

University of Warwick institutional repository: <http://go.warwick.ac.uk/wrap>


A Thesis Submitted for the Degree of PhD at the University of Warwick

<http://go.warwick.ac.uk/wrap/63807>

This thesis is made available online and is protected by original copyright.

Please scroll down to view the document itself.

Please refer to the repository record for this item for information to help you to cite it. Our policy information is available from the repository home page.



Computer Simulation of Mesocrystals

by

Štěpán Růžička

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Department of Physics

July 2014

THE UNIVERSITY OF
WARWICK

Contents

Acknowledgments	iv
Declarations	v
Abstract	vi
Chapter 1 Introduction	1
1.1 Mesocrystals	1
1.2 Atomistic scale	3
1.3 Colloidal scale	3
1.4 Colloidal gelation	5
1.5 Scale of colloidal aggregates	8
1.6 Classical crystallization	13
1.7 Non-classical crystallization	14
1.8 Mesocrystallization	15
1.9 Simulation methods	17
1.10 Dynamical fidelity	18
Chapter 2 Collective moves for attractive potentials	21
2.1 Introduction	21
2.2 Definitions and basic relations	24
2.3 Link formation probabilities	30
2.4 Free cluster selection	32
2.5 Selection of the cluster under high acceptance probability	35
2.6 Controlling the cluster size	37
2.7 Collective rotational Monte Carlo move	40
2.8 Selection of the cluster under low acceptance probability	42
2.9 Recursive selection of the cluster	44
2.10 Explicit formulation of the superdetailed balance condition	46

2.11	Results	49
2.12	Discussion	53
2.13	Conclusion	58
Chapter 3	Kinetically slowed down phase separation	59
3.1	Introduction	59
3.2	Methods	60
3.2.1	Simulation model	60
3.2.2	Analysis	62
3.3	Equilibrium fluids	65
3.4	Undercooled attractive systems	67
3.4.1	Coarsening in single-particle MC	67
3.4.2	Displacement size in single-particle MC	70
3.4.3	Phase separation under collective motion	72
3.4.4	Generation and acceptance densities	84
3.5	Undercooled charged attractive systems	90
3.5.1	Stabilization of the aggregates	90
3.5.2	Scaling of the dynamics	95
3.5.3	Generation and acceptance densities	97
3.6	Discussion	98
3.7	Conclusion	106
Chapter 4	Collective moves for general potentials	108
4.1	Introduction	108
4.1.1	Linking probabilities	112
4.1.2	Early rejection scheme	113
4.2	Results	114
4.2.1	Equilibrium sampling	114
4.2.2	Enhancement of acceptance probability	115
4.2.3	Scaling of the dynamics	115
4.3	Discussion	116
4.4	Conclusion	118
Chapter 5	Renormalized long-range repulsion	122
5.1	Simulation model	122
5.2	Results	125
5.3	Discussion	126
5.4	Conclusion	131

Chapter 6 Discussion	132
6.1 Alignment in the Wigner glass	132
6.1.1 Formation of monodisperse aggregates	132
6.1.2 Alignment of partitions	135
6.1.3 Compression into a mesocrystal	136
6.2 Alignment via oriented attachment	137
6.3 Micelle as a building block	139
Chapter 7 Conclusion	141
Appendix A The inconsequential nature of the asymmetric core in Sec. 2.5	145
Appendix B Derivation of the acceptance probability of the algorithm of Sec. 2.8	147
Appendix C Validation of the early rejection schemes	152
C.1 Algorithm of Sec. 2.4	152
C.2 Algorithm of Sec. 2.6	153
C.3 Algorithm of Sec. 2.8	154

Acknowledgments

I would like to thank my supervisor Prof. Michael P. Allen for his great patience and care which made this thesis possible. I would like to thank my second supervisor Dr. David Quigley for critical feedback and for giving me plenty of freedom during my research. I thank Yuriy Bushuev, Helmut Cölfen, Jasmine Desmond, Jonathan Doye, James Elliott, Aaron Finney, James Grant, Robert Jack, Patrick Kiley, Bartosz Marzec, Aleks Reinhardt, Gil Rutter, Anna Schenk, Adam Swetman, Andreas Verch, Bart Vorselaars, and Louise Wright for discussions. In particular, I would like to thank Yi-Yeoun Kim for discussions about the experimental approach to mesocrystallization, and Trung Nguyen and Stephen Whitlam for correspondence explaining the renormalized repulsion model and the Virtual Move Monte Carlo algorithm. I thank the Centre for Scientific Computing of the University of Warwick for computational resources. By name, I thank Matthew Ismail and Magnus Lewis-Smith from the Centre for Scientific Computing. I would like to thank everyone from Warwick University who supported a relaxed social environment which turned my stay at Warwick into a stimulating and enjoyable experience. I acknowledge financial support received from the Warwick Postgraduate Research Scholarship and from the EPSRC grant ‘Hard-soft matter interfaces: from understanding to engineering’ with reference number EP-I001514.

Declarations

The single particle Monte Carlo code was provided by M. P. Allen. The code calculating bond orientational parameters was written by D. Quigley.

Abstract

The mesoscale is a thousand times larger than the atomistic scale with colloidal particles, rather than atoms or molecules, forming the constituent building blocks for organized structures. Nanotechnology has recently started interpreting colloidal units as colloidal molecules, and a lot of interest emerged in their assembly into self-organized structures called colloidal crystals. A mesocrystal is a special type of colloidal crystal, where constituent colloidal units are crystallographically registered nanocrystals. Computer simulation of colloidal self-assembly requires coarse-graining, where short-range attraction is usually the dominant force on the colloidal scale. In this thesis, moderately deep quenches of short-range attractive spherical colloids at low packing fractions are performed, because strong clustering into liquid drops preceding crystallization is known to follow quenching and is reminiscent of mesocrystallization. Moreover, variation of simulation parameters in those models allows an easy investigation of the competition between the fundamental processes of multiscale assembly such as coarsening, crystallization, gelation, and phase separation interfering with the dynamical slowing down. Since colloidal building blocks are usually characterized by complex anisotropic interactions we use Monte Carlo rather than dynamical simulations. The thesis is focused on the development of the Virtual Move Monte Carlo which is a Monte Carlo cluster algorithm selecting the moving clusters according to the local energy gradients. The algorithm allows one to control the rates of local crystal evolution and a larger scale cluster aggregation. The thesis investigates the conditions at which the crystallization precedes the aggregation and vice versa. It is confirmed by the Monte Carlo simulations that the properties of kinetically slowed down structures are independent of the microscopic dynamics, and that three different linear growth regimes, correlated to the local order, are present in systems where structure evolves mainly via single particle exchange. It is also shown that long-range repulsions are necessary to stabilize the phase separated aggregates, and that a renormalization of the repulsion leads to a significant decrease of polydispersity of the liquid or recrystallized drops. The efficiency of rotational cluster moves simulating the alignment into crystallographic register is enhanced by several orders of magnitude in this thesis. The results are discussed in the language of recent colloidal physics and related to a wider range of coarsening and self-assembly phenomena observed during the nonclassical crystallization or mesocrystallization processes.

Chapter 1

Introduction

1.1 Mesocrystals

Mesocrystals are solids composed of nanocrystalline units which are crystallographically aligned [Cölfen and Antonietti, 2008]. A soft organic layer separating the nanocrystals can be required in the definition to distinguish the mesocrystal from a single crystal. The presence of a layer is, however, not necessary, and a single crystal can be referred to as a mesocrystal, provided it formed via a self-assembly of smaller crystals [Leite and Ribeiro, 2012]. The constituent nanocrystalline units have typical size ranging from a nanometre to a micrometer, and mesocrystals are sometimes considered to be a special class of colloidal crystals. Mesocrystals are commonly present in nature, in particular, as different structures and crystal polymorphs of calcium carbonate. The sea urchin spine and tooth are well studied examples of calcium carbonate mesocrystals with a soft organic layer [Seto et al., 2012; Killian et al., 2009]. Mesocrystals have remarkable material properties, similar to those of composite materials or superalloys. The crystalline content makes them brittle, the soft organic content ductile, and the bulk of the material is usually tough.

