted by one another's presence such that the pressure and chemical potential are altered from their values in a bulk crystal, and can equal that of the metastable liquid. Either the gas *or* the liquid could then coat the crystal. In agreement with the idea that the lattice boundary mismatch is related to the possibility of liquid coating the crystal, simulations with templates 15 particles wide in y (rather than 5), so that more of the crystal interfacial area is further away from the boundary, did not show any such liquid-coating.

9.3.2 Direct visualisation of crystal layers

To see some possible growth and defect patterns in detail, I use the visualisation software OVITO [Stukowski, 2010] to observe the structure of the first two layers of crystal growth, in Setups 1a and 1b discussed above. The following observations are made at $t = 4368$, at which time around 3 layers of growth have been deposited onto the templates. When interpreting these images, we should be careful to remember that they are instantaneous snapshots: in reality, the whole crystal is fluctuating to some degree, *particularly* when a large number of defects are present, as here.

First deposited layers

First, I compare images of the first layers deposited onto the templates. FIG. 9.5 shows snapshots from Setup 1a and Setup 1b, looking onto the first deposited layer from the high-x side – the template particles are therefore in the background. Two main features are apparent. Firstly, it is clear by inspection of the deposited particles' colours that those deposited onto the smaller particle template (left hand side in FIG. 9.5) are on average smaller than those on the large particle template. This indicates that the templates are being selective to some extent as to which particles they incorporate: one would of course expect an effect of this kind to be especially strong for the first deposited layer, in which the particles incorporated sit alongside completely immobilised template particles.

Secondly, there is an interesting difference between the two trajectories in terms of how many particles occupy each column of the first deposited layer. The columns lying on the boundaries between the templates (whose numbers are shaded green in the figure) in both cases contain 10 particles, rather than 11. This can be taken to indicate that it is easier for the crystal to impose an effective vacancy on the small particle template (which would 'prefer' 11 particles) than to impose an interstitial defect onto the large particle template by incorporating 11 particles into the column. Next, in Setup 1b, the columns *either side* of this boundary column have the correct number of particles for their respective templates. However, in Setup 1a, the column fourth from the left contains 10 particles, whereas the template it sits on contains 11. Note that this is a purely random feature of the trajectory. Perhaps, at the time when this layer was being deposited, no particle small enough to fit into this gap was available nearby – observe that the gap is surrounded in this plane by relatively large particles. We will see below that this missing particle is associated with (but may not necessarily be the cause of) poor subsequent growth of the crystal in Setup 1a.

Second deposited layers

We now move on to the second deposited layers in these two sample trajectories. In order to aid visualisation, I show the second layer against the background of the template particles, with the first deposited layer excluded: this is illustrated from a side-on view in FIG. 9.6. Snapshots of the second layers are shown in FIG. 9.7. The overall impression is that the layer in Setup 1b exhibits significantly better crystalline order than that of Setup 1a. The registration between the second deposited layer and the template is much clearer in the former case, whereas in the latter there are significant distortions and a rather large number of defects.

It is tempting to speculate that the vacancy in the fourth column of Setup 1a's first layer (discussed in the previous section) is somehow responsible for the disorder in its second layer. The region around the vacancy is, in the second layer, occupied by a cluster of rather large particles, which are circled in FIG. 9.7. These large particles are more able to occupy this area than they would be if the vacancy did not exist and, being large, distort the crystal in front of the small-particle template. However, in truth, there is not enough information to make so bold a claim about the growth history of the crystal – perhaps both the vacancy and the poor

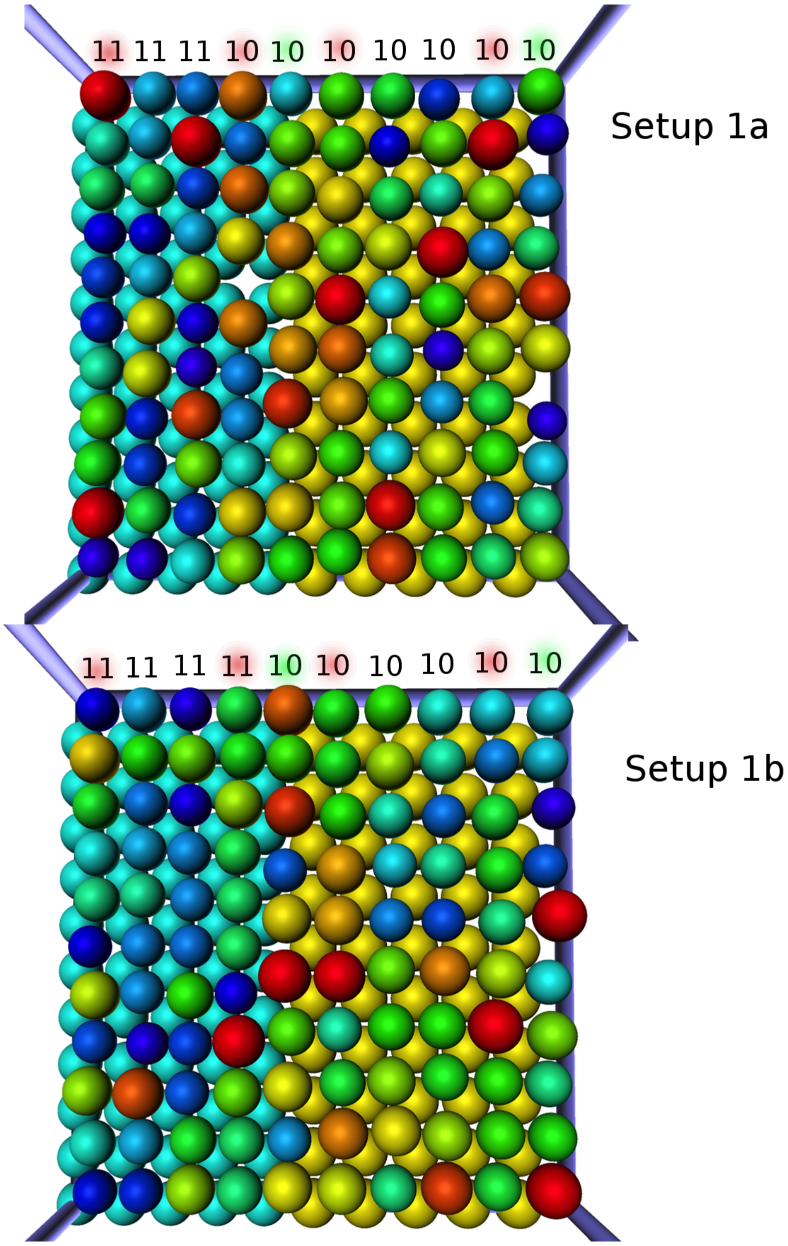

Figure 9.5: Snapshots from two independent trajectories using Setup 1, looking onto the first deposited crystal layer from the high-x side, so that the templates appear in the background. Colours indicate particle size, from $d \leq 0.9$ (blue) to $d \geq 1.1$ (red). Above each column, the number of particles in the z dimension of the first deposited layer is indicated. Green shaded numbers indicate that the column is on a boundary between the templates, and red shaded numbers indicate the columns either side of a boundary.

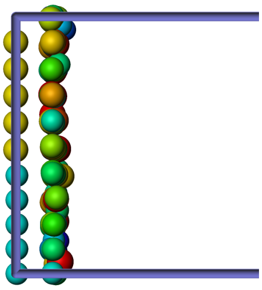

Figure 9.6: Side-on view of which particles are included when visualising the second deposited layer with the template in the background. In this image, y is the vertical axis and x the horizontal one. The particles in the first deposited layer are excluded, while those in the templates are shown in order to make the positions of those in the second deposited layer more clear.

ordering of the second layer are both just symptoms of e.g. a particularly unfavourable local size distribution at the time the crystal layers were forming.

Finally, I note that the apparently poor ordering in the crystal of Setup 1a relative to 1b is corroborated by the instantaneous crystallinity measures of the system at the time when these snapshots were taken. Setup 1a shows 3.5% crystallinity (as measured by the bond order algorithm, described in Section 3.2.5), while Setup 1b, with its better ordering, shows 5.6% crystallinity (the systems are of the same size, $N = 5000$).

9.3.3 Setup 2

As may have been expected, the much wider, bimodal distribution used in Setup 2 results in very little crystal growth on the simulated timescale, with only one complete layer typically growing. The presence of roughly equal amounts of both component distributions seems sufficient to ‘poison’ the system against crystal growth – in contrast to Setup 1, the two distributions are far enough apart that the presence of a particle from one in the crystal forming on the template corresponding to the other results in considerable distortion. As seen in FIG. 9.8, which shows the first and only layer of crystal deposition from a sample trajectory, this is particularly apparent in the small-particle template (on the left of the figure). Any particle in the first deposited layer on the small-particle template which is large enough to be coloured slightly yellow is associated with vacancies of the smaller, ‘correctly’ sized particles on nearby lattice points. As would also be expected, this results in a far clearer selectivity of the templates in terms of particle size, in contrast to Setup 1, where particles of the wrong size for the template could still be accommodated without inducing a defect in the first deposited layer.

This result shows that, although template control can induce strong selectivity in a polydisperse bimodal mixture like that in Setup 2, the resultant crystal growth is likely to be extremely slow for parameters of the type used here.

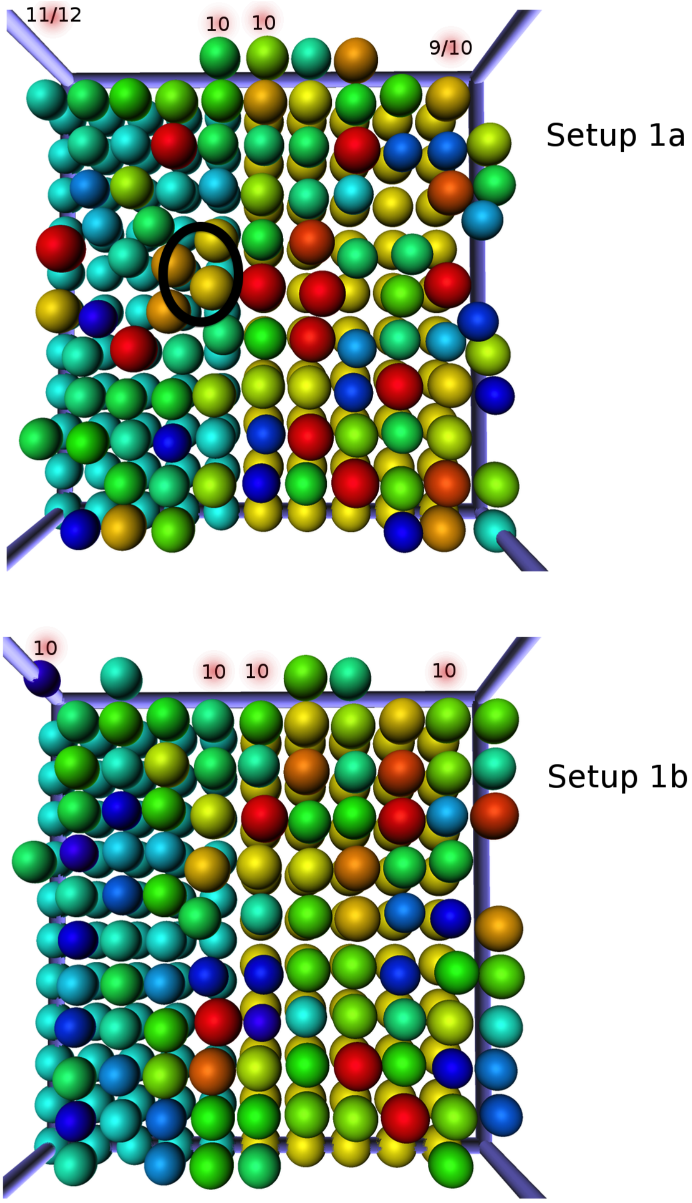

Figure 9.7: Snapshots from two independent trajectories using Setup 1, looking onto the second deposited crystal layer from the high-x side. The templates appear in the background, but the first deposited layer is excluded, as shown in FIG. 9.6. Colours indicate particle size, as in FIG. 9.5. The numbers shaded in red indicate the number of particles in the columns either side of a template boundary. The black ellipse in the top panel marks the area around the vacancy in the first deposited layer, as discussed in the text.