It is believed that to understand the mesocrystallization may provide a key to the design of new synthetically produced nanostructures [Zhou and O'Brien, 2012]. Moreover, since mesocrystals form at ambient temperatures, there is a possibility of artificially creating supermaterials in an ecological way, which would replace high temperature treatments in furnaces. To have a deeper insight into the formation of biological mesocrystals may also have biomedical applications such as the design of treatments for osteoporosis.

There are several reasons why mesocrystals are sometimes thought to be ill-defined. First, a single crystal can be termed a mesocrystal, if it is self-assembled

from smaller crystals via oriented attachment of crystalline surfaces [Yuwono et al., 2010]. This definition is related to the history of crystal formation rather than the structure, and once the mesocrystal forms it cannot be distinguished from a single crystal. Second, mesocrystals can also refer to structures which are different from single crystals in that the crystallographically aligned units are separated from other units by a soft organic layer of macromolecules or other soft material. Third, the mechanisms behind the alignment of nanocrystals can be different. For example, the mechanism behind the formation of the sea urchin spine [Seto et al., 2012] is different from the mechanism leading to the mesocrystal of cerium oxide [Dalmaschio et al., 2010] in that the self-organization of soft amorphous colloidal blocks precedes the crystallization on the atomistic scale.

The mesocrystals thus turn out to span a wide range of materials and processes, which have little in common. The only property linking these structures and mechanisms is the presence of a larger number of crystallographically aligned nanocrystals before or after the crystallization. The success behind the concept of mesocrystals seems to stem from the unification of a wider class of unexplored, usually biological, materials with potentially useful material properties.

This thesis is about the computer simulation of mesocrystals or more specifically about the process leading to mesocrystals called the *mesocrystallization*. Since mesocrystallization can be any crystallization pathway leading to the mesocrystal [Cölfen and Antonietti, 2008], there are potentially many processes which could be analysed. The thesis has two primary objectives. It turns out that collective motion [Whitelam et al., 2009a], and collective rotational motion in particular, are degrees of freedom which are crucial for hierarchical self-assembly, where building blocks are crystallographically aligned [Leite and Ribeiro, 2012; Hedges and Whitelam, 2011]. It is believed that mesocrystallization is a crystallization process, where multiscale modelling is necessary [Harding et al., 2008], and that the preservation of kinetics is essential to model nonclassical crystallization [Whitelam et al., 2009a]. To capture these hierarchical phenomena efficiently with a possibility of extending to complex interactions of nanobuilding units, we focus on the developments and tests of the Virtual Move Monte Carlo (VMMC) algorithm [Whitelam et al., 2009a; Whitelam and Geissler, 2007]. The second aim of this work is to apply the VMMC algorithm on a minimum system which may form structures similar to a mesocrystal. We use particles with short-range attractions and long-range repulsions (SALR), because these systems are expected to form modulated cluster crystals [Sciortino et al., 2004; Xia et al., 2011]. The definition of mesocrystals also involves atomistic alignment. This thesis does not distinguish between scales. The basic building block

is a spherical monomer, which primarily represents a colloid, but may also represent a molecule or an atom. This is possible because the physics at those scales can be similar. For instance, liquid-liquid coexistence usually studied on the colloidal scale [Anderson and Lekkerkerker, 2002; Lu et al., 2008] has also been reported on the atomistic scale of calcium carbonate [Wallace et al., 2013].

1.2 Atomistic scale

Consideration of individual atoms is fundamental to model the epitaxial alignment between the mesocrystalline units. Nevertheless, the atomistic scale of solutions forming the mesocrystals can be rich. It may include the simulation of nucleation [Demichelis et al., 2011], or the growth of the crystal [Stack et al., 2011], but may also involve the molecular driven polymorph selection [Fichthorn, 2014], and the attachment of molecules to surfaces [Wright and Walsh, 2013], which may further mediate the self-assembly [Harding et al., 2008]. A usual problem with atomistic simulations is the system size versus the achievable simulation timescale. In mesocrystallization, the amorphous or crystalline units constituting the mesocrystal may be micron-sized and may contain thousands up to billions of atoms or molecules. It is impossible to simulate the self-assembly of a reasonable number of such crystals, while considering the atomistic detail. Moreover, these systems are often formed of ionic solutions, where the ions precipitate, and where solvent slows down the diffusion to an extent that the simulations are not computationally tractable. In this thesis spherical monomers represent colloids but may also represent atoms or molecules. The small colloidal crystals may then be regarded as small atomistic nanocrystals, and the epitaxial alignment may be defined by the lattice of those crystals. The effect of solvent molecules, ions, or biomolecules mediating the self-assembly is then either neglected or is included in the pairwise interaction.

The thesis would not be complete without mentioning the work of Alimohammadi and Fichthorn [2009]. The authors of that paper use molecular dynamics to simulate the oriented attachment of two titanium dioxide nanocrystals to show that the alignment is due to the undercoordinated surface atoms of the crystals rather than the dipole-dipole interactions of the nanocrystals.

1.3 Colloidal scale

Colloids comprise materials composed of mesoscopic particles, which are immersed in a solvent whose molecules are considerably smaller, and have typically an atom-

istic size. The solvent may include short polymeric chains, causing short-range entropic attractions between the particles or salts dissociated into ions, which can be responsible for long-range repulsive Yukawa interactions. A prototypical example of a colloidal particle is a hard sphere, such as chemically synthesized polystyrene, silica or polymethyl methacrylate (PMMA). Over the last decades, the term colloid has been extended to a wider range of particles such as soft polymeric chains, polyelectrolytes, star polymers, or dendrimers with mesoscopic dimensions [Likos, 2013].

Physics provides theories to describe and model interactions at various length scales, ranging from microscopic quantum or atomistic fields to macroscopic gravitational forces between stellar objects. To observe phenomena which are of interest in the colloidal assembly only some of these interactions are important, and others may be neglected. This approximation, known as coarse graining, often represents a colloidal particle by an isotropic spherically symmetric potential field. Anisotropy may be introduced by changing the geometry of the potentials, or by decorating the colloidal particles with attractive or repulsive patches [Sciortino, 2010; Zhang and Glotzer, 2004], representing bio-molecular forces or other specific directional forces mediating the self-assembly [Haxton and Whitelam, 2012].

The most common interactions between colloidal particles are entropic short-range attractions, which are typically induced by short non-adhesive polymer chains or by smaller sized colloids in binary colloidal mixtures. Another common source of short-range attractions in colloidal systems are van der Waals attractions resulting from dipole fluctuations, and may be a combination of three or more terms such as the Keesom force, Debye force, dipole-charge interactions, or London dispersion forces. To guarantee the stability of colloidal suspensions against aggregation, long-range repulsive forces are usually present in the system. Electrostatic interaction resulting from charge on the colloids or from ions in solutions is the first common cause of the repulsion known as charge stabilization. The second common cause for long-range repulsion is steric stabilization, where colloids are grafted with polymer chains providing protective coating.

Assigning anisotropy to the particles may entirely change the phase diagram. The simplest example is a system of ellipsoids, where different crystal morphologies can be observed, depending on the geometry of the ellipsoids [Odriozola, 2012]. A system of hard cubes is another common example of simple anisotropic systems, and has recently received considerable attention due to the vacancy induced stability of the simple cubic crystal phase [Smallenburg et al., 2012]. Systems of convex polyhedra have recently been categorized by Damasceno et al. [2012] into a remark-

ably complex range of structures, having a potential significance in the theory of oriented attachment or mesocrystallization. This thesis only considers isotropically interacting particles to develop MC methods which are suitable for the simulation of self-assembly in systems of anisotropic particles.

1.4 Colloidal gelation

Gelation is a consequence of kinetics interfering with the thermodynamics. The fundamental kinetics of first order transitions can be studied on the Ising model [Landau and Binder, 2009, chap. 10.4]. It is known that if the system is rapidly quenched to low temperatures, the phase separation driven by spinodal decomposition is kinetically slowed down, and the system gets arrested in a long-living arrested structure percolating the system, called a gel. In a three dimensional short-range attractive system, the situation is richer, because of the possibility to move the gas liquid critical point below the fluid-solid coexistence line, thus generating density fluctuations and temporary stable clusters [Perez et al., 2011].

In what follows, we will discuss physical gelation which is reversible, depending on the strength of the pairwise interaction, as opposed to chemical gelation, which is considered to be irreversible [Zaccarelli, 2007]. Physical gelation is commonly studied using mixtures of colloids and nonadsorbing polymers [Lu et al., 2008], where the interaction strength is controlled by the concentration of polymers. Other systems are possible such as polymer grafted silica spheres, where the attraction is controlled by temperature affecting the solvent quality of polymer chains.

A gel is identified as the state of a system, where particles are attached to each other such that they form a network of attractive bonds percolating the system. It is required that the structure of this network does not change for a sufficiently long time, which is typically equal to, or greater than 10^2 s, similarly to glasses. If the bond structure changes, while the system percolates, we talk about transient percolation, which is not considered to be a gel. It is also expected that the stability of the network, given by the zero gradient of mean square displacement, is due to the strength of attraction forces, rather than the arrest caused by repulsive forces, which typically happens at high densities. In the latter case, the system is considered to be a repulsive glass, rather than a gel [Zaccarelli and Poon, 2009].

There are two main classes of gels. Gels which are formed via arrested phase separation are known to be non-equilibrium gels, because their thermodynamically stable state is expected to be different from the gel state [Charbonneau and Reichman, 2007a]. These gels are typically formed by fast quenches of short-range at-

tractive colloids both with and without the long-range repulsion [Foffi et al., 2005a; Charbonneau and Reichman, 2007b]. The second type of gel, referred to as equilibrium gels, occurs in short-range attractive systems with a suitably chosen long-range repulsion. The quench is usually slow and the typical equilibrium structure is the Bernal spiral [Sciortino et al., 2005].

The definition of arrested state is closely related to diffusion. Diffusion in different parts of the phase diagram can be analysed in terms of isodiffusivity lines [Foffi et al., 2002], i.e. loci in the phase diagram with the same diffusion constant. To find the kinetically arrested part of the phase diagram, isodiffusivity curves corresponding to low diffusion need to be examined by performing long simulations at each state point. The isodiffusivity line, forming a boundary between the states for which the mean square displacement is lower than the particle dimension over 10^2 s, is called the glass line. Simulations of such long timescales are not achievable, but extrapolation of the other, computationally tractable, lines can be made, and the glass line can be inferred from extrapolation of the isodiffusivity lines. In simulations above the critical point, the extrapolated glass line lies within the high density region, suggesting that the arrest is due to repulsive or excluded volume rather than attractive interactions. The temperature at which the glass line intersects the gas-liquid spinodal is denoted as T_c^{sp} , and determines whether or not the complete phase separation happens on the timescale of 10^2 s. If the temperature of the quench is such that $T_c^{\text{sp}} < T$, the system undergoes a full phase separation into a gas, liquid or a crystal. If $T < T_c^{\text{sp}}$, the arrest or slowing down may occur in the locally dense regions. In Sec. 3.4 we will perform quenches of low density systems close to the temperature T_c^{sp} (Fig. 1.1). This allows us to study the competition between the aggregation of the dense regions and the evolution of local crystallinity within them.

The structure of the gel can be affected by parameters of the system as well as the protocol used to prepare the gel. For example, it is known that low temperatures or strong repulsive amplitudes lead to the formation of gels with thin branches, whereas increasing the temperature or lowering the attraction leads to the thickening of branches [Zhang et al., 2012], and crystallinity within these branches may develop. The gel formation may also proceed slowly such that small aggregates form first. Again, depending on the screening length, and the strength of the attraction or repulsion, the clusters may be spherical or elongated at the early stages of aggregation, and self-assemble into a linear gel of clusters [Sciortino et al., 2005] or hypothetically into a crystal of clusters at the later stages [Sciortino et al., 2004], as will be discussed in Sec. 3.6. The range of attraction may be another parameter affecting the structure of the kinetically slowed down aggregates. Generally, longer

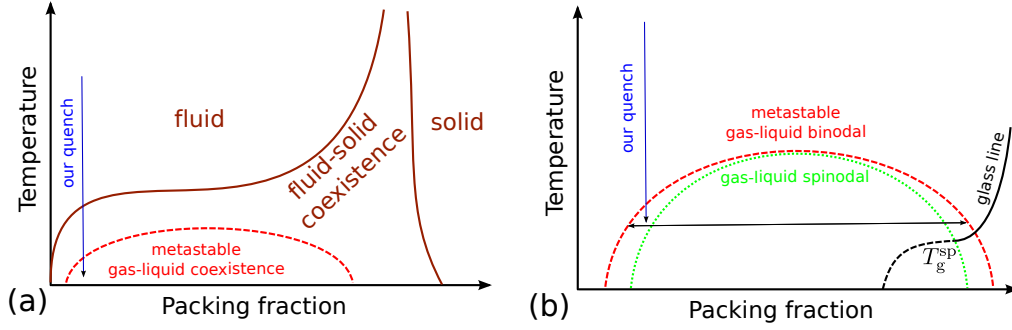


Figure 1.1: Thermodynamic and kinetic phase diagrams of short-range attractive systems. The quenches done in this thesis (Sec. 3.4) will be performed to low densities into (or close to) gas-liquid coexistence region which initiates spinodal decomposition into high and low density regions [Charbonneau and Reichman, 2007a]. (a) The range of attraction of the system considered in this thesis is such that the gas-liquid binodal is hidden below the fluid-solid coexistence line, and metastable density fluctuations precedes the phase separation into stable coexistence of a fluid and a solid [Vliegenthart et al., 1999]. (b) Kinetic phase diagram shows that the quench is done such that the temperature is close to the intersection of the spinodal and the attractive glass line. This means that the bond lifetime is comparable to the lifetime of metastable drops, and we can study how local crystallinity evolves within the amorphous regions [Perez et al., 2011], and how it competes with aggregation of the drops into a single domain.

range of pairwise attraction leads to easier local rearrangements and relaxation, larger phase separation region of the phase diagram, and a more gradual gelation crossover [Zhang et al., 2013]. Nevertheless, it is worth emphasizing that even attractive ranges as wide as 100% of particle diameter can exhibit kinetic arrest at finite temperatures as has been reported by Sastry [2000].

Cates et al. [2004] distinguish between two low packing fractions at which gelation does and does not occur. They suggest that a slow quench below T_c^{sp} performed at a very low density, leads to a local aggregation into unconnected pieces of attractive glass. These aggregates are initially frozen, and remain isolated due to extremely long timescales in further coarsening. The stability of those aggregates was also confirmed in a simulation by Perez et al. [2011], who used fast quenches to temperatures both above and below T_c^{sp} . Results of our simulations for this system and for fast quenches will be discussed in Sec. 3.4. Cates et al. [2004] also discussed products of slow quenches at slightly higher densities, where there are enough particles need to span the entire system. Their prediction of gel formation is based on Mode Coupling Theory [Kroy et al., 2004], which neglects local structure

evolution leading to the phase separation. This approximation, however, turned out to be too crude to predict a realistic behaviour [Lu et al., 2008; Zaccarelli, 2007]. It was thus concluded that it is the phase separation not percolation, which is responsible for gelation at those densities in short-range attractive colloids [Lu et al., 2008].

Finally, we mention more recent insights into the gelation. It has been shown by Whitlam et al. [2009a], that the gelation in undercooled systems does not happen in the absence of collective motion, and is thus a direct consequence of collective dynamical phenomena. The interplay between single particle and collective motion in weakly quenched system is extensively discussed in this thesis. Section 3.6 then addresses the conditions under which collective motion can and cannot be neglected to reproduce a realistic product of the quench. Related to this, we mention that temperature T_c^{SP} lies slightly below the optimum temperature region for self-assembly obtained from fluctuation-dissipation ratios [Jack et al., 2007; Klotsa and Jack, 2013]. In this thesis the temperature of the quench, which roughly corresponds to the low temperature end of the optimum self-assembly region. We also mention, that long timescale simulations of Testard et al. [2011] made the definition of gels more exact, in that the kinetic arrest is not understood as a complete arrest, but rather as a kinetically slowed down phase separation. This thesis uses these two notions interchangeably.