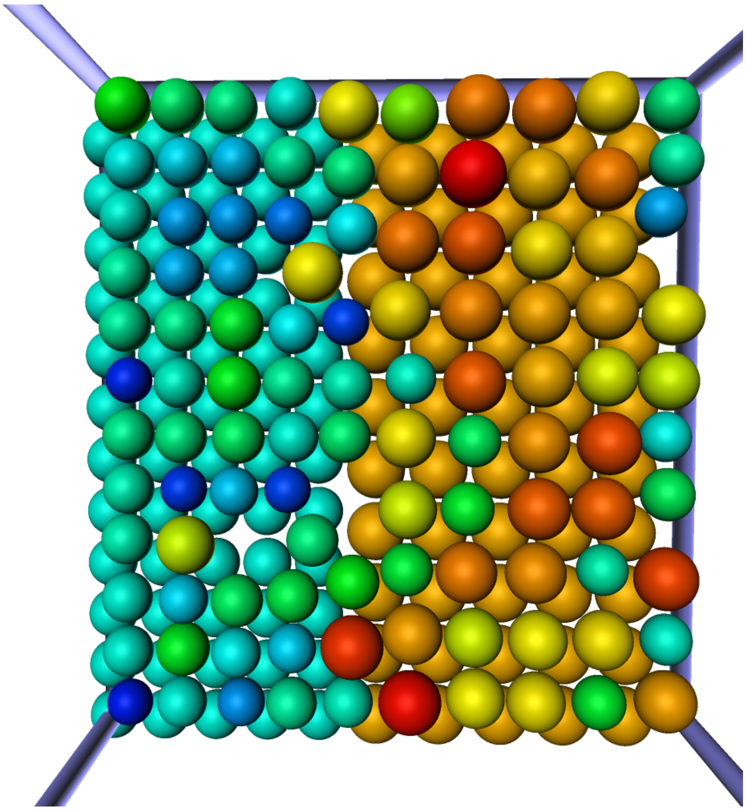

Figure 9.8: Snapshot of the first and only deposited layer in a sample trajectory from Setup 2, looking onto the layer from the high-x side, so that the templates appear in the background. Colours indicate particle size from $d \leq 0.82$ (blue) to $d \geq 1.18$ (red).

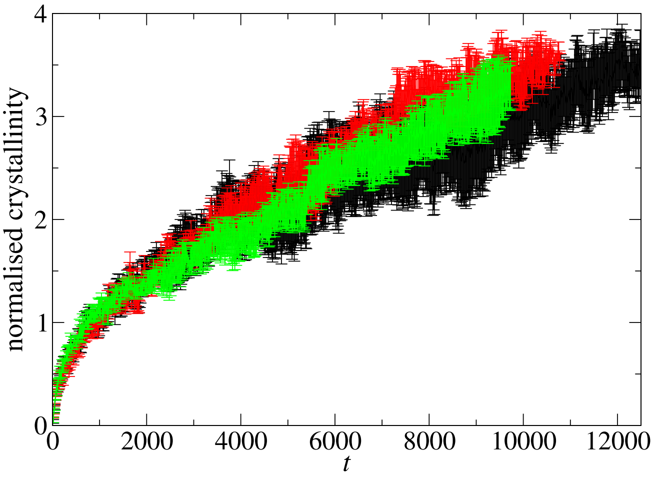

Figure 9.9: Crystal growth curves for the dual template Setup 1 parameters, with different y extents of the templates. In the black series, each template is 5 particles wide in y, red 10, green 15. The crystallinity is normalised by multiplication by N and division by $L_y L_z$, in order to enable comparison when the templates have different areas. The multiplication by N is not strictly necessary as all simulations on this plot contain $N = 5000$ particles.

9.4 Quantitative observations

From direct observation, discussed in the previous section, essentially no crystal growth was seen to take place in the simulations using Setup 2. For this reason, the quantitative analysis of the crystal is restricted to the lower-σ Setup 1 simulations.

9.4.1 Crystal growth rate

I now examine the crystal growth rate in Setup 1, under 3 different choices of the number of y lattice points: 5, 10 and 15. As discussed in Section 9.2.3, this is done in order to check for any finite-size dependence in these terms due to the mismatch along the boundary between the dual templates – if the width of each template is increased, then the ‘amount of crystal’ directly affected by the mismatch is proportionally smaller. When comparing crystallinity through time between simulations of different L_y, it is necessary to correct for differences between simulations in the interfacial area $L_y L_z$ – a crystal with greater interfacial area can add more particles to its surface per unit time just by virtue of the increased area, which is not the effect that is of interest here.

In FIG. 9.9, I therefore plot the crystallinity multiplied by a factor $N/(L_y L_z)$ as a function of time. It can be seen that, within error, there is no difference in crystallinity between the different template sizes, on the simulated timescale. There could perhaps be a cancelling-out between two effects behind this. On the one hand, when the templates are wider in y, more of the crystal is further away from the mismatched boundary (enhancing growth), but particles that are being fractionated onto one or the other template must diffuse further to reach their ‘correct’ template, slowing growth.

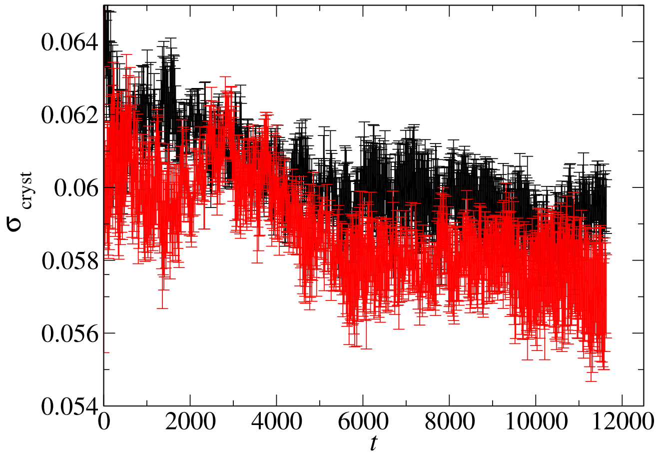

Figure 9.10: Crystal polydispersity through time for the A template (smaller particles, shown in black) and B template (larger particles, red) regions, where particles deposited onto each template have been identified based on their y coordinate. There is a measurable reduction from the parent $\sigma = 0.062$ through time, but the error bars are too large to distinguish between the individual templates. The data shown are for templates each 5 particles wide in y – the data for choices of 10 and 15 for this value are indistinguishable within error so are omitted for clarity.

9.4.2 Fractionation

I now investigate the composition of the crystals formed on each of the dual templates. The measurement of crystalline statistics is essentially the same as in Chapters 6 and 8. However now, as well as dividing the particles into those on either side of the crystal interface, they are divided into two groups based on their y coordinate, to distinguish those deposited onto one of the dual templates from those deposited onto the other. As earlier in this chapter, the templates will be referred to as A (composed of smaller particles with smaller lattice spacing) and B (larger particles).

Polydispersity

The behaviour of the crystal polydispersity σ_{cryst} is unremarkable, in light of the results in Chapter 6. As shown in FIG. 9.10, the crystals show as before a slight reduction in polydispersity as they grow. Within statistical error, the particles deposited on to both the A and B templates, for all choices of the y lattice width, show the same behaviour.

Layer-by-layer mean diameter

I now perform a spatial analysis of the distribution of particle size within the crystals. The results are shown in FIG. 9.11. Since the crystal 'interfaces' are now extremely spread-out due to the poor ordering and the apparent possibility of liquid coating (discussed in Section 9.3.1), I have indicated the maximum extent of crystal-flagged particles (as identified by the bond order parameter). It is clear that there is marked fractionation between the dual templates in the first few layers of crystal growth, which are more or less complete, but that, as the edge of the crystal

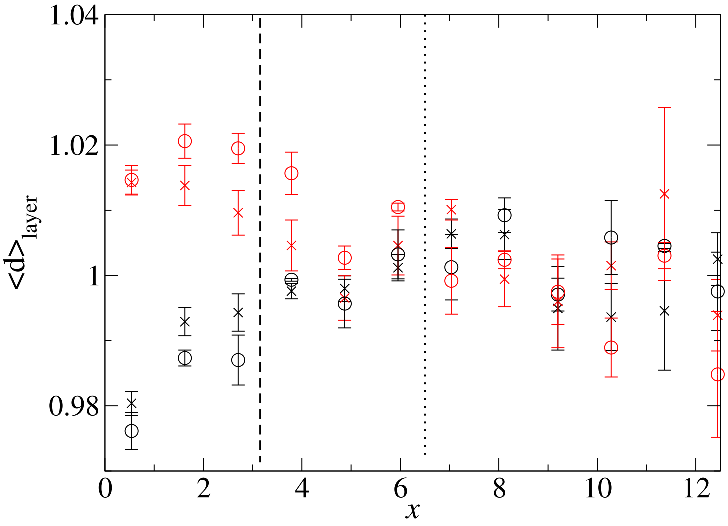

Figure 9.11: Layer-by-layer mean diameter in the low-x end of the simulation cell, with templates 5 particles wide in y (crosses) and 15 particles wide (circles). Red shows the region deposited onto the large particle B template, and black the small particle A template. The data points indicate the start of the bin for each layer. The dashed line indicates the minimum furthest extent (across each independent simulation) of crystalline ordering (as identified by the bond order parameter) for the cross data series, the dotted line for the circles. The data are taken from a late time in the simulation ($t \approx 14000$ for the crosses, $t \approx 11000$ for the circles). Data are averaged over 6 independent simulations for the crosses, and 3 for the circles.

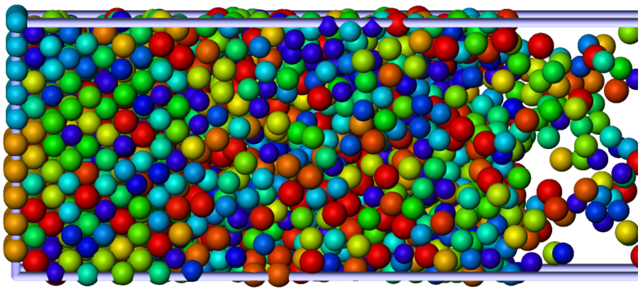

Figure 9.12: Snapshot of trajectory Setup 1b at the time $t \approx 14000$ when the corresponding layer-by-layer fractionation measurements in FIG. 9.11 were taken. The crystal now appears to be coated with a liquid region (discussed in the text).

is reached, the size preference becomes less strong so that the mean diameters in each region become equal within error.

Since those outer layers are typically rather incomplete and disordered, it may be that the crystal will be able to retain the strong fractionation of the first few layers in the outer layers also, once they are complete. However, the simulations would need to run for much longer to allow this to be confirmed. As it is, the data seem to indicate a gradual weakening of the templates' size preferences with increasing x, suggesting that the particle selectivity imposed by the template properties may have some finite 'persistence length'.

9.5 Conclusions

In this chapter I have investigated the effect of a somewhat unusual templating method on crystal growth in a polydisperse system. Choosing a parent size distribution made of two closely spaced component distributions, I customised a pair of dual templates to each of these component distributions respectively. The aim was to examine the resultant growth, particularly in terms of selectivity induced by the differing template properties.

As may have been expected, it appears that two narrow 'stripe'-like templates, each a few particle diameters wide, with mismatched lattice parameters, do not constitute a substrate onto which high-quality crystals can be grown. Any crystals which do form must have a lattice mismatch along the boundary between the templates, which leads to significant distortion and poor ordering of the crystal region. If each template were made many times wider than the particle diameter, the effect of this mismatch would be proportionally smaller, but this could also result in reduced ability to fractionate the particles since, having arrived at one template, a particle must then diffuse further away to reach the other template. In the regime of template widths studied here, there was no measurable benefit to growth rate in making the templates wider. However, the dual templates provided the opportunity for a direct look at some of the defect structure possible when (in a polydisperse system) two crystals of significantly different lattice parameter are brought side-by-side. For instance, it appears far more likely for the boundary between the templates to accommodate an effective vacancy defect for the lower-lattice parameter template, than to have an effective interstitial for the larger one.

The inter-particle potential was chosen so that, based on the results of Chapter 6, the crystal should be unable to locally coexist with the liquid and therefore be coated by a gas layer. Perhaps the most surprising result in this chapter, therefore, is that for some trajectories here the crystal seemed to become coated by a liquid layer instead. I argued that the most likely cause of this is the significant distortion of the crystal by the mismatched side-by-side templates, which could render the bulk crystal free energy considerations in Chapter 6 inapplicable. The exact reason is unclear, but I suggested that the crystal region (being extremely distorted by the presence of two mismatching crystals) may be unable to reach its 'true' bulk free energy minimum, effectively raising the crystal free energy minimum and allowing a common tangent to be drawn with the liquid minimum. Some amount of support comes from the fact that simulations with templates 15 particles (rather than 5) wide in the y direction, so that more interfacial area lies further from the distorting mismatched boundary, did not show liquid-coating of the crystal – in those cases, the *bulk* crystal free energy (from which the lack of a crystal-liquid common tangent was derived in Chapter 6) may be more applicable since the free energy cost of the mismatch is proportionally smaller. This would require much further work to verify, however.