1.5 Scale of colloidal aggregates

When a system consists of monomers, which are divided into isolated partitions (clusters), and if those partitions are stable such that they do not merge for a sufficiently long time, the partitions can be seen as new monomers or as new rigid building blocks of the system, and we may talk about the scale of colloidal aggregates.

Probably the simplest way of realizing such a long-living state of aggregates has been mentioned in Sec. 1.4, and proceeds via a quench of a short-range attractive system [Cates et al., 2004; Lu et al., 2006; Perez et al., 2011] at low densities. The resulting state is guaranteed not to be thermodynamically stable [Charbonneau and Reichman, 2007a], and the stability is only due to the extremely long timescales needed for the partitions to meet and merge [Cates et al., 2004]. A long-range repulsion is believed to be required to ensure their stability [Charbonneau and Reichman, 2007a; Whitlam et al., 2009a]. In fact, recent experiments by Zhang et al. [2012] have suggested that very little is known about the thermodynamic stability

of even dilute short-range attractive and long-range repulsive colloidal dispersions with small aggregates. In this thesis, we provide more evidence that addition of long-range repulsion enhances temporal stability of aggregates and protects them against their coalescence. We show this in Sec. 3.5. We equip our well studied short-range attractive particles with a weak long-range repulsion, and we quench the system to the same conditions as before to study the differences in the kinetics of phase separation (Fig. 1.2).




Figure 1.2: Kinetic and thermodynamic phase diagram of a short-range attractive system with a long-range repulsion based on the diagram of Zaccarelli [2007], and more recent results of Toledano et al. [2009] and Mani et al. [2014]. Long-range repulsion stabilizes the phase separated drops at low densities, which may lead to the formation of thermodynamically stable (spherical) clusters, or at lower temperatures to non-equilibrium long living (elongated) clusters. The non-equilibrium clusters need some time to reorganize into the stable isolated (spherical) clusters [Mani et al., 2014]. The diagram assumes that the transition between the two is related to the intersection of the glass line and the spinodal. The quenches in this thesis (Sec. 3.5) are performed close but above the boundary where equilibrium and non-equilibrium clusters meet. This is because the clusters will be spherical and we will observe the evolution of local crystallinity in our long-living clusters (the quench is above T_g^{sp}). However, since the simulation timescale will not be long enough to verify the thermodynamic stability of the clusters, the quench may also be performed to the temperature region close to the boundary where equilibrium and transient clusters meet.

The first prototype of thermodynamically stable colloidal clusters at low density Stradner et al. [2004] used concentrated lysozyme solutions and charged PMMA particles to show that the size of the aggregates is proportional to the packing fraction, and extended this proposition [Groenewold and Kegel, 2001] to

larger densities. More importantly, the paper of Stradner et al. [2004] showed that the position of the first peak in the static structure factor, corresponding to the average inter-cluster distances, is invariant with respect to the packing fraction. This means that upon a slow reduction of volume of the system, also referred to as crunching [Mossa et al., 2002; Royall and Malins, 2012], the clusters merge and rearrange into larger clusters. The aggregates are highly polydisperse and spherical or elongated, depending on the strength of the attraction or repulsion [Zhang et al., 2012]. The clusters can be considered to constitute isotropic or anisotropic building blocks for a further hierarchical self-assembly.

There are other systems and potential forms exhibiting clustering [Dinsmore et al., 2011]. The ones which are worth mentioning in the context of this thesis are micelles [Wu et al., 1992], star polymers [Stiakakis et al., 2005] or star-linear polymer mixtures [Camargo et al., 2012]. This is because copolymer micelles seem to have a role in mediating the formation and self-assembly of both the single calcite crystal [Kim et al., 2011] and the mesocrystal of the sea urchin spine [Seto et al., 2012]. The link between the cluster phase and mesocrystals will be discussed in Chap. 6.

Simulations of SALR systems are computationally expensive, because they require consideration of all particles within the long-range cutoff. Moreover, to model hierarchical assembly a large number of aggregates is needed, which requires even a larger number of particles in the system. A first systematic simulation study of SALR systems was carried out by Sciortino et al. [2004], who reported a stable fluid of clusters, and arrested low density states of matter at high densities called the Wigner glasses. Toledano et al. [2009] then systematically described the kinetic phase diagram, and established the boundary between the Wigner glass and the gel.

Sciortino et al. [2004] also suggest that the aggregates in the cluster fluid may be spherically renormalized by purely repulsive particles, which would imply the existence of modulated cluster crystals formed of the monodisperse aggregates. However, the formation of modulated cluster solids might not be straightforward due to the intrinsic property of the system to form aggregates that are not entirely monodisperse [Stradner et al., 2004; Charbonneau and Reichman, 2007b]. Formation of a high density modulated phase or a cluster crystal in SALR systems has not been observed experimentally [Zhang et al., 2012] nor in a simulation apart from the work of Xia et al. [2011], who used similar systems where the long-range repulsion was based on a cluster volume dependent renormalization. A similar renormalization will be analysed in Chap. 5. Related to this, we point out that, similarly to the Wigner glass, the building blocks for the cluster crystal do not necessarily need to be monodisperse or spherical aggregates, and that a re-entrant percolation region

in these system exists [Sciortino et al., 2005] allowing clusters to be stabilized at a higher density and thus produce structures with densities closer to the densities of mesocrystals.

Basic properties of phase diagrams in the SALR systems have been investigated by MC and Molecular Dynamics simulation. The existence of inhomogeneous phases at intermediate densities such as columnar and lamellar phases has been reported by de Candia et al. [2006]. A whole region of inhomogeneous phases along with the respective phase transitions have been more closely characterized by Archer and Wilding [2007]. Mossa et al. [2004] used basin hopping techniques to find the ground states of small clusters formed of particles with SALR. Toledano et al. [2009] then produced a collection of small clusters via slow quenches, and showed that their local structure is affected by the long-range forces of other clusters present in the system. This was later investigated in detail by Malins et al. [2011], who suggest that the long-range forces at high packing fractions disturb the internal structure of the clusters so strongly that the aggregates can no longer be considered isolated systems. Charbonneau and Reichman [2007b] showed that fast quenches do not necessarily lead to the formation of stable clusters but rather to their kinetic arrest.

An interesting question is how the range of attraction in SALR affects the properties of the cluster phase. Longer range attraction generally facilitates local relaxation, which implies that thermodynamically stable structures form easier, and non-equilibrium structures live for only a short period of time. For example, de Candia et al. [2006] report formation of columnar and lamellar phases for the range of attraction equal to 20% of the particle diameter. These phases have not been observed when the range is made ten times shorter [Mani et al., 2014]. The SALR systems where the attraction strength is comparable to the particle size are studied only marginally, because the longer-range attraction is more difficult to realize experimentally [Sweatman et al., 2014], and is simulationally more expensive, since the corresponding repulsion range must also be larger. Sweatman et al. [2014] shows that the equilibrium cluster phase is composed of spherical rather than elongated aggregates, which is in agreement with that the relaxation within the aggregates is complete, and that they reach the equilibrium state [Mani et al., 2014]. A distinct and fast relaxation transition from elongated to spherical clusters was also observed in SALR system interacting via standard Lennard-Jones attraction [Xia et al., 2011], but can also be observed in systems interacting via shorter range attraction [Mani et al., 2014], provided the attraction strength is within the range of $6-8k_B T$. This is an energy barrier, which can be occasionally overcome on the timescale of the simulation. Nevertheless, Sastry [2000] suggests that the kinetic arrest may happen

even in systems interacting via a relatively long-range attractions, and it is reasonable to extrapolate this results and expect a kinetic arrest in locally dense regions of respective SALR systems after suitably deep quenches to finite temperatures. These systems are not treated in this thesis.