Turning to the compositions of the crystals grown, I found that, as in Chapter 6 for similar parent polydispersity as here, the crystals slightly reduce their polydispersity as they grow. In terms of mean diameter, there was clear selectivity between the dual templates in the first few layers of crystal growth: each template seemed to some extent to be able to 'pick out' smaller or larger particles, according to its own properties. However, the crystals remained significantly polydisperse and, for Setup 1 (in which the component distributions of the parent overlap significantly), many larger particles still ended up being incorporated into the small-particle template, and vice versa. For the higher-polydispersity Setup 2, there was much clearer selectivity: a particle from one part of the (now truly bimodal) distribution was extremely unlikely to end up on the 'wrong' template. However, crystal growth was so slow here that only 1 or 2 layers formed on the simulated timescale. In both cases, the data were not sufficient to say for sure how far into the grown crystal the size preference of the template asserted itself – particularly as the outer crystal layers were typically incomplete – but seemed to be consistent with a decreasing strength of preference further away from the templates.

Overall, the results show that fine-tuning of an epitaxial substrate's properties can induce some selectivity of particles from a size-polydisperse bulk. However, the still-significant polydispersity of the crystals grown, along with visual inspection, show that they still remain quite impure in terms of their composition. When the parent size distribution was fully bimodal, so that the corresponding templates had very different lattice parameters, much stronger selectivity was observed, but crystal growth was too slow for any quantitative measurements. Longer simulations or, preferably, precisely controlled experiments on epitaxial selectivity in polydisperse systems, may help to fill the gaps in this respect. I have also illustrated in detail some effects on crystal structure of bringing two crystals of differing lattice parameter together, and have found that free energy considerations implying gas-coating of the crystals in Chapter 6 may fail to apply when, as here, extreme distortion in the form of mismatching dual templates is imposed on the crystal region.

Chapter 10

Results VII – Diffusion of concentration and particle size fields in a polydisperse hard sphere fluid

In this chapter, I study diffusive relaxation in a polydisperse hard sphere fluid within a framework devised by R. M. L. Evans, using the moment-based BMCSL free energy to find the relaxation eigenmodes of coupled fields describing local concentration and local particle size. The motivation is broadly similar to that of Ref. [Warren, 1999] – using a moment-based method to reduce the description of diffusion to a manageable number of variables (as opposed to having one for each of the infinite number of particle species), analogous to the moment free energy method of rendering polydisperse phase equilibrium calculations tractable [Sollich and Cates, 1998; Warren, 1998].

10.1 Introduction

Diffusion in colloidal fluids has been discussed in Section 2.3, and some simulation measurements specific to polydispersity were shown in Chapter 4. From a theoretical point of view, polydisperse diffusion is, like polydisperse phase equilibria, extremely complicated. In both cases, a naive extrapolation from the case of a finite mixture to the polydisperse case results in disaster (if the reader will permit some melodrama), since we are confronted with an infinite number of variables even to *describe* the system, let alone perform calculations. In the context of phase equilibria, I have discussed various means of overcoming this problem already – see Section 2.2.3. The most successful non-perturbative approach is the moment free energy method of Sollich and others [Sollich and Cates, 1998; Warren, 1998]. It takes advantage of the BMCSL free energy (a polydisperse generalisation of the Carnhan-Starling free energy), the excess part of which depends on only the first four moments of a polydisperse size distribution. By this means, the central stumbling block – an infinite number of conserved species variables – is overcome.

In this chapter, I introduce (from a starting point due to R. M. L. Evans) a theory of diffusive relaxation in a polydisperse fluid via an analogous approach: a description of diffusion in terms of moment densities rather than individual species densities. Using the BMCSL free energy and a perturbative framework, with hydrodynamics neglected, I find expressions for the coupled relaxation of fields describing local concentration and local (average) particle size. Determining the eigenmodes of relaxation, I find a mode of pure concentration diffusion, and a mixed mode which becomes purely 'size diffusion' at a particular value of the volume fraction ϕ. The work constitutes a relatively brief but concise illustration of the overall shape and possible benefits of such an approach to diffusion in polydisperse systems. Even irrespective of its role in phase transition kinetics (which have been the focus of most of this thesis), diffusive transport in the presence of polydispersity is an important problem, and one on which little theoretical work has so far been performed. A special exception is the work of Warren [Warren, 1999], which is similar in motivation to the present work but focuses on phase transitions, and is couched in terms of the Flory-Huggins model of polydisperse polymers.

10.2 Derivation

I now give the detailed derivation to find expressions for the relaxation of the coupled fields.

10.2.1 Species flux

The derivation begins with an expression for the flux of a particle size species labelled i – note that i does not now refer to a particular particle, but to a species. The species flux $\mathbf{J}_i$ is given by:

$$\mathbf{J}_i = -\sum_j M_{ij} \boldsymbol{\nabla} \mu_j \tag{10.1}$$

in which M_{ij} is a mobility matrix describing the flux of species i induced by gradients in local chemical potential μ_j of species j. The chemical potentials may be split into ideal and excess parts, like so:

$$\mu_j = \ln\rho_j + \mu_j^{\mathrm{ex}} \tag{10.2}$$

where ρ_j is a local species density and μ_j^{ex} depends implicitly on the local densities of *all* size species. Densities are assumed to be made dimensionless, by multiplication by the volume of a sphere of unit diameter.

Mobility matrix

For the mobility matrix M_{ij}, a full treatment would require knowledge of how each species responds to chemical potential gradients in all other species. The non-diagonal elements of such a matrix would be due to hydrodynamic interactions (HI) between the particles. For simplicity, I now make the approximation of ignoring HI. Such a step is often taken in simulations of soft matter (including those in this thesis). HI cannot ever be eliminated in a real system, but may be screened to a greater or lesser extent, as has already been discussed in Section 5.2. In the present context, such an approximation can be motivated from the point of view of exploring the consequences for diffusion of the BMCSL free energy on its own, so that its contribution can be isolated from that of HI. I therefore write that $M_{ij} = \delta_{ij}\rho_i D_i$, i.e. is diagonal, where D_i is a diffusion constant, which could be chosen to depend on concentration. For now, though, I choose instead to keep D_i as the ideal Stokes-Einstein diffusion coefficient D_0 for the particle size species i.

Excess chemical potentials

The BMCSL free energy gives an expression for a polydisperse fluid which depends only on the first four moment densities of the size distribution $\phi_0, \phi_1, \phi_2, \phi_3$. A moment density ϕ_n is defined as $\phi_n = \sum\limits_i d_i^n \rho_i$. The excess part of the BMCSL free energy density is:

$$f^{\mathrm{ex}} = \left(\frac{\phi_2^3}{\phi_3^2} - \phi_0\right)\ln(1-\phi_3) + \frac{3\phi_1\phi_2}{1-\phi_3} + \frac{\phi_2^3}{\phi_3(1-\phi_3)^2} \ . \tag{10.3}$$

The gradients of excess chemical potentials may then be written, using the chain rule, as:

$$\boldsymbol{\nabla}\mu_i^{\mathrm{ex}} = \sum_\alpha \left(\frac{\partial \mu_i^{\mathrm{ex}}}{\partial \phi_\alpha}\right)\boldsymbol{\nabla}\phi_\alpha \tag{10.4}$$

Substituting the above into Equation 10.1, the flux of species i is given by:

$$-\mathbf{J}_i = D_i \boldsymbol{\nabla}\rho_i + \rho_i \sum_\alpha \left(D_i \frac{\partial \mu_i^{\mathrm{ex}}}{\partial \phi_\alpha}\right)\boldsymbol{\nabla}\phi_\alpha \ . \tag{10.5}$$

10.2.2 Moment density dynamics

Now, using that $-\boldsymbol{\nabla}\cdot\mathbf{J}_i = \frac{\partial \rho_i}{\partial t}$ and the definition of the moment densities, the local rate of change of a moment density is:

$$\frac{\partial \phi_n}{\partial t} = \sum_i d_i^n D_i \nabla^2 \rho_i + \sum_i d_i^n \boldsymbol{\nabla}\cdot\left(\rho_i \sum_\alpha D_i \frac{\partial \mu_i^{\mathrm{ex}}}{\partial \phi_\alpha}\right)\boldsymbol{\nabla}\phi_\alpha \ . \tag{10.6}$$

I now isolate the particle size dependence of the diffusion coefficient, writing $D_i = s d_i^{-1}$. The chemical potential derivatives can be rewritten as:

$$\frac{\partial \mu_i^{\mathrm{ex}}}{\partial \phi_\alpha} = \sum_\beta d_i^\beta \frac{\partial^2 f^{\mathrm{ex}}}{\partial \phi_\beta \partial \phi_\alpha} \tag{10.7}$$

since $\mu_i^{\mathrm{ex}} = \sum\limits_\beta \frac{\partial f^{\mathrm{ex}}}{\partial \phi_\beta}\frac{\partial \phi_\beta}{\partial \rho_i}$ by the chain rule, and $\frac{\partial \phi_\beta}{\partial \rho_i} = d_i^\beta$ by definition. The derivates of the free energy may be more compactly written as a Hessian matrix $A_{\alpha\beta} = \frac{\partial^2 f^{\mathrm{ex}}}{\partial \phi_\beta \partial \phi_\alpha}$. We then arrive at:

$$\frac{\partial \phi_n}{\partial t} = s\nabla^2 \phi_{n-1} + s\boldsymbol{\nabla}\cdot\left(\sum_{\alpha\beta} \phi_{n-1+\beta} A_{\alpha\beta}\right)\boldsymbol{\nabla}\phi_\alpha \tag{10.8}$$

which relates the rate of change of one moment density to the spatial derivates of the others.

10.2.3 First order expansion

I now introduce again the variable ϵ, which is a particle's deviation from mean size in units of the mean, and is familiar from e.g. Chapter 5. I define a spatially-varying field of 'average local size deviation' ('local' in the same sense as the moment densities are local), $\langle\epsilon\rangle$. Assuming $\langle\epsilon\rangle$ to be small, to first order in this field, the moment densities can be related to one another via $\phi_n = \phi_3\,(1 + [n-3]\langle\epsilon\rangle)$. Therefore, one can work in terms of the concentration field (described by the third moment density ϕ_3 which is equivalent to the volume fraction). For simplicity, I write $\phi_3 = \phi$. Further, I expand to first order in concentration, writing $\phi = \widehat{\phi} + \delta\phi$, where $\delta\phi$ is the time- and space-varying portion and $\widehat{\phi}$ is a constant background. For consistency and clarity, I also rewrite the analogous perturbative size deviation fluctuations as $\langle\epsilon\rangle = \delta\epsilon$. The task is now to use Equation 10.8 to produce coupled equations for the dynamics of the concentration ($\delta\phi$) and size ($\delta\epsilon$) fields.

10.2.4 Dynamics of concentration field

Setting $n = 3$ in Equation 10.8, and with the above first order expansion to express all other moments in terms of $\delta\epsilon$ and $\phi \equiv \phi_3$, the dynamics of the concentration field are given by:

$$\frac{\partial\delta\phi}{\partial t} = s\nabla^2(\delta\phi - \widehat{\phi}\delta\epsilon) + s\widehat{\phi}\sum_{\alpha\beta} A_{\alpha\beta}\nabla^2(\delta\phi + \widehat{\phi}[\alpha-3]\delta\epsilon) \tag{10.9}$$

where terms of order $(\delta\epsilon)^2$, $(\delta\phi)^2$, $\delta\phi\delta\epsilon$, $\delta\phi\boldsymbol{\nabla}\delta\phi$, etc. have been dropped.

10.2.5 Dynamics of size field

To find the dynamics of the $\delta\epsilon$ from Equation 10.8, I use that, to first order, $\delta\epsilon \equiv \langle\epsilon\rangle = \frac{\phi_1}{\phi_0} - 1$. (This can be shown by writing both ϕ_0 and ϕ_1 in terms of ϕ_3). Then:

$$\frac{\partial\delta\epsilon}{\partial t} = -\frac{\phi_1}{\phi_0^2}\frac{\partial\phi_0}{\partial t} + \frac{1}{\phi_0}\frac{\partial\phi_1}{\partial t}\ . \tag{10.10}$$

Finding the $\frac{\partial\phi_n}{\partial t}$ from Equation 10.8, and throwing away terms higher than first order, the dynamics of the size field is then given by:

$$\frac{\partial\delta\epsilon}{\partial t} = s\nabla^2\delta\epsilon\ . \tag{10.11}$$

As expected, a region of positive local $\delta\epsilon$ will relax towards $\delta\epsilon = 0$, however, to first order, there is no dependence on concentration in this relaxation. Note that a dependence would be introduced if s were chosen to depend on concentration.