The phase behaviour of SALR systems has also been analysed analytically. Groenewold and Kegel [2001] used a basic micelle theory to describe clustering of surfactants and showed that the size of the stable clusters can reach several hundreds of particles, and is proportional to the packing fraction at least for low densities. Sear and Gelbart [1999] observed peaks in the static structure factor at the low values of the wave vector using a basic mean field approximation. This gave some evidence of the propensity of SALR systems to form stable cluster fluids, which was later confirmed by comparison of MC simulations and integral equation theories [Broccio et al., 2006]. Pini et al. [2000] showed by a more accurate approximation in the integral equation theory that addition of long-range repulsion to the short-range attraction suppresses the liquid-vapour critical point to lower temperatures, thus increasing the region, where the system experiences large density fluctuations. Moreover, the radial distribution function in their system at intermediate densities suggests the presence of large clusters with size being up to two orders of magnitude larger than the particle size. These results were later successfully compared by Archer et al. [2007], who used mean field density functional theory, and different integral equation theories. Archer et al. [2007], however, were also able to predict the phase behaviour of the SALR systems, where the first peak in the static structure factor diverges at positive values of the wave vector, indicating the presence of phase separation into inhomogeneous phases. These results were then summarized and confirmed by Monte-Carlo simulations [Archer and Wilding, 2007]. The most accurate integral equation theories so far have been developed by Bomont et al. [2010], who have also reported jumps in the structural properties indicating a transition from the uniform fluid to a locally inhomogeneous cluster phase [Bomont et al., 2012]. Finally, Sweatman et al. [2014] used an advanced micelle theory to describe the phase behaviour of SALR system at low densities. The authors report coexistence of a low density cluster vapour phase, and a high density cluster liquid phase, which is a key ingredient for the hierarchical self-assembly of a cluster crystal discussed in Sec. 3.6 and in the final discussion of this thesis. The above analytical models, are not fully comparable with the models used in this study, and can only indicate qualitative trends. The work of Archer and Wilding [2007] uses a longer range repulsion and higher packing fractions. Sweatman et al. [2014] use similar packing fractions but longer range attractions, and the model of Bomont

et al. [2010] has a similar potential form but higher packing fractions.

1.6 Classical crystallization

Classical crystallization theory describes a many body system interacting via complex interactions. The system is assumed to be under thermodynamic conditions corresponding to a first order phase transition. Typically, one phase is a fluid phase, the other is a crystalline phase. The conditions are usually such that the system is drawn out of equilibrium either by undercooling, compression, or by supersaturation of the crystallizing particle species. The crystallizing particles are usually called the monomers. The system then strives for equilibrium by phase separation or crystallization of the monomers, which is a process reducing the free energy. The free energy of formation of a small spherical nucleus with radius r is given by

$$\Delta G(r) = -4/3\pi r^3 \Delta G_V + 4\pi r^2 \gamma, \quad (1.1)$$

where ΔG_V is the difference between the free energy per volume in a crystal and in the fluid phase, and γ is the free energy per unit surface between the two phases. We assume that ΔG_V and γ are positive; this can follow from the definition or from the properties of the system. The free energy of the system is reduced by creating a nucleus of radius r only if $\Delta G(r) < 0$. However, $\Delta G(r)$ is positive for small r with a peak at r_c called the critical radius of the nucleus, and defined as $d\Delta G(r)/dr|_{r=r_c} = 0$. Since $d\Delta G(r)/dr$ is negative for $r < r_c$ and positive for $r > r_c$, a subcritical nucleus tends to dissolve, and nuclei which are larger than the critical size tend to grow.

The classical crystallization theory expects a random fluctuation within the system leading to the generation of critical sized nuclei. The free energy is then reduced by subsequent growth of these nuclei into a single crystal. The theory implicitly assumes that the nuclei or the growing crystals are homogeneous, spherical, and that they adopt the same crystal structure as the final crystal, usually the structure with the lowest free energy. In experiments, the nucleation and growth of the crystal are usually separated by a control parameter with the aim to achieve the desired crystal polymorph. Another important approximation of the classical crystallization theory is that ΔG_V and γ are equilibrium quantities, even though the system is driven significantly out of equilibrium during the crystallization. Quantities ΔG_V and γ can be obtained by simulation [Frenkel and Smit, 2002; Landau and Binder, 2009] or experimentally [Leite and Ribeiro, 2012].

The growth can be described by the time dependence of the monomer flux

through a spherical volume as

$$J(t) = -4\pi r^2 D dC/dr, \quad (1.2)$$

where D is the diffusion coefficient, r is the distance from the centre of the nucleus, and C is the monomer concentration in the solution. This differential equation has an exact solution, if the boundary condition assumes that $C_{ss} - C_s = C_0$ is constant with time. Here C_{ss} is the supersaturated monomer concentration, and C_s is the concentration of particles at the surface of the nucleus. This condition is plausible in highly supersaturated systems or generally in systems where the driving force for nucleation is high. The result of these boundary conditions is that the crystal growth decays in time exponentially, showing that the control of the early stages of growth is crucial to obtain the desired polymorph, shape and size of the crystal. The simulations performed in this thesis aim to capture these stages, even if the spinodal decomposition is the main driving force for phase separation.

Other deviations from the classical crystallization theory can happen, if any of the above assumptions are violated. In particular, if the surface of the growing crystal is not spherical and has well developed crystal surfaces, the free energy per particle is different for particles in the bulk, surfaces, edges and corners of the growing crystal [Barnard and Zapol, 2004]. It will be seen in Sec. 3.5 that the presence of crystal surfaces affects the kinetics of crystallization, and is closely related to the system size of the simulated system.

1.7 Non-classical crystallization

Any realistic crystallization pathway must be kinetically achievable, and must result in a thermodynamically stable state. The competition between kinetics and the thermodynamics is traditionally described by the Ostwald step rule [Ostwald, 1897] or the rule of Stranski and Totomanov [1933]. The former states that the metastable state closest in free energy to the fluid will emerge first, and eventually convert through a series of weak phase transitions to the most stable state. The latter rule is similar and suggests that the state which appears first is the one with the lowest nucleation barrier. Although these rules can be valid in experiments, they are without theoretical foundations and have a limited range of validity in that the liquid phase is likely to precede the crystallization only if the particles rotate slowly, if they must align before they crystallize, and if the specific interactions are just strong enough to stabilize the crystal [Hedges and Whitlam, 2011].

The modern day approach to study nonclassical crystallization is to use the

properties of short-range attractive systems. These systems are easily realisable both experimentally and by simulation, they allow one to vary the position of the gas-liquid critical point with respect to the position of the fluid-solid binodal, by changing the range of attraction. It has been shown by ten Wolde and Frenkel [1997] that for suitably short-ranged attractions, the free energy barrier to crystal formation is lower than in a longer range attractive system. This is due to the presence of density fluctuations near the critical point, which enhance the formation of dense drops in which the crystallinity can easily evolve afterwards. This two-step crystal formation has become a common and popular model to analyse nonclassical crystallization. The kinetics of phase separation in short-range attractive systems is dependent on the protocol of the quench [Royall and Malins, 2012], or on the temperature at which the phase separation occurs [Perez et al., 2011]. The position of the gelation line, introduced earlier, determines the character of the phase transition and the topology of the dense regions. If the temperature is below T_g^{sp} , the phase separation through the dense region is slowed down, and the theory of gelation and kinetic arrest needs to be taken into account. If the temperature is above T_g^{sp} , the liquid drops have a chance to fully recrystallize, and the two step crystallization theory can be studied by computer simulation. In this thesis we will attempt to study the nonclassical crystallization pathways in the vicinity of T_g^{sp} , because crystallization at this temperature has analogies with experimentally observed mesocrystallization pathways in that the evolution of local order competes with aggregation on a larger scale.

Mesocrystallization or oriented attachment constitutes a special class of non-classical crystallization model and proceeds through mechanisms many of which have analogies with the non-classical crystallization observed in short-range attractive models at multiple different scales [Sweatman et al., 2014]. This class of nonclassical crystallization is described in what follows.

1.8 Mesocrystallization

We will now present the main mechanisms leading to mesocrystals, and we shall assume that oriented attachment is a mesocrystallization mechanism. We will attempt to relate those to recent colloidal physics [Bechinger et al., 2013], and kinetically slowed down phase separation [Zaccarelli, 2007], in particular.

First, the collision theory of oriented attachment assumes that the crystalline clusters collide similarly as in a gas. The collisions are of two kinds: effective and non-effective collisions, depending on whether the crystalline lattices are or are not

aligned. The effective collision is highly energetic and irreversible. The non-effective collision is low energetic and reversible. This form of interaction is believed to be responsible for crystal alignment in mesocrystals such as one-dimensional rods of titanium dioxide [Leite and Ribeiro, 2012]. The theory assumes that long-range repulsions are strong enough and compete with short-range attractions (dispersed systems), the phase of stable or metastable clusters can be observed (Sec. 3.5). We also assume that these clusters are identical, i.e. monodisperse and with the same shape and crystal structure. In Sec. 6.1.1 we will discuss possible ways of achieving this assumption by our computer simulation. In experiments, this problem concerns the field of synthesis of nanoparticle colloids, polymorph selection and crystal growth.

The second mechanism of mesocrystal formation is derived for systems which are weakly attractive, and where the long-range repulsion is not dominant or is negligible. Such systems are expected to form an amorphous gel-like aggregate first (Chap. 3), and the crystallinity emerges within that aggregate. The reason for crystallographic alignment of mesocrystals formed within the gels is different for different systems [Cölfen and Antonietti, 2008]. Tang et al. [2006] showed that the combination of dipole moment, positive charge, and hydrophobic interactions drive the system of CdTe into 2D structures resembling the self-assembly of S-layer protein [Whitelam, 2010]. Crystallization of the sea urchin spine [Seto et al., 2012] may also fall into this category, because an amorphous network of calcium carbonate develops first, and the crystallinity and self-organization at the atomistic level develops afterwards within this network.

Another common mechanism of oriented attachment is observed in systems with a good balance between short-range attraction and long-range repulsion. If two particles collide in such a system, they remain attached for a sufficiently long time, such that they can rotate with respect to each other until they align and become a single highly stable crystal. This mechanism is referred to as the grain-rotation-induced grain coalescence mechanism [Moldovan et al., 2002; Moore et al., 2010; Bjerre et al., 2013]. In Chap. 4 we will show how to enhance rotations of large partitions by our simulation, however, we will see that large rotations within a single domain are unlikely, and that the system considered in this thesis might not be suitable to represent this mechanism.

Other possible mechanisms of mesocrystal formation are extensively discussed by Cölfen and Antonietti [2008].

1.9 Simulation methods

Traditional simulation techniques to model the self-assembly and gelation are the Brownian or Molecular dynamics, which are simulation techniques easily achievable by simple adaptation of existing codes which can be parallelized and used to simulate large systems. Such codes are usually not designed to deal with non-spherical potentials, which may be a challenge to overcome [De Michele, 2010].

Standard single-particle Monte Carlo (SPMC) with small displacements is then usually easier to implement, however, the information about the time is not exact and can only be approximated. The approximation is known to work well in an equilibrium fluid [Sanz and Marenduzzo, 2010; Romano et al., 2011], but its behaviour is unexplored in non-equilibrium simulations of phase separation or gelation. One of the main reasons for this is that non-equilibrium regions of the phase diagram yielding good assembly are usually associated with strong clustering, and the whole cluster motion, which is characteristic to this region, cannot be captured by single-particle Monte-Carlo simulation.

This thesis extensively discusses a way of designing efficient Monte Carlo (MC) moves, called the Virtual Move Monte Carlo (VMMC) moves, with the aim to reproduce a kinetically realistic motion observed in such systems. Other similar techniques have been designed, such as the static cluster algorithm [Troisi et al., 2005], or cluster Monte Carlo [Babu et al., 2008]. These techniques do not allow for a type of large collective rearrangements, which seem to play a key role in the relaxation at certain stages of self-assembly. Although not observed in this thesis, the grain-rotation-induced grain coalescence mechanism [Leite and Ribeiro, 2012], leading to epitaxial alignment between the mesocrystalline units is an example.

Apart from easier simulation of complex potentials, other advantages of the VMMC simulation over Brownian dynamics is that it is several times faster [Whitelam and Geissler, 2007], and that it allows one to approximate the hydrodynamic Brownian and Stokesian collective diffusion effects via controlling the size distribution of various sized moving clusters (Sec. 2.6). Hydrodynamics has, indeed, been proven to promote colloidal gelation [Furukawa and Tanaka, 2010].

We stress that this thesis only considers physical Monte-Carlo moves. Those include displacements of single particles and clusters of particles by reasonably sized translations and rotations. Other collective physical moves, such as the chain move [Bernard et al., 2009; Bernard and Krauth, 2012] are not considered here although the relaxation via these moves is crucial in hard sphere systems, and may also be important in dense glassy regions of phase separated systems. Non-physical

moves such as the reflection moves [Liu and Luijten, 2005; Sinkovits et al., 2012] or aggregation volume bias moves [Chen and Siepmann, 2000] are efficient in the exploration of equilibrium energy landscape, but the states sampled in this way do not necessarily need to be kinetically accessible. This thesis only treats kinetically accessible states, and non-kinetic moves which can lead to unphysical states are not considered.

For simplicity, the properties of VMMC are tested and analysed on spherically symmetric potentials in this thesis, although the main use of the algorithm remains in dynamically realistic simulations of anisotropically interacting particles. The spherical potentials can also be easily simulated with standard dynamical methods. A comparison with those methods would be valuable but has not been done here.

1.10 Dynamical fidelity

As mentioned in the previous section, a nearly quantitative match between SPMC and Brownian dynamics can be achieved, if the timestep is scaled with respect to the acceptance probability [Sanz and Marenduzzo, 2010]. Approximation of the dynamics by SPMC is known to work well in equilibrium systems, but fails in undercooled systems at low densities [Whitelam and Geissler, 2007; Whitelam et al., 2009a]. This is because a strong clustering or condensation into liquid drops precedes the crystallization, and SPMC does not capture motion of those drops.

Collective MC moves are thus necessary to generate kinetically realistic trajectories. Approximation of the dynamics by MC simulation in systems drawn out of equilibrium can be crude, but may still be the best and only available approach in non-equilibrium systems interacting via anisotropic potentials and may even lead to realistic products of self-assembly. Examples are self-assembling trajectories of Archaeal Chaperonin [Whitelam et al., 2009b] or S-layer protein [Whitelam, 2010]. An advantage of VMMC is that it selects the moving clusters according to the local forcefields, such that it captures both large scale aggregation, and the evolution of local crystallinity. Another advantage is that VMMC allows to approximate hydrodynamic effects according to the Stokes law which states that differently sized clusters diffuse differently. Basic analytical considerations [Whitelam and Geissler, 2007] lead to the fact that the translational diffusion constant of a spherical cluster of size n_c is proportional to n_c^{-1} in Brownian and to n_c^{-2} in Stokesian dynamics. The generation and acceptance rate of the moving clusters in the VMMC should thus have a corresponding distribution. Nevertheless, the size of proposed clusters

may vary during the course of a non-equilibrium simulation, and is unknown due to the iterative nature of selecting the clusters even in an equilibrium simulation. One of the aims of this thesis is to shed more light on this distribution, and show how it depends on the size of the maximum displacement. Our results imply that the properties of such distribution are crucial in understanding and interpretation of the attempts to associate a timestep with a MC cycle composed of virtual moves [Whitelam and Geissler, 2007; Haxton and Whitelam, 2012]. These attempts are outside the scope of this thesis, and the time through the rest of the thesis is determined by the number of MC sweeps.

We note that the VMMC even allows to capture different diffusive rates of anisotropic clusters in different directions. Moreover, in short-ranged strongly attractive systems, numerical instability of Brownian dynamics does not allow large timesteps, which is a problem not present in the VMMC.