The Hessian matrix $A_{\alpha\beta}$

At all stages in the present derivation, the Hessian matrix $A_{\alpha\beta}$ appears in between two $\boldsymbol{\nabla}$ operators (see e.g. Equation 10.9). When the first order approximation is made, the rightmost $\boldsymbol{\nabla}$ picks out only terms of order δ, since these are the only terms allowed to vary spatially. Therefore the $\delta\phi$ and $\delta\epsilon$ dependence of $A_{\alpha\beta}$ disappears, and we are left with only the constant part.

10.3 Eigenmode analysis

10.3.1 Fourier transforming

To find the eigenmodes of the system, I first transform into Fourier space. The perturbations become functions of wavenumber q. Noting that all terms in Equations 10.9 and 10.11 are linear in s and contain only ∇^2 spatial derivatives, I now express the dynamics in terms of $\tau \equiv sq^2 t$. In Fourier space, the coupled equations become:

$$\frac{\partial \delta\phi_q}{\partial \tau} = \widehat{\phi}\delta\epsilon_q - \delta\phi_q - \sum_{\alpha\beta} A_{\alpha\beta}\widehat{\phi}(\delta\phi_q + \widehat{\phi}[\alpha - 3]\delta\epsilon_q) \tag{10.12}$$

$$\frac{\partial \delta\epsilon_q}{\partial \tau} = -\delta\epsilon_q \tag{10.13}$$

The coupled equations can be written vectorially as:

$$\frac{\partial}{\partial \tau}\begin{pmatrix}\delta\phi_q \\ \delta\epsilon_q\end{pmatrix} = \begin{pmatrix} -1 - \sum\limits_{\alpha\beta} A_{\alpha\beta}\widehat{\phi} & \widehat{\phi} - \sum\limits_{\alpha\beta} A_{\alpha\beta}\widehat{\phi}^2(\alpha - 3) \\ 0 & -1 \end{pmatrix}\begin{pmatrix}\delta\phi_q \\ \delta\epsilon_q\end{pmatrix} . \tag{10.14}$$

The matrix in the above equation (denoted $\boldsymbol{B}$) has normalised eigenvectors $\boldsymbol{e}_1$ and $\boldsymbol{e}_2$, given by:

$$\boldsymbol{e}_1 = \frac{1}{\sqrt{1+y_1^2}}\begin{pmatrix}1 \\ y_1\end{pmatrix} \; , \; \boldsymbol{e}_2 = \frac{1}{\sqrt{1+y_2^2}}\begin{pmatrix}1 \\ y_2\end{pmatrix} \tag{10.15}$$

with:

$$y_1 = \frac{\sum\limits_{\alpha\beta} A_{\alpha\beta}\widehat{\phi}}{\widehat{\phi} - \widehat{\phi}^2 \sum\limits_{\alpha\beta} A_{\alpha\beta}(\alpha - 3)} \; , \; y_2 = 0 \tag{10.16}$$

and associated eigenvalues such that $\boldsymbol{B} \cdot \boldsymbol{e}_n = \lambda_n \boldsymbol{e}_n$, given by:

$$\lambda_1 = -1 \; , \; \lambda_2 = -1 - \sum_{\alpha\beta} A_{\alpha\beta}\widehat{\phi} \; . \tag{10.17}$$

The values of $-\lambda_n$ give effective diffusion coefficients for the relaxation of each of the eigenmodes. In the eigenvectors $\boldsymbol{e}_n$, the top element describes how much of the associated eigenmode is $\delta\phi$-like, and the bottom how much is $\delta\epsilon$-like. Note that the eigenvectors and eigenvalues are functions of just one variable, the background (i.e. constant) concentration $\widehat{\phi}$ against which the fluctuations $\delta\phi$ and $\delta\epsilon$ are taking place.

It can be seen from the above that $\boldsymbol{e}_1$ corresponds to a 'mixed mode' of local concentration and size which relaxes with a constant diffusion coefficient, while $\boldsymbol{e}_2$ represents a purely concentration-like mode. This is expected, since Equation 10.13 tells us that a perturbation in $\delta\phi$ *only* does not excite fluctuations in $\delta\epsilon$.

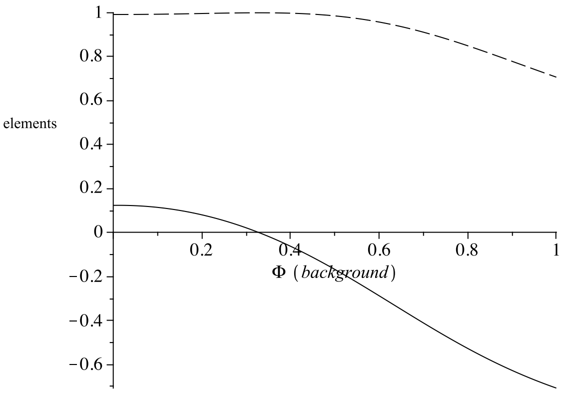

Figure 10.1: The normalised elements of the mixed eigenvector $\boldsymbol{e}_1$ as functions of background concentration $\widehat{\phi}$. The solid line shows the top element (representing concentration) given by $1/\sqrt{1+y_1^2}$, while the dashed line shows the bottom element (representing local size) $y_1/\sqrt{1+y_1^2}$.

10.3.2 Dependence on background concentration

In FIG. 10.1, I plot the elements of the mixed eigenvector $\boldsymbol{e}_1$ as functions of $\widehat{\phi}$. It can be seen that the mode is mostly dominated by its size-like component and, at a critical concentration $\widehat{\phi} \approx 0.33$, the mode represents purely size diffusion. Beyond that, the elements have opposite sign to each other, i.e. the mode now represents perturbations towards higher local concentration accompanied by smaller local size (or vice versa). Note that I have shown the behaviour up to $\widehat{\phi} = 1$ – as discussed for the Carnahan-Starling free energy in Chapter 7, the BMCSL free energy has a divergence (i.e. what should correspond to random close packing) at $\phi_3 = 1$. Therefore, $\widehat{\phi} = 1$ marks the point above which, *within the BMCSL free energy*, the fluid cannot be packed any more densely.

It is worth noting in passing that, for the colloid-polymer mixtures in Chapter 7, a significant improvement in the regime of short attraction range was obtained by replacing the Carnahan-Starling hard sphere free energy with one that had a divergence at the correct value $\phi_{\text{RCP}} \approx 0.64$. Therefore, in cases where the BMCSL form is used as a polydisperse hard sphere reference system for e.g. the perturbative inclusion of an additional attractive potential, its divergence at $\phi_3 = 1$ may also cause problems in the same way that the Carnahan-Starling form did in Chapter 7.

FIG. 10.2 shows the behaviour of the non-constant diffusion coefficient $-\lambda_2$. It increases with background concentration, showing that $\boldsymbol{e}_2$ represents a mode of collective diffusion in which the relaxation rate of local concentration perturbations increases with the background concentration.

10.3.3 Real space

Transforming back to real space, with the normalised eigenvectors in hand, it is now possible to write any vectorial perturbation to the concentration and size fields as a linear sum of the

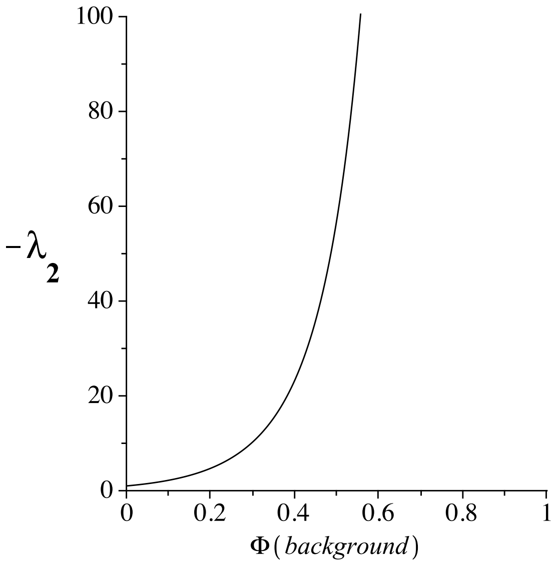

Figure 10.2: Behaviour of the eigenvalue λ_2, which is the negative diffusion coefficient of the eigenmode corresponding to the pure concentration eigenvector $\boldsymbol{e}_2$, as a function of background concentration $\widehat{\phi}$.

eigenvectors, like so:

$$\begin{pmatrix} \delta\phi \\ \delta\epsilon \end{pmatrix} = a\boldsymbol{e}_1 + b\boldsymbol{e}_2 \tag{10.18}$$

where the coefficients a and b obey:

$$\frac{\partial a}{\partial t} = -s\lambda_1 \nabla^2 a \ , \ \frac{\partial b}{\partial t} = -s\lambda_2 \nabla^2 b \ . \tag{10.19}$$

With the above, the interdependent relaxation of local concentration and local particle size may be solved by decomposition into two independent diffusion equations.

10.4 Conclusions

In this chapter, I have briefly introduced a method, within a first order framework and without hydrodynamics, of modelling the interdependent diffusive relaxation of fields describing fluctuations in local concentration and local particle size in a polydisperse hard sphere fluid. This involved treating the diffusion of moment densities (through the truncatable BMCSL free energy, which depends on only four moments), rather than individual particle size species, in order to get around the problem of the infinite number of density variables required to describe a polydisperse system.

For simplicity, I made a number of assumptions. Probably most importantly, hydrodynamic interactions (HI) were neglected, by assuming the mobility matrix M_{ij} to be diagonal. Despite this, the results are valuable in that they show the consequences for diffusion of the underlying chemical potential landscape of the system (assumed here to be given by the BMCSL free energy). Further, HI may be screened to some extent, depending on the details of the real system under study, so that results found in their absence constitute at least an interesting baseline case.

The calculations were restricted to first order in local concentration and size. This approximation was introduced in Section 10.2.3, in which all moment densities were related to the third via the perturbative variable $\langle\epsilon\rangle$, describing local average particle size. Note that, previous to that, Equation 10.8 is 'exact' inasmuch as all the independent moment densities are treated fully – a higher-order treatment could thus begin from this point.

I used the BMCSL free energy, which is the best accepted model for a polydisperse hard sphere fluid but, like the Carnahan-Starling form on which it is based, has a divergence at an unphysically high volume fraction $\phi_3 = 1$. The implications of this aspect of the Carnahan-Starling form for phase coexistences were discussed in Chapter 7. However, in the present case in which we are concerned with a homogeneous system and all the physics is somehow 'contained within' the BMCSL free energy (i.e. there is no polymer free energy term or other free energy branch involved), it can be expected that the results are qualitatively reasonable. Nevertheless, an alternative free energy (as one becomes available – note that the WKV3 form in Chapter 7 was for a *monodisperse* system) could be substituted into the theory easily by incorporation into the Hessian matrix $A_{\alpha\beta}$.

The first order expansion enabled the system to be instructively described in terms of only two fields, describing local concentration and particle size. I calculated the dynamics of these fields, and isolated their eigenmodes and associated eigenvalues. One eigenmode describes a collective mode of pure concentration diffusion, with a diffusion coefficient increasing strongly with background concentration $\widehat{\phi}$. The other eigenmode is mixed, but is dominated by its particle size component. Its diffusion coefficient was constant, at a value less than that of the pure concentration mode for all $\widehat{\phi} > 0$. This mode corresponds mostly to the self diffusion-like rearrangement by which changes in mean local particle size must take place. It is as intuitively expected that, for finite $\widehat{\phi}$, the relaxation of concentration fluctuations is much quicker than that of those in local particle size. At a certain value of $\widehat{\phi} \approx 0.33$ (which would be subject to quantitative change if the BMCSL free energy were replaced), the latter mode becomes completely size-like, so that local particle size and local concentration relax independently. Beyond this, towards the 'random close packing' point of the BMCSL free energy at $\widehat{\phi} = 1$, the concentration element becomes larger in magnitude than the size one, and now has the opposite sign.

The findings are interesting in their own right, but more value lies in their demonstrating (like the previous work in Ref. [Warren, 1999]) how a moment-based approach to polydisperse diffusion can – as well as providing clearer insight into the behaviour of the system – provide a manageable and systematic way of incorporating polydispersity into some wider dynamical theory. For instance, the theory tested in Chapter 8 [Evans and Holmes, 2001] relied on forcing the (infinite) diffusion matrix to be diagonal, so that each particle species responded only to concentration gradients in its own species. Contrastingly, the present theory (although the *mobility* matrix is assumed diagonal) incorporates the response of species i to concentration gradients in species j through j's effect on i's chemical potential, as specified by the BMCSL free energy. Couching the insights of [Evans and Holmes, 2001] in terms of moment diffusion would enable progress to be made without as drastic a diffusion approximation as is used there – this would provide some interesting further work and a natural application of the theory which has begun to be explored here. The results in Chapter 8 showed definitively that equilibrium crystal fractionation may be very far from being achieved in a dynamical system, and that dynamically-induced fractionation may therefore be more important for real applications (although this latter point could not be confirmed or refuted in that chapter). A simple but well-grounded way of treating polydisperse diffusion, as here, could have great utility as a part of future dynamical theories of polydisperse phase transitions and aggregation processes.