This thesis focuses on how the size of the MC displacement affects the dynamical evolution in supercooled systems. A possible way to determine the optimum size of the maximum displacement is to consider variously sized spherical aggregates with permanent bonds and fixed size, and compare their diffusion constants in the VMMC simulation with the corresponding diffusion constants obtained from Stokesian equations [Haxton and Whitelam, 2012]. The maximum displacement in the VMMC is then chosen such that both diffusion constants match. However, the match may be reached for a wider range of displacements. The optimum maximum displacement is thus chosen as the smallest maximum displacement such that the diffusion is within the agreement of 10% of diffusion obtained from Stokesian equations. This should allow a reasonable acceptance rate of single particles, making local relaxation possible. We note that very small displacements lead to the dominance of SPMC motion and underrepresentation of large scale collective motion within the system, which is why the displacement should be sufficiently large. We also note, that these calibrations are done at equilibrium (similarly to Sanz and Marenduzzo [2010]), and the resulting optimum displacement is used in simulations of non-equilibrium systems [Whitelam and Geissler, 2007; Whitelam et al., 2009a; Haxton and Whitelam, 2012], despite the fact that the optimum may change during the course of the simulation. Sec. 3.4.4 provides details about this change and how it affects the dynamics.

Dynamical pathways of crystallizing systems out of equilibrium generated with VMMC have been compared with Brownian dynamics [Whitelam and Geissler, 2007], in a simulation where a two-dimensional low density system of short-range attractive disks is undercooled below the fluid-solid binodal. The comparison was

merely based on visualizing the configurations in the middle and at the end of the self-assembly, and a good agreement between the VMMC and the Brownian dynamics was found, when compared to the SPMC simulation where the dynamics evolves in an entirely different way. To our knowledge, other direct comparison of the VMMC with dynamical methods has not been done, apart from unpublished tests of Ouldrige et al. [2010], who used molecular dynamics simulation to validate that the VMMC algorithm samples from the Boltzmann distribution.

Chap. 3 investigates the evolution of the trajectory in a qualitative way for different sizes of the displacement in order to investigate the reliability of the dynamical trajectories generated by VMMC. Concomitantly, the analysis provides an insight into the distribution of proposed clusters, which is essential in the association of the VMMC sweep with a physical time. We only perform MC simulations in this thesis, although a direct comparison with dynamical simulations would be desirable. We emphasise that, even if such a comparison is done, it may not provide a reliable answer about the dynamical fidelity of the VMMC, since the trajectories of undercooled systems may be affected by the choice of the timestep in analogy to the choice of the maximum displacement in the VMMC.

In Sec. 3.6 we briefly mention the results of our VMMC simulations, done on a low density system which was well examined by Brownian dynamics [Perez et al., 2011]. The result confirms that the size of the displacement affects the collective motion and the approximation of the dynamics. The trajectories are thus not fully reliable, however, it is shown in Sec. 3.6, that despite these differences, the resulting structures can be similar. This is in agreement with previous simulations of kinetically slowed down phase separation [Foffi et al., 2005a], which show that the resulting structures are independent of whether the molecular or Brownian dynamics are used to generate the trajectories.

Chapter 2

Collective moves for attractive potentials

Virtual move Monte Carlo is a Monte Carlo (MC) cluster algorithm forming clusters via local energy gradients, and approximating the collective kinetic or dynamic motion of attractive colloidal particles. In this chapter (see also Růžička and Allen [2014]) the VMMC algorithm is carefully described, analysed, and tested. To formally validate the algorithm through highlighting its symmetries, alternative and compact ways of selecting and accepting clusters are presented. These additional schemes illustrate the formal use of abstract concepts in the design of biased MC techniques: the superdetailed balance and the early rejection scheme. A brief and comprehensive summary of the algorithms is provided, which makes them accessible without needing to understand the details of the derivation. The algorithms presented in this chapter are not new in their essence, apart from the modification presented in Sec. 2.6, which may be several orders of magnitude faster under specific circumstances.

2.1 Introduction

Virtual move Monte Carlo (VMMC) is a sequel to other cluster algorithms [Swendsen and Wang, 1987; Niedermayer, 1988; Wolff, 1989; Wu et al., 1992; Liu and Luijten, 2004, 2005] for the computer simulation of atomic systems. Its basis is to apply a Monte Carlo (MC) move map to all particles in a cluster. The links between particles, defining the cluster, are created *a priori* by considering the application of the same map to individual particles in the relevant particle pairs. The moving cluster in this algorithm can then be physically interpreted as a region of parti-

cles to which a fluctuation propagates from an initially disturbed particle. VMMC was designed by Whitelam and Geissler [2007], later corrected by Whitelam and Geissler [2008], and a two-step link formation was proposed by the same authors [Whitelam et al., 2009a]. One of its advantages is the ability to approximate, in a coarse-grained fashion, the kinetic evolution of the system, and hence the ability to capture physical (and avoid non-physical) kinetic traps. This can make it resemble Brownian or Stokesian dynamics, if the parameters within the method are chosen with regard to physical quantities, such as size-dependent diffusion coefficients. Although VMMC has begun to be used in the simulation of various aggregation or self-assembly processes [Whitelam et al., 2012; Haxton and Whitelam, 2012; Grant et al., 2011; Russo et al., 2010; Romano et al., 2012; Bieler et al., 2012; Šulc et al., 2012; Fusco and Charbonneau, 2013], we believe that there is a need to present the algorithm in a completely clear form, especially in its symmetrized version. For clarity, and for reasons described below, the Brownian or Stokesian dynamical aspects are not considered here. VMMC can, indeed, be used just to investigate the thermodynamic properties [Romano et al., 2012].

Selection of a MC cluster generally proceeds via a recursive linking of particles that are not yet part of the cluster to the particles that are already in the cluster. As opposed to previously designed MC cluster schemes linking particle pairs with probability depending only on pairwise energy, VMMC creates links with probability depending on pairwise energy gradients corresponding to the direction of the move. This linking approach allows us to approximate a dynamically realistic cluster, but formulation of the microscopic reversibility condition is associated with a problem, because the pairwise energy change in the forward direction between particles in the moving cluster is not the same as the pairwise energy change in the reverse direction. The *unsymmetrized* version of VMMC [Whitelam and Geissler, 2007] solves this issue by recording the forward and reverse linking probabilities between all pairs to which a link is proposed during the cluster selection, and uses them to calculate a complex factor [Whitelam and Geissler, 2008] which biases the probability with which the cluster is accepted. A two-step symmetric linking procedure was later proposed in a so-called *symmetrized* VMMC [Whitelam et al., 2009a], which reduced the set of pairs in the biasing factor to only a subset of pairs in the boundary of the cluster. Both the unsymmetrized and symmetrized versions of the VMMC algorithm were validated through imposing the microscopic reversibility condition on link patterns, called the *realizations* of the cluster. In this chapter, we use the main ideas of the original proofs, but treat them with slightly modified algebra and definitions. We define the realization of the cluster through the set of all pairs to which a

link is proposed, and we discuss whether the same pattern can be selected under the reverse move. Our definitions, basic algebraic relations between pairwise and virtual energies together with the discussion of the order in which the links are proposed, provide an alternative insight into the validation of this abstract algorithm, and can simplify its implementation.

We present three ways of selecting the cluster. The first, general way, called the free cluster selection (Sec. 2.4), attempts to include in the moving cluster all particles of the system that are not yet part of the cluster. This selection procedure is not efficient, and is presented only for theoretical purposes, because its cluster acceptance probability is simple, and can be easily derived. The second way of cluster selection (Secs. 2.5 and 2.6) is newly presented in this chapter. As opposed to that of Sec. 2.4, it only considers certain particles of the system, and is efficient in simulations, especially if proposed clusters are likely to be accepted. The form of the acceptance probability is again simple, and allows us to efficiently control the properties and distribution of proposed clusters. Finally, Sec. 2.8 addresses the original way of selecting the moving cluster [Whitelam et al., 2009a] from an even smaller number of particles. This version of the algorithm is still equivalent to other versions presented here, and is particularly efficient if the acceptance probability of proposed clusters is low.

To rapidly implement the VMMC one may want to skip the lengthy derivations of acceptance probabilities, and focus on the definitions of pairwise and virtual energies in Sec. 2.2, linking probabilities in Sec. 2.3, and use one of the algorithm summaries presented in Secs. 2.5, 2.6 or 2.8. If the acceptance probability is expected to be low, the original symmetrized VMMC [Whitelam et al., 2009a] described in Sec. 2.8 should be fastest. If the acceptance probability is expected to be high, the version described in Sec. 2.5 should be simpler to implement with a negligible difference in efficiency compared to the version of Sec. 2.8. If properties of proposed clusters are controlled by some auxiliary variables, and if unsuitable clusters are expected to have some chance to be accepted, the version presented in Sec. 2.6 may increase the overall acceptance probability, and become more efficient than that presented in Sec. 2.8. Finally, Sec. 2.9 distinguishes between two orders of selecting the cluster. Random ordering is conventional and should work well in most simulations. Isotropic ordering is newly presented here, and may better approximate dynamical clusters, or increase the cluster move acceptance probability, thereby improving the efficiency.