Chapter 11

Conclusions

11.1 Conclusions

This work has been concerned with understanding the kinetics of soft matter phase transitions through the study of model colloidal systems, particularly where size-polydispersity is present. It can therefore be thought of as an attempt to bring together two of the 'usual suspects' complicating soft matter physics – phase transitions kinetics and polydispersity – and help towards a more practical understanding of their effects. This is important, given that our present understanding of polydispersity is predominantly equilibrium-based and therefore often provides somewhat indirect insight into the behaviour of real systems (which are often not at equilibrium and, indeed, seem *less likely to be so* when even mild polydispersity is present [Liddle *et al.*, 2011]). The results presented here indicate that further work in this direction is both necessary and feasible.

After providing a basic introduction to the thermodynamics and dynamics of colloidal systems (Chapter 2), I took calibration results from the simulation code (introduced in Chapter 3) in order to illustrate key concepts and check agreement with existing work (Chapter 4). The primary purpose of these results was to lay the ground for later chapters, but they served a secondary purpose: anecdotally, I am aware of relatively fewer Kinetic Monte Carlo studies of colloidal phase behaviour and dynamics than studies involving Molecular Dynamics, Brownian Dynamics, etc. Therefore, the realistic physics displayed in the calibration results, along with the results of later chapters, serve to give the reader some insight into where this conceptually and computationally simple approach might be usefully employed.

The key theme running through most later chapters was fractionation of particle size between separating phases. Chapter 5 showed that fractionation may take place in the very early stages of gas-liquid separation, in the same direction but with a smaller magnitude than at equilibrium. That direction was found to depend qualitatively on a trivial-seeming detail of the inter-particle interaction, which was explained using Evans' perturbative theory of polydispersity. In accordance with that theory, the dependence of mean diameter fractionation on parent variance and of variance fractionation on parent skew remained approximately linear up to $\sigma \approx 0.4$.

Chapter 6 concerned crystal growth in the presence of gas-liquid separation. Fractionation was key again, being enhanced if the crystal was coated by a gas layer, and thereby reversing the effect of that gas layer on growth from the monodisperse case. The observed phase transition kinetics agreed with equilibrium and dynamical theory, and should be rather general, since the underlying free energy landscape responsible for the gas layer could occur in a variety of systems with relatively short attraction range. Local size ordering, in addition to overall fractionation, in bulk polydisperse phases, was observed for the first time and, importantly, was present even where polydispersity was so low that no *overall* fractionation was observed (Chapter 8). The local ordering was qualitatively explained here, but it would be fascinating to study this in a more quantitative way for a variety of polydispersities and crystal densities, perhaps with input from e.g. density functional theory, in order to determine and control the effects of polydispersity on crystal structuring.

In Chapter 7, I explored gas-liquid fractionation in an experimental colloid-polymer system, at times much later than those in Chapter 5. The fractionation here would be expected to be equilibrium-like in its magnitude, since the phases have had time to ripen in terms of their compositions. Within Evans' perturbative theory, applied to the Mean-Field Asakura-Oosawa (MFAO) free energy, qualitative fractionation agreement with some emerging experimental data was found. This relied on a modification to the MFAO theory, replacing the Carnahan-Starling (C-S) hard sphere free energy (which has a divergence at $\phi = 1$, not the expected $\phi_{RCP} \approx 0.64$) and giving rise to the newly-christened 'MFAO+WKV3' free energy. The fact that the MFAO free energy can be so markedly improved for small polymer size – just by replacing the C-S expression – is relevant to the moment free energy approach to polydispersity as well, and indeed to any theory of phase behaviour which takes a colloid-polymer free energy as input.

I studied, in Chapter 8, a prediction of fractionation induced not by a thermodynamic driving force, but by diffusive particle transport towards a hard sphere crystal. A fully diffusion-limited growth regime, required to properly test the prediction, was not reached on the simulated timescale. However, I did find that the *equilibrium* fractionation, against which the diffusive effect should take place, was never achieved either. The results suggest, within extremely small error bars, that a hard sphere system of polydispersity $\sigma = 0.03$ may grow a crystal with essentially no thermodynamically-driven fractionation at all.

Chapter 9 involved a speculative templating strategy, simulating the effects of an epitaxial substrate composed of two slightly different and precisely controlled templates, in order to see e.g. how much the templates induce fractionation. Some selectivity was induced, but it was unclear from the data how far this might persist into the crystals formed. The bulk free energy considerations from Chapter 6, which predict a gas coating of the crystal, seemed not to apply due to significant distortion of the crystal region by the dual templates.

Finally, in Chapter 10, I theoretically studied diffusion in a polydisperse hard sphere fluid. Use of the moment-based BMCSL free energy and a perturbative expansion allowed significant physical insight to be retained while using only two variables to describe the system: local concentration and local particle size. The main value of this work was in showing how polydispersity might be incorporated into diffusion, systematically and using a small number of physically meaningful variables. This could be useful for future theories of polydisperse phase transitions in which dynamics are included.

In summary, I have made a number of novel observations which complement and extend our present, mostly equilibrium, understanding of polydispersity. Importantly, polydispersities which might normally be considered 'near-monodisperse' were found to be associated with qualitative effects, e.g. a reversal of the role of metastability in crystal growth, or a reversal of fractionation direction depending on the choice between inter-particle potentials which are *only distinct by virtue* of the presence of polydispersity. Taken along with the existing literature, it becomes increasingly apparent that ignoring polydispersity can mean failure to understand or predict its dramatic and often qualitative consequences for a system's behaviour, even if the polydispersity is so small that it is difficult to measure or characterise experimentally. The overall message of the present work is the extent to which a dynamical approach – dynamical simulations and theory combined with a nuanced attitude to the equilibrium phase diagram (viz. whether/how it will be reached) – yields insights into polydispersity's role in the phase transition kinetics of real systems.

11.2 Limitations and future work

I have stated already the importance of future work along the general lines of this thesis, i.e. considering polydispersity's effects on phase transitions from both dynamical and equilibrium standpoints. Here I will briefly suggest some specific possibilities for future work, some prompted by limitations of the approaches I have used.

The neglect of hydrodynamics here is not uncommon in comparable dynamical simulations, but is important (unlike in equilibrium simulations where hydrodynamics are strictly irrelevant). It would therefore be valuable, and in line with the aim of more closely representing real systems, to perform simulations incorporating hydrodynamics while keeping them large and long-running enough to measure e.g. fractionation effects to the desired precision. This is not a trivial task, but should become ever more feasible with time. Nonetheless, simulations without hydrodynamics will remain important. For one thing, they enable structural and hydrodynamical effects to be disentangled from one another. Also, their much lighter computational load enables parameter space to be explored more easily, highlighting phenomena (such as those found in this thesis) which would be interesting to study with hydrodynamics included.

The free energies used e.g. in Chapters 5, 7 and 8 to predict fractionation, and in Chapter 6 to predict phase transition kinetics, are not exact. This means that some quantitative features of the results, e.g. exactly how far off equilibrium the fractionation in Chapter 5 is, could be subject to change if more precise expressions are used in the future. I note, however, that the qualitative predictions of the free energies used were generally well supported by simulation results, and that more complex expressions may bring their own problems, e.g. in differentiating them in order to predict fractionation. Further, the expressions used represent uncontroversial choices, being widely studied and characterised in terms of both their usefulness and shortcomings. In particular, the Carnahan-Starling free energy was adequate for the purposes of Chapters 5 and 6, but its quantitatively poor behaviour in Chapter 7 should prompt caution in any future, more quantitative, studies of the physics displayed in those chapters.

The 'boiled-egg' crystal growth in Chapter 6 could be the subject of much future work, since the free energy landscape precluding crystal-liquid coexistence could be present in many other systems with short range attractions (e.g. preliminary calculations for the colloid-polymer mixture in Chapter 7 appear to show such a landscape). It could perhaps be directly observed in experiment using confocal video microscopy. The truly monodisperse case in which metastable G-L separation was found to slow down growth (although it enhances nucleation) could be tested in experiments on protein crystallisation. Experiments in which a single crystal is studied (e.g. by the introduction of a template, as in my simulations), would be particularly valuable, since they would allow growth to be studied independently of nucleation. A key complicating feature of experiments, without perfect colloid-solvent density matching, is that there would be a buoyancy force acting to pull the gas layer away from the crystal. Numerical modelling incorporating gravity would therefore be extremely helpful in determining how the effects of the gas layer are further influenced by this factor.

Chapter 7 highlights the need for more experimental work in which polydisperse systems and phases are characterised precisely, along the lines of Ref. [Liddle *et al.*, 2011] and forthcoming work on the same system. This would enable more detailed testing of theory, and in particular would highlight instances where nonequilibrium effects are present. The finding that the MFAO model of colloid-polymer mixtures can be improved significantly at small size ratios just by replacing the Carnahan-Starling free energy should prompt a reevaluation of instances where the MFAO model has been found lacking, in order to determine to what extent observed inaccuracies are due to the use of the Carnahan-Starling form.

Concerning Chapter 8, larger simulations could be performed to try and reach the fully diffusion-limited regime, which is important in experiment but, in simulation, seems as though it could take a very long time to access. The fact that equilibrium fractionation was not achieved at any stage means it would be worthwhile to visit the theoretical prediction again without requiring local equilibration. Also, that theory may be a good first application for the moment-based approach to polydisperse diffusion described in Chapter 10. Both these considerations apply not just to the prediction in Ref. [Evans and Holmes, 2001], but to any future theory that seeks to incorporate polydisperse dynamics into an aggregation process.

The above is certainly not exhaustive, but serves to give some idea of the variety of future work which, on the basis of the findings in this thesis, should prove interesting and valuable.

Appendix A

Code samples

In this appendix I present code samples to illustrate the core functionality and concepts of the main simulation code. I also present some examples of analysis code, to illustrate the general approach to reading and manipulating trajectory information.

The entire code repository is available at:

`http://bitbucket.org/johnjosephwilliamson/phd-code-clone-bitbucket`

The repository contains a README file which describes the folders in the top-level directory, as well as a short MANUAL file detailing compilation of the code and the running of a simple simulation.

In the samples presented here, I have tried to contextualise each extract and include enough comments to make the meaning of variables and outside functions reasonably clear. Any directory paths given in this chapter are relative to the top-level directory of the repository. Where integers are used for flagging truth values, I have used the convention 1 = true and 0 = false.

A.1 Core simulation code

The core simulation code consists of a main file which is 'wrapped' by either a console-based or graphical wrapper which includes it and calls its functions. The file is called:

`newsimulation.cpp`

It exists in a number of versions for e.g. scalable or non-scalable dependence of the square well range upon particle size (see Chapter 5). In this section, I present extracts from the scalable square well code, contained in:

`SQWell_Simulation/ScalableCompilation/newsimulation.cpp`

A.1.1 'Particle' data structure

Each particle is represented by an instance of the 'particle' data structure, which groups together the various properties a particle has.

```
// Structure to represent a 'particle' and its various properties.
typedef struct {
  vector position;
  double radius;
  double seedradius; // A dimensionless seed radius used to 'grow' particles up to the
      required radius.
  double r0sq; // Used in the crystal checking algorithm.
  double delta; // The individual Stokes-Einstein stepsize parameter which is size-dependent.
  int xcell; // Indices of the subcell (or 'zone') the particle is in, for linked-cell list
      method.
  int ycell;
  int zcell;
  int xcross, ycross, zcross; // Count periodic boundary crosses for use in postprocessing (
      MSD etc).
  int cellindex; // Where is the particle in its subcell's occupants list.
  int mobile; // For immobilising the particle if required.
  int crystal; // Flagging as crystalline.
  int initlattice; // Is a member of the initial template lattice?
  int overlap; // Does the particle have any overlaps?
  double overlapvol;

 } particle;
```

A.1.2 Particle interaction algorithm

This part of the code operates when a stochastic trial particle move has been proposed. It either rejects the move, if a collision of hard particle cores occurs, or returns a difference in the number of square wells occupied, so that the calling function can accept the move probabilistically based on the well depth. The essence of the Kinetic Monte Carlo algorithm is that both the particle to move and the magnitude/direction of the move are chosen using a random number generator, with the proviso that the moves are small enough to represent the Brownian 'kicks' of particles by the solvent. The dependence of acceptance probability on the change in the number of wells occupied models the effect of the attractive square wells. See Chapter 3 for further details.

```
// Sepcheck. Takes a particle and its label i as input, and operates on the 'collision'
// variable to tell the calling function whether the particle can move,
// how many square wells are escaped/entered, etc.
inline int sepcheck_2(const particle &subject, int i, int &collision) {
  int xcell, ycell, zcell;
  int howmany, a, b, c, h, j;
  double sepsq_current,sepsq_new, sepsq_well, sumradius;

  // If NOINTERACT is switched on, ignore collisions completely:
  if (NOINTERACT) {
    return 0;
  }

  // Zero counter for the number of square well the particle is involved in.
  int delwells = 0;
  subcell interest_cell;

  // Search through adjacent zones (subcells of the sim box) which could contain interacting
      particles.
  for (a = -2; a < 3; a++) {
    xcell = periodiccellx(subject.xcell + a);
    for (b = -2; b < 3; b++) {
      ycell = periodiccelly(subject.ycell + b);
```

```
   for (c = -2; c < 3; c++) {
    zcell = periodiccellz(subject.zcell + c);

    interest_cell = cell[zcell + ZNUM * (ycell + YNUM * xcell)];
    howmany = interest_cell.occupants;
    // Go through the subcell's list of occupants checking for interacters.
    for (h = 1; h <= howmany; h++) {
     j = interest_cell.particle[h];
     // Don't compare a particle with itself!
     if (i == j)
      continue;
     sepsq_new = sepsquared(subject.position, point[j].position);
     // Check for collision:
     sumradius = subject.radius + point[j].radius;
     if ((sepsq_new < sumradius * sumradius)) {
      collision = 1;
      return 0; // This position would cause a collision, so return.
     }
     // If no collision, check whether point[i] is
     // curently within j's well, in order to
     // be able to give a value of delwells.
     sepsq_current = sepsquared(point[i].position, point[j].position);
     sepsq_well = sumradius * sumradius * RANGE_PLUS_ONE_SQUARED;
     // 25 July 2011: This version now uses scalable version of definition:
     // 2 Oct 2011: speed up by using precalculated range+1 squared and unbracketing.
     if (sepsq_new < sepsq_well) {
      if (sepsq_current > sepsq_well) {
       ++delwells;
      }
      continue;
     }
     if (sepsq_new > sepsq_well) {
      if (sepsq_current < sepsq_well) {
       --delwells;
      }
     }
    }
   }
  }
 }
  // This sets x=0 as a hard wall so that in highly polydisperse cases, particles intruding
  // through the template encounter a hard flat wall.
 // If FLATRIGHTWALL = 1, we want to treat the RHS of the 'bumpy wall' of
 // immobilised particles as flat. So, set x=0 as a hard wall.
 if ((subject.position.x < subject.radius)) {
  if (FLATRIGHTWALL - EQUILSTAGE == 1) {
      collision = 1;
      return 0; // This position would cause a collision, so return
    }

  }
 // If we've got this far without returning, this subject particle has
 // no collisions. Return change in well number.
 return delwells;
}
```

Having determined the change in the number of square wells occupied, the acceptance probability for the move is determined and implemented as shown in the following extract from the function:

`particlemove(...)`

A considerable advantage of the step-like nature of the square well potential is that the energy contribution from each well is a known constant, which makes the code significantly faster than if energy differences have to be calculated each time based on the exact separation of particles.

```
// Normal case, no hardcore overlaps exist.
  // March 2012 -- *** This now uses sepcheck2() function instead.
    if (type == 1 && point[i].overlap == 0) {
        canmove = 0;
    collision = 0;
        // Check the new position for collisions and change in number of wells.
        // NB delwells is being sign-changed here because the code below to do with probability
        // is sign-inverted.
        // Historical quirk.
    delwells = (-1)*sepcheck_2(newpoint, i, collision);

    // If there is a collision, don't update point[i].

    if (collision)
      return 0;


    // Accept with some probability, depending on change in wells occupied.
        double pfac = 1;
    if (delwells < 0) {
      // don't want to have to divide in this loop. store divisor in tmp_pfac instead.
      double tmp_pfac = 1;
      for (k = 0; k > delwells; k--) {
        // BOLTZ is the appropriate exponential, pre-calculated to save time.
        tmp_pfac *= BOLTZ;
      }
      pfac /= tmp_pfac;
    }
    if (delwells > 0) {
      for (k = 0; k < delwells; k++) {
        pfac *= BOLTZ;
      }
    }
    canmove = (uni() < (1.0 / (1.0 + pfac)));
    }
```

A.1.3 Crystal template creation

Where a crystal template is required, particles are taken from the left (low-x) end of the simulation cell and placed at $x = 0$ in the appropriate face of the FCC lattice. Further details appear in Chapter 3. Shown below is the function in the core simulation code which handles this process:

`makelattice(...)`

```
// Function to create a lattice at the left-side of the simulation cell.
// Argument determines whether the latticed particles are immobilised or not.
// To make the initial conditions around the lattice as realistic as possible,
// the lattice is placed at an x-coordinate equal to the mean x-coord
// of all the particles that were removed from the bulk to create the lattice.
// Equivalently, move all non-lattice particles left (decreasing x) by that distance.
int makelattice(int immobile) {
  int i, j, k, l, m;
  double xlatsep, minx, meanx; // Lattice point separation in x direction.
```

```
int latticen;
double meanlatticex;
int count, onlattice;
vector newpoint;
// Lattice point typedef.
typedef struct {
  int inhabited;
  vector position;
} latticepoint;

// Array of lattice points:
latticepoint lattice[MAXLATTICE][MAXLATTICE][MAXLATTICE];
latticen = XLATTICE*YLATTICE*ZLATTICE;
if (latticen == 0) {
  return 0;
}
// First initialise the lattice points:
xlatsep = sqrt(LATSEP*LATSEP*0.5);
for (i = 0; i < XLATTICE; i++) {
  for (j = 0; j < YLATTICE; j++) {
    for (k = 0; k < ZLATTICE; k++) {
      lattice[i][j][k].inhabited = 0;
      lattice[i][j][k].position.x = i * xlatsep;
      if (i == 0 || i%2 == 0) {
        lattice[i][j][k].position.y = j * LATSEP;
        lattice[i][j][k].position.z = k * LATSEP;
      } else {
        lattice[i][j][k].position.y = (j + 0.5) * LATSEP;
        lattice[i][j][k].position.z = (k + 0.5) * LATSEP;
      }
      // Periodic boundaries:
      lattice[i][j][k].position.x = periodicx1(lattice[i][j][k].position.x);
      lattice[i][j][k].position.y = periodicy1(lattice[i][j][k].position.y);
      lattice[i][j][k].position.z = periodicz1(lattice[i][j][k].position.z);
    }
  }
}
// Lattice points are now initialised.
// Now place particles on the lattice, always taking the leftmost of the
// particles (i.e.\ those closest to the 'nucleation site').
meanx = 0;
for (count = 0; count < latticen; count++) {
  minx = LX; // i.e.\ large.
  // Find the leftmost particle:
  for (l = 0; l < N; l++) {
    if (point[l].initlattice == 0) {
      if (point[l].position.x < minx) {
        minx = point[l].position.x;
        m = l;
      }
    }
  }
  onlattice = 0;
  // Set newpoint to the first available lattice point.
  for (i = 0; i < XLATTICE; i++) {
    for (j = 0; j < YLATTICE; j++) {
      for (k = 0; k < ZLATTICE; k++) {
        if (lattice[i][j][k].inhabited == 0) {
          newpoint = lattice[i][j][k].position;
          lattice[i][j][k].inhabited = 1;
          onlattice = 1;
          break;
        }
      }
      if (onlattice == 1)
```

```
          break;

      }
      if (onlattice == 1)
        break;
    }
    // Now move the leftmost particle ’m’ to the chosen lattice point.
    meanx += point[m].position.x;
    particlemove(m,2, newpoint.x, newpoint.y, newpoint.z);
    point[m].initlattice = 1;

    if (immobile == 1)
      point[m].mobile = 0;
  }
  // Move all non-lattice particles left by meanx - meanlatticex.
  // This is roughly equiv to bringing the lattice *forward*
  // by mean x co-ord of those particles used to create it.
  // i.e.\ trying to preserve local concentration near lattice.
  meanx /= latticen;
  meanlatticex = (XLATTICE - 1) * (xlatsep / 2);
  for (i = 0; i < N; i++) {
    if (point[i].initlattice == 0) {
      newpoint = point[i].position;
      newpoint.x -= meanx;
      newpoint.x += meanlatticex;
      particlemove(i, 2,newpoint.x, newpoint.y, newpoint.z);
    }
  }
  return 0;
}
```

Moving the template forward, by the mean x coordinate of the particles removed from the bulk to create it, helps to prevent a spuriously low initial concentration in the region in front of the template. Overlaps between bulk and template particles may be created, however. While they exist, the particles involved are only allowed to move in ways which reduce their overlap with the lattice. This results in the removal of all overlaps in the first few t_d of the simulation.

A.1.4 Measurement of local volume fraction

A key visualisation technique for crystal growth (see e.g. Chapter 6) is the ‘area fraction’ method of determining local volume fraction. This involves taking planar slices through the system in (say) the $y - z$ plane, along the x axis, and calculating how much of each plane’s area is taken up by (circular) slices through particles. One thereby measures the local volume fraction in an infinitesimal interval δx. If the consecutive planes along the x axis have a spacing significantly smaller than a particle diameter, a useful measurement of y and z averaged volume fraction through the system is made. There is a considerable speed benefit in only calculating areas, as opposed to volumes, and the resolution can be increased arbitrarily simply by bringing the planes closer together.

As far as I am aware, this is a novel way of observing, without the need for binning, the variation in concentration along an axis. It is especially simple where the particles are spherical (so that the intersection area is easily calculated), and is naturally suitable for cases where the variation along one axis in particular is more important than the others, as in my crystal growth simulations. It averages over fluctuations in the y and z dimensions parallel to the crystal interface, allowing the crystal's growth to be mapped onto a single spatial axis. This enables the construction of plots like those in Section 6.4.1, in which the entire growth trajectory of the crystal through time and (1-dimension of) space can be visualised on a 2-dimension + colour plot.

Note also that, in polydisperse systems, arguably the most useful overall measure of 'amount of stuff' for comparative purposes *is* volume fraction, as opposed to number density. See for example Ref. [Wilding and Sollich, 2004], in which the position of the polydisperse gas-liquid binodal appears far more similar to the monodisperse case when the volume fraction representation is used.

The function that calculates this quantity is:

`areadensity(...)`

Because of the relatively inexpensive calculation, this is one of the few analysis functions that is actually run during the simulation, as opposed to being performed in post-processing.

```c
// Return the area of ’matter’ in a plane at a given x co-ordinate. This is quite clever.
// Axis: x. Operates directly on areadensity, crystdensity pointers,
// to tell the caller what is the area density, and the area density of only crystalline
// particles.
double areadensity(double pos, double *area, double *cryst) {
  int poscell, xcell, ycell, zcell, planecell;
  int minacell, maxacell, minbcell, maxbcell, minccell, maxccell;
  int howmany;
  int count;
  int i, a, b, c, h;
  double sep, radius;
  double areadensity, crystareadensity;
  // The x cell of the plane:
  planecell = pos / XCELL;
  maxacell = planecell + 3;
  minacell = planecell -2;
  maxbcell = YNUM;
  minbcell = 0;
  maxccell = ZNUM;
  minccell = 0;
  areadensity = 0;
  crystareadensity = 0;
  count = 0;
  for (a = minacell; a < maxacell; a++) {
    xcell = periodiccellx(a);
    for (b = minbcell; b < maxbcell; b++) {
      ycell = periodiccelly(b);
      for (c = minccell; c < maxccell; c++) {
        zcell = periodiccellz(c);

        howmany = cell[zcell + ZNUM * (ycell + YNUM * xcell)].occupants;
        // NB h starting at 1
        for (h = 1; h <= howmany; h++) {
          i = cell[zcell + ZNUM * (ycell + YNUM * xcell)].particle[h];
          sep = point[i].position.x - pos;
          sep = periodicx2(sep);

          radius = point[i].radius;

          sep *= sep;
          radius *= radius;
          if (sep < radius) {
            areadensity += (M_PI * (radius - sep));
            count++;
```

```
            if (point[i].crystal == 1) {
              crystareadensity += (M_PI * (radius - sep));
            }
          }
        }
      }
    }
  }

  areadensity /= (LY * LZ);
  crystareadensity /= (LY * LZ);
  *area = areadensity;
  *cryst = crystareadensity;
  return 0;
}
```

A.2 Trajectory files

The initialisation properties of the simulation (cell geometry, particle size distribution, volume fraction etc.) are set by the file:

`preinit.txt`

while properties such as the square well interaction parameters are set in the file:

`simconfig.txt`

After the simulation has been configured through these files and has been run, the necessary data describing a simulation trajectory for analysis purposes is contained in two files. The file:

`DAT_static_props.dat`

contains the ‘static’ properties of the simulation, i.e. those which do not change through time. An example static properties file is shown below:

```
RANDOMSEED = 500
N = 5000
VOLFRAC = 0.34000000000000002442
LX = 66.40048685359379021520
LY = 10.82503886715439733734
LZ = 10.82503886715439733734
DELTA = 0.08000000000000000167
TIMESTEP = 0.00020000000000000001
XLATTICE = 1
YLATTICE = 10
ZLATTICE = 10
LATSEP = 1.08250388671543973373
FRACT = 0.58950000000000002398
STATICLATTICE = 1
WALLWIDTH = 1.77160717396533096135
XPERIODIC = 0
maxrad = 0.59053572465511028344
minrad = 0.40195331004775580741
meanrad = 0.50000000000000000000
meanrad^2 = 0.25087629109113401826
meanrad^3 = 0.12631371135273947170
polydisp = 0.05920442858888252252
```

```

0 0.53714621538851226035  60.43276993382191619730 6.94684812847467103580
     3.07619593477889763378  1
1 0.49633163352350928577  37.79542565694896438799 3.53764099190837599807
     4.14574706813097382252  1
2 0.52258006650135480076  1.68012068345685694482  0.22060360680406057576
     9.84655924144557914701  1

[PARTICLE LIST CONTINUES...]
```

As can be seen, the static properties file contains all the overall simulation properties which may be required by an analysis program, followed by a particle list describing the starting configuration, in the form:

`index, radius, x, y, z, is-mobile`

The dynamic properties of the trajectory, i.e. the particle positions through time, are contained in numbered files named:

```
0_DAT_dynamic_props.dat
1_DAT_dynamic_props.dat
...
```

The static and dynamic 'DAT' files form the essential input for any post-processing/analysis code which must be run. Importantly, very little analysis is actually performed during the simulation, allowing trajectories to be generated quickly and new analyses to be performed as they become necessary.

A.3 Analysis code

I now briefly introduce some examples of analysis code which is used after the simulation has been run. The production simulation trajectories are stored in the form of consecutively-numbered files containing the time, particle index, radius, x, y, z coordinates, etc. Analysis codes usually take these files as input or, in the case of the 'Interface finder' described below, take in the time-dependent volume fraction profile of the system.

A.3.1 Interface finder

For the purposes of measuring e.g. correlation functions either inside or outside the crystalline region (as in Chapter 6), it is necessary to identify the position of the crystal interface. Another method would be simply to include (or exclude) only those particles flagged as crystalline by the flagging algorithm described in 3.2.5. However, in doing so, one would be exposed to the possibility of biases introduced by that algorithm. For instance, in a polydisperse crystal, the flagging algorithm seems to favour the identification of large particles as crystalline on the basis of their bond orientational order, whereas a measurement of *all* particles behind the crystal interface may yield no apparent preference. Using a particle's position relative to the interface therefore seems the best approach: all particles well within the crystal (i.e. far from the interface) will be recognised without bias. Particles occupying the 'grey area' around the interface may or may not be included depending how one chooses to identify the interface, but again there should be no overall bias, and the contribution to any statistics of these particles will become negligible as the crystal increases in extent. Most importantly, there is no exposure to the possibility of dependence of a particle's bond orientational order parameter on the other properties (in this case, particle size) that are being measured.

Due to the geometry of the system and the templating procedure, it is known that crystal growth will proceed along the x axis from the template at $x = 0$ and that, typically, the crystal interface will appear quite flat, being comparatively short in the y and z dimensions. Therefore, if we imagine moving through the system along the x axis, the crystal interface will be characterised by a sharp drop in local volume fraction at some value of x, since the crystal is the densest phase.

Information on local volume fraction is available from the output of the code outlined in Section A.1.4. The interface finding program scans through the x-dependent profile, calculating the average volume fraction in an averaging window (typically of size 1.5 particle diameters). The point at which the volume fraction first drops below some specified threshold is then identified as the crystal interface at that time. For consistency, the threshold is usually set to the parent volume fraction ϕ_p – as soon as the crystal has grown one or two layers, the fluid volume fraction immediately next to the interface will be depleted below ϕ_p. However, the results appear quite insensitive to the choice of any sensible value of the threshold, and are checked visually against the concentration profiles in order to ensure sanity.

When more precise information on phase interfaces and their properties (e.g. width, curvature) is required, one might instead proceed by fitting a tanh curve to a time-averaged density profile. The method described here is sufficient for my purposes, given that the interface is moving quickly and that the contribution to any measurements from particles within the interface region is small in comparison to that of the much larger number of particles either fully inside or fully outside the interface.

The code sample shown below is from the file:

`Interface_finder/interface.cpp`

beginning at line 70.

```
  // Scan through input files area density points:
    while (fscanf(fin, "\n%lf", &t) == 1) {
        fscanf(fin, "\t%lf\t%lf", &x, &aread);
        // If we've moved into the next timestep:
        if (t > tcurrent) {
            printf("\nTime = %lf", t);
            tcurrent = t;
            n=0;
            interface_found = 0;
        }
        // Just keep going through the file if we've found the interface already
        // in order to get through to the next timestep:
        if (interface_found)
            continue;

        // Populate the arrays of xposns and their associated area densities.
        // If the arrays are unfilled (i.e.\ n < navg) just plonk the values right in.
        // Else, shift the values along (oldest one falls off the end).
        if (n < navg) {
            averaging[n] = aread;
```

```
            xposns[n] = x;
        } else {
            for (i = 0; i < (navg-1); i++) {
                averaging[i] = averaging[i+1];
                xposns[i] = xposns[i+1];
            }
            averaging[navg-1] = aread;
            xposns[navg-1] = x;
        }
        // Increment the counter here:
        n++;

        // If the array is now full, calculate the average area density for this avging
            window:
        // NB -- this is inefficient, but easier to understand, leave it like this.
        if (n >= navg) {
            total = 0;
            // Find average area density:
            for (i = 0; i < navg; i++) {
                total += averaging[i];
            }
            // If the averaged concentration in the window falls below the threshold, mark
                the midpoint of it as the xposition of the interface.
            if (total < navg_times_phi) {
                interface_found = 1;
                printf("\nFound Interface..");
                fprintf(fout, "\n%lf\t%lf", t, 0.5*(xposns[navg-1] + xposns[0]));
            }
        }
    }
  printf("\nDone.\n");
  return 0;
}
```

A.3.2 Crystal region statistics

To illustrate the general approach to making fractionation measurements, I now present code from:

`MeanRadius_Through_Time_Interfaceposn/meanrad_interface.cpp`

This code uses the output of the interface finder (described in the previous section) and the main trajectory files containing particle positions and size information to measure the properties of particles in the crystalline region (i.e. behind the crystal interface). This was used e.g. to measure the crystal polydispersity through time in Chapter 6. The output file contains four columns: the time, number of crystalline particles, mean radius, and mean-squared radius (in order to calculate polydispersities).

N.B. In the below, the construction within the function:

`fscanf(...)`

to discard a read-in variable has been rendered as:

`%\*`

i.e. with a backslash before the asterisk, in order to avoid errors in typesetting the code. Normally, the backslash before the asterisk would not be present.

```
    while (1) {
        path = stringify(dynfile) + "_DAT_dynamic_props.dat";

        if ( (fin=fopen(path.c_str(), "r")) == NULL) {
            printf("\nReached final dynamic_props file, prefix %d", dynfile);
            break;
        }
        // Next loop -- go through all timesteps.
        notfinished = 1;
        while (notfinished) {
            n_cryst = 0;
            total = 0;
            total2 = 0;
            // Read DAT_x_interface_posn.dat to find interface position.
            fscanf(fin2, "\n%\*lf\t%lf", &interface_posn);
            // Go through all particles for this timestep.
            for (i=0; i < N; i++) {
                // Scan in particle info.
                if (fscanf(fin, "\n%lf\t%\*d\t%lf\t%lf\t%\*lf\t%\*lf\t%d\t%\*d\t%\*d\t%\*d\t
                     %\*d", &t, &radius, &xpos, &crystflag) == EOF) {
                    notfinished = 0;
                }
                // If the read failed.
                if (notfinished != 1) {
                    printf("\nReached EOF\n.");
                    if (i != 0) {
                        printf("\nEOF not at i=0, something's wrong.");
                        return -1;
                    }
                    break;
                }
                // Add any crystal particle to statistics for this timestep.
                //if (crystflag) {
                if (xpos < interface_posn) {
                    n_cryst++;
                    total += radius;
                    total2 += radius\*radius;
                }
            }
            // If the above didn't result in EOF and a loop break, print results.
            if (notfinished) {
                meanrad = total / n_cryst;
                meanrad2 = total2 / n_cryst;
                fprintf(fout,"\n%lf\t%d\t%lf\t%lf", t, n_cryst, meanrad, meanrad2);
                printf("\nTime = %lf", t);
            }
        }
        dynfile++;
    }
    return 0;
}
```

References

Ackerson, B. J., and K. Schätzel (1995), Phys. Rev. E **52**, 6448.

Ackland, G. J., and A. P. Jones (2006), Phys. Rev. B **73**, 054104.

Alder, B. J., and T. E. Wainwright (1957), J. Chem. Phys. **27**, 1208.

Allard, M., and E. H. Sargent (2004), Appl. Phys. Lett. **85**, 5887.

Amini, M., and B. B. Laird (2006), Phys. Rev. Lett. **97**, 216102.

Anderson, V. J., and H. N. W. Lekkerkerker (2002), Nature **416**, 811.

Asakura, S., and F. Oosawa (1958), J. Polym. Sci. **33**, 183.

Auer, S., and D. Frenkel (2001), Nature **413**, 711.

Auer, S., and D. Frenkel (2005), in *Advanced Computer Simulation*, Advances in Polymer Science, Vol. 173, edited by C. Holm and K. Kremer (Springer Berlin / Heidelberg) pp. 130–130.

Bailey, A. E., W. C. K. Poon, R. J. Christianson, A. B. Schofield, U. Gasser, V. Prasad, S. Manley, P. N. Segre, L. Cipelletti, W. V. Meyer, M. P. Doherty, S. Sankaran, A. L. Jankovsky, W. L. Shiley, J. P. Bowen, J. C. Eggers, C. Kurta, T. Lorik, P. N. Pusey, and D. A. Weitz (2007), Phys. Rev. Lett. **99**, 205701.

Bartlett, P. (1997), J. Chem. Phys. **107**, 188.

Bartlett, P. (1998), J. Chem. Phys. **109**, 10970.

Bartlett, P., and P. B. Warren (1999), Phys. Rev. Lett. **82**, 1979.

Batchelor, G. K., and J. T. Green (1972), J. Fluid Mech. **56**, 375.

Berthier, L. (2011), Physics **4**, 42.

van Blaaderen, A., R. Ruel, and P. Wiltzius (1997), Nature **385**, 321.

van de Bovenkamp, J., and J. K. G. Dhont (1994), Physica A **212**, 239.

Brader, J. M. (2001), *Statistical Mechanics of a Model Colloid-Polymer Mixture*, Ph.D. thesis (Bristol).

Brady, J. F., and G. Bossis (1988), Annu. Rev. Fluid Mech. **20**, 111.

Burke, E., J. Q. Broughton, and G. H. Gilmer (1988), J. Chem. Phys. **89**, 1030.

Butter, K., K. Kassapidou, G. Vroege, and A. Philipse (2005), J. Colloid Interface Sci **287**, 485.

Buzzacchi, M., I. Pagonabarraga, and N. B. Wilding (2004), J. Chem. Phys. **121**, 11362.

Byelov, D. V., M. C. D. Mourad, I. Snigireva, A. Snigirev, A. V. Petukhov, and H. N. W. Lekkerkerker (2010), Langmuir **26**, 6898.

Carnahan, N. F., and K. E. Starling (1969), J. Chem. Phys. **51**, 635.

Cates, M. E., K. Stratford, R. Adhikari, P. Stansell, J.-C. Desplat, I. Pagonabarraga, and A. J. Wagner (2004), J. Phys.: Condens. Matter **16**, S3903.

Choi, Y., T. Ree, and F. H. Ree (1991), J. Chem. Phys. **95**, 7548.

Cichocki, B., and B. U. Felderhof (1991), J. Chem. Phys. **94**, 556.

Cichocki, B., and K. Hinsen (1990), Physica A: Statistical Mechanics and its Applications **166** (3), 473 .

Colvin, V. L. (2001), MRS Bulletin **26**, 637.

Dijkstra, M., J. M. Brader, and R. Evans (1999), J. Phys.: Condens. Matter **11** (50), 10079.

Dinsmore, A. D., M. F. Hsu, M. G. Nikolaides, M. Marquez, A. R. Bausch, and D. A. Weitz (2002), Science **298**, 1006.

Dixit, N. M., and C. F. Zukoski (2000), J. Colloid Interface Sci. **228**, 359.

Doye, J. P. K., A. A. Louis, I.-C. Lin, L. R. Allen, E. G. Noya, A. W. Wilber, H. C. Kok, and R. Lyus (2007), Phys. Chem. Chem. Phys. **9**, 2197.

Erné, B. H., E. van den Pol, G. J. Vroege, T. Visser, and H. H. Wensink (2005), Langmuir **21**, 1802.

Evans, R. M. L. (2001), J. Chem. Phys. **114**, 1915.

Evans, R. M. L. (2005), J. Phys. A **38**, 293.

Evans, R. M. L., and D. J. Fairhurst (2004), Colloid. Polym. Sci. **282**, 766.

Evans, R. M. L., D. J. Fairhurst, and W. C. K. Poon (1998), Phys. Rev. Lett. **81**, 1326.

Evans, R. M. L., and C. B. Holmes (2001), Phys. Rev. E **64**, 011404.

Evans, R. M. L., and W. C. K. Poon (1997), Phys. Rev. E **56**, 5748.

Evans, R. M. L., W. C. K. Poon, and M. E. Cates (1997), Europhys. Lett. **38**, 595.

Evans, R. M. L., W. C. K. Poon, and F. Renth (2001), Phys. Rev. E **64**, 031403.

Fasolo, M., and P. Sollich (2003), Phys. Rev. Lett. **91**, 068301.

Fasolo, M., and P. Sollich (2004), Phys. Rev. E **70**, 041410.

Fasolo, M., and P. Sollich (2005), J. Chem. Phys. **122**, 074904.

Filion, L., M. Hermes, R. Ni, and M. Dijkstra (2010), J. Chem. Phys. **133**, 244115.

Foffi, G., G. D. McCullagh, A. Lawlor, E. Zaccarelli, K. A. Dawson, F. Sciortino, P. Tartaglia, D. Pini, and G. Stell (2002), Phys. Rev. E **65**, 031407.

Fortini, A., E. Sanz, and M. Dijkstra (2008), Phys. Rev. E **78**, 041402.

Gast, A., C. Hall, and W. Russel (1983), J. Colloid Interface Sci. **96**, 251.

Ginoza, M., and M. Yasutomi (1999), Phys. Rev. E **59**, 2060.

Haas, C., and J. Drenth (1999), J. Cryst. Growth **196**, 388.

Haas, C., and J. Drenth (2000), J. Phys. Chem. B **104**, 368.

Hall, K. R. (1972), J. Chem. Phys. **57**, 2252.

Hansen, J., and I. McDonald (2006), *Theory of Simple Liquids*, Academic Press (Elsevier Academic Press).

Hiemenz, P. C., and R. Rajagopalan (1997), *Principles of Colloid and Surface Chemistry, Third Edition, Revised and Expanded*, Undergraduate chemistry (Taylor & Francis).

Hiramatsu, H., and F. E. Osterloh (2004), Chem. Mater. **16**, 2509.

Hoogenboom, J. P. (2002), *Colloidal epitaxy : a real-space analysis*, Ph.D. thesis (Utrecht (Netherlands)).

Hoogenboom, J. P., A. K. van Langen-Suurling, J. Romijn, and A. van Blaaderen (2004), Phys. Rev. E **69**, 051602.

Ilett, S. M., A. Orrock, W. C. K. Poon, and P. N. Pusey (1995), Phys. Rev. E **51**, 1344.

Jarzynski, C. (1997), Phys. Rev. Lett. **78**, 2690.

Jodrey, W. S., and E. M. Tory (1985), Phys. Rev. A **32**, 2347.

Lekkerkerker, H. N. W., W. C. K. Poon, P. N. Pusey, A. Stroobants, and P. B. Warren (1992), Europhys. Lett. **20**, 559.

Leocmach, M., and H. Tanaka (2013), Soft Matter **9**, 1447.

Liddle, S. M., T. Narayanan, and W. C. K. Poon (2011), J. Phys. Condens. Matter **23**, 194116.

Lin, K.-H., J. C. Crocker, V. Prasad, A. Schofield, D. A. Weitz, T. C. Lubensky, and A. G. Yodh (2000), Phys. Rev. Lett. **85**, 1770.

Liu, H., S. Garde, and S. Kumar (2005), J. Chem. Phys. **123**, 174505.

Lo Verso, F., R. L. C. Vink, D. Pini, and L. Reatto (2006), Phys. Rev. E **73**, 061407.

Martin, P. A. (1977), J. Stat. Phys. **16**, 149.

Martin, S., G. Bryant, and W. van Megen (2003), Phys. Rev. E **67**, 061405.

Martin, S., G. Bryant, and W. van Megen (2005), Phys. Rev. E **71**, 021404.

Mattsson, J. (2012), Private communication.

Medina-Noyola, M. (1988), Phys. Rev. Lett. **60**, 2705.

Middleton, T. F., and D. J. Wales (2004), J. Chem. Phys. **120**, 8134.

Mulero, A., C. Galãn, and F. Cuadros (2001), Phys. Chem. Chem. Phys. **3**, 4991.

Mullins, W. W., and R. F. Sekerka (1963), Journal of Applied Physics **34** (2), 323.

Murarka, R. K., and B. Bagchi (2003), Phys. Rev. E **67**, 051504.

Novotny, M. A. (2001), eprint arXiv:cond-mat/0109182 arXiv:cond-mat/0109182 .

O'Malley, B., and I. Snook (2005), J. Chem. Phys. **123**, 054511.

Onsager, L. (1949), Ann. N.Y. Acad. Sci. **51**, 627.

Pagonabarraga, I., M. E. Cates, and G. J. Ackland (2000), Phys. Rev. Lett. **84**, 911.

Phan, S.-E., W. B. Russel, J. Zhu, and P. M. Chaikin (1998), J. Chem. Phys. **108**, 9789.

Poon, W. C. K. (1997), Phys. Rev. E **55**, 3762.

Poon, W. C. K. (2002), J. Phys.: Condens. Matter **14**, R859.

Poon, W. C. K. (2012), Private communication.

Poon, W. C. K., F. Renth, and R. M. L. Evans (2000), J. Phys.: Condens. Matter **12**, A269.

Poon, W. C. K., F. Renth, R. M. L. Evans, D. J. Fairhurst, M. E. Cates, and P. N. Pusey (1999), Phys. Rev. Lett. **83**, 1239.

Renth, F., W. C. K. Poon, and R. M. L. Evans (2001), Phys. Rev. E **64**, 031402.

Riese, D. O., G. H. Wegdam, W. L. Vos, R. Sprik, D. Fenistein, J. H. H. Bongaerts, and G. Grübel (2000), Phys. Rev. Lett. **85**, 5460.

Romano, F., C. D. Michele, D. Marenduzzo, and E. Sanz (2011), J. Chem. Phys. **135**, 124106.

Rotenberg, B., J. Dzubiella, J.-P. Hansen, and A. A. Louis (2004), Molec. Phys. **102** (1), 1.

Russel, W. (1987), *The dynamics of colloidal systems*, Olaf A. Hougen lectures (University of Wisconsin Press).

Russel, W. B. (2003), Nature **421**, 490.

Russel, W. B., D. A. Saville, and W. R. Schowalter (1992), *Colloidal Dispersions*, Cambridge Monographs on Mechanics and Applied Mathematics (Cambridge University Press).

Sanz, E., and D. Marenduzzo (2010), J. Chem. Phys. **132**, 194102.

Sanz, E., C. Valeriani, D. Frenkel, and M. Dijkstra (2007), Phys. Rev. Lett. **99**, 055501.

Scheffold, F., and T. G. Mason (2009), J. Phys.: Condens. Matter **21**, 332102.

Schmidt, N., S. Barr, A. Udit, L. Gutierrez, T. Nguyen, M. G. Finn, E. Luijten, and G. Wong (2010), in *APS Meeting Abstracts*, p. 19011.

Scholz, J., J. Kelso, and G. Schöner (1987), Phys. Lett. A **123**, 390.

Schöpe, H. J., G. Bryant, and W. van Megen (2007), J. Chem. Phys. **127**, 084505.

Segrè, P. N., O. P. Behrend, and P. N. Pusey (1995), Phys. Rev. E **52**, 5070.

Segrè, P. N., and P. N. Pusey (1996), Phys. Rev. Lett. **77**, 771.

Shilpi, S., A. Jain, Y. Gupta, and S. K. Jain (2007), Crit. Rev. Ther. Drug Carrier Syst. **24**, 361.

Sollich, P. (2002), J. Phys.: Condens. Matter **14**, R79.

Sollich, P., and M. E. Cates (1998), Phys. Rev. Lett. **80**, 1365.

Sollich, P., and N. B. Wilding (2010), Phys. Rev. Lett. **104**, 118302.

Sollich, P., and N. B. Wilding (2011), Soft Matter **7**, 4472.

Stein, A., B. E. Wilson, and S. G. Rudisill (2013), Chem. Soc. Rev. 10.1039/C2CS35317B, (Advance article, DOI: 10.1039/C2CS35317B).

Stell, G. (1963), Physica **29**, 517.

Stukowski, A. (2010), Modelling Simul. Mater. Sci. Eng. **18**, 015012.

Suresh, S. (2006), Nat. Mater. **5**, 253.

Tadros, T. F. (2007), *Colloid stability and application in pharmacy*, Colloids and interface science series (Wiley-VCH Verlag).

Tadros, T. F. (2008), "Colloid aspects of cosmetic formulations with particular reference to polymeric surfactants," in *Colloids in Cosmetics and Personal Care* (Wiley-VCH Verlag GmbH & Co. KGaA) pp. 1–34.

Tegze, G., L. Gránásy, G. I. Tóth, F. Podmaniczky, A. Jaatinen, T. Ala-Nissila, and T. Pusztai (2009), Phys. Rev. Lett. **103**, 035702.

Tirado-Miranda, M., C. Haro-Pérez, M. Quesada-Pérez, J. Callejas-Fernández, and R. Hidalgo-Álvarez (2003), J. Colloid Interface Sci. **263**, 74.

Van Blaaderen, A., and A. Vrij (1992), Langmuir **8**, 2921.

Velikov, K. P., and A. van Blaaderen (2001), Langmuir **17**, 4779.

Vrij, A. (1976), Pure Appl. Chem. **48**, 471.

Wang, W., M. K. Khoshkbarchi, and J. H. Vera (1996), Fluid Phase Equilib. **115**, 25.

Warren, P. B. (1998), Phys. Rev. Lett. **80**, 1369.

Warren, P. B. (1999), Phys. Chem. Chem. Phys. **1**, 2197.

Wette, P., I. Klassen, D. Holland-Moritz, T. Palberg, S. V. Roth, and D. M. Herlach (2009), Phys. Rev. E **79**, 010501.

Weysser, F., A. M. Puertas, M. Fuchs, and T. Voigtmann (2010), Phys. Rev. E **82**, 011504.

Whitelam, S., and P. L. Geissler (2007), J. Chem. Phys. **127**, 154101.

Wilding, N. B., and P. Sollich (2004), Europhys. Lett. **67**, 219.

Williams, S. R., C. P. Royall, and G. Bryant (2008), Phys. Rev. Lett. **100**, 225502.

Williamson, J. J., and R. M. L. Evans (2012), Phys. Rev. E **86**, 011405.

Williamson, J. J., and R. M. L. Evans (2013), Soft Matter , (Advance article, DOI: 10.1039/C3SM27627A).

ten Wolde, P.-R., M. J. Ruiz-Montero, and D. Frenkel (1996), Faraday Discuss. **104**, 93.

Yethiraj, A. (2007), Soft Matter **3**, 1099.

Yethiraj, A., and A. Van Blaaderen (2003), Nature **421**, 513.

Yiannourakou, M., I. G. Economou, and I. A. Bitsanis (2010), J. Chem. Phys. **133**, 224901.

Yu, Y.-X., J. Wu, Y.-X. Xin, and G.-H. Gao (2004), J. Chem. Phys. **121**, 1535.

Zaccarelli, E. (2007), J. Phys.: Condens. Matter **19**, 323101.

Zaccarelli, E., C. Valeriani, E. Sanz, W. C. K. Poon, M. E. Cates, and P. N. Pusey (2009), Phys. Rev. Lett. **103**, 135704.