Before implementing the VMMC it should be considered whether the aim is to approximate a physical dynamics, or to enhance sampling of an equilibrium

distribution; also whether collective translational moves are sufficient, and more complicated collective rotational moves can be omitted. It should also be considered whether, and at which stage of self-assembly or aggregation, the collective motion becomes important. In fluids, for example, single particle Monte Carlo is expected to be sufficient both for approximation of the dynamics and for sampling the equilibrium distribution. In simulation of gelation in quenched attractive systems (Fig. 1.1), collective motion becomes dominant only in a narrow simulation window (Sec. 3.4.3). If one has a working VMMC code, some time should be allowed for optimizing the parameters of the VMMC for a specific system. In particular, the optimum choice of the maximum translational or rotational displacement significantly affects the efficiency of the algorithm (Chap. 3). If the dynamics is of interest, properties of proposed clusters need to be controlled, and approximation of the timestep associated with a MC cycle composed of collective moves (appendix of Haxton and Whitlam [2012]) also needs to be taken into account.

Ideas and summaries presented in this chapter can help with the understanding of VMMC. We do not provide any major efficiency enhancements compared to the previous schemes. The aim is to formally describe the algorithm, and provide a unification of underlying concepts which is helpful in the generalization of this algorithm (Sec. 4) to any form of pairwise interaction.

2.2 Definitions and basic relations

In what follows, \mathcal{S} will denote a system of N interacting particles with states being defined by N position (and orientation) vectors of the particles. The interaction between particles is defined by an attractive pairwise potential $-\infty < V(r) \leq +\infty$ with a cut-off distance R_c . The typical interaction is a repulsive core and short-ranged attraction. The inverse temperature of the system is denoted by $\beta = 1/k_B T$, where k_B is the Boltzmann constant. In analogy with other cluster methods, VMMC works with particle pairs (or links between them), denoted by (i, j) , where i and j denote two different particles in \mathcal{S} . The set of all pairs in \mathcal{S} will be denoted as \mathcal{A} . For the purposes of this chapter, it is enough to assume that the set \mathcal{A} does not distinguish between (i, j) and (j, i) . We say that (i, j) is an *interacting pair*, if particles i and j are separated by a distance less than R_c . Throughout the text, μ and ν denote two different states of \mathcal{S} , with energies E_μ and E_ν respectively. Since VMMC is a single-cluster algorithm [Liu and Luijten, 2005], we shall assume that μ and ν differ in the position of a single group of particles called the *moving cluster* or simply *cluster* and denoted by \mathcal{C} . To approximate realistic kinetics, we will

further assume that the difference in the position of \mathcal{C} , is given by the application of a small translational or rotational move map applied to all the particles in the cluster, representing a rigid motion of the cluster. The move map (or briefly map) will be denoted as M and its inverse as M^{-1} . We note that the moving cluster or cluster was referred to as a pseudocluster in the unsymmetrized formulation of the scheme [Whitelam and Geissler, 2007], which used the term physical cluster or simply cluster not for the moving cluster but for what is conventionally referred to as the partition. In this chapter, there is no need to distinguish between the partition and the moving cluster, and the term pseudocluster is thus not used.

The selection of particles from the system \mathcal{S} to the cluster \mathcal{C} proceeds recursively. The first particle of \mathcal{C} is selected randomly from \mathcal{S} . Successive particles of \mathcal{C} are selected by recursive linking of those particles in \mathcal{S} that are not yet part of \mathcal{C} , i.e. we check the existence of a pair (i, j) in state μ such that

$$i \in \mathcal{C}, j \notin \mathcal{C}, \quad (2.1a)$$

$$\text{a link has not yet been proposed to } (i, j), \quad (2.1b)$$

$$\text{an optional condition imposed on } (i, j) \text{ is satisfied.} \quad (2.1c)$$

If there is no pair (i, j) with properties (2.1), the cluster selection is complete. If a pair with properties (2.1) exists, a link is proposed to that pair. The link forms with a probability which will be denoted as $p_{ij}^{(\mu)}$ and defined later. If the link forms, j becomes a new member of \mathcal{C} . The existence of pairs with properties (2.1) is checked recursively until there are no such pairs (the cluster is selected). If the selection of the cluster only follows conditions (2.1a) and (2.1b), we say that the selection is *free*. If there is an extra condition (2.1c) on top of the conditions (2.1a) and (2.1b), we term it *restricted* selection. After the cluster selection, the set of all pairs in \mathcal{S} can then be uniquely decomposed as

$$\mathcal{A} = \mathcal{L}^{(\mu)} \cup \mathcal{F}^{(\mu)} \cup \mathcal{B}^{(\mu)} \cup \mathcal{X}^{(\mu)}, \quad (2.2)$$

where $\mathcal{L}^{(\mu)}$ is the set of pairs to which a link was proposed and formed, $\mathcal{F}^{(\mu)}$ is the set of pairs (i, j) , $i \in \mathcal{C}$, $j \in \mathcal{C}$ to which a link was proposed and failed to form, $\mathcal{B}^{(\mu)}$ is the set of pairs (i, j) , $i \in \mathcal{C}$, $j \notin \mathcal{C}$ to which a link was proposed and failed to form, and $\mathcal{X}^{(\mu)}$ are all other pairs in \mathcal{A} . Set $\mathcal{X}^{(\mu)}$ clearly includes pairs (i, j) , $i \in \mathcal{C}$, $j \in \mathcal{C}$, and depending on the condition (2.1c), pairs (i, j) , $i \notin \mathcal{C}$, $j \notin \mathcal{C}$, or other pairs (i, j) , $i \in \mathcal{C}$, $j \notin \mathcal{C}$, to which a link was not proposed. The triple of sets $(\mathcal{L}^{(\mu)}, \mathcal{F}^{(\mu)}, \mathcal{B}^{(\mu)})$ is called the *realization* of the cluster \mathcal{C} in state μ under the map M , and is denoted

as $R_C^{(\mu)}$. The set $\mathcal{B}^{(\mu)}$ will be called the boundary (of the cluster \mathcal{C}). An example situation for restricted cluster selections is illustrated in Fig. 2.1.

Let us now show for the case of free cluster selection that if $R_C^{(\mu)}$ has a non-zero probability in state μ under M , then there exists a realization of \mathcal{C} in state ν under M^{-1} , denoted as $R_C^{(\nu)}$, such that $R_C^{(\nu)} = R_C^{(\mu)}$. This will allow us to impose the microscopic reversibility condition on generating the realization in states μ and ν . We will need to show not only that each pair in $R_C^{(\mu)}$ can form or fail in state ν , but will also need to show that the recursive selection is such that a link is proposed to that pair under M^{-1} in state ν . Recall that $p_{ij}^{(\mu)}$ denotes the probability of forming a link between pair (i, j) in state μ under M . We will require that $p_{ij}^{(\mu)}$ is defined such that

$$0 \leq p_{ij}^{(\mu)} < 1 \quad \forall (i, j) \ i, j \in \mathcal{S}, \quad (2.3a)$$

$$p_{ij}^{(\mu)} = p_{ij}^{(\nu)} \quad \forall (i, j) \ i, j \in \mathcal{C}, \quad (2.3b)$$

where $p_{ij}^{(\nu)}$ denotes the link formation probability in state ν under M^{-1} . Clearly, if $(i, j) \in \mathcal{L}^{(\mu)}$, then $i, j \in \mathcal{C}$, $p_{ij}^{(\mu)} > 0$, and property (2.3b) implies that the link between (i, j) can also form in state ν , and hence $\mathcal{L}^{(\nu)}$ such that $\mathcal{L}^{(\nu)} = \mathcal{L}^{(\mu)}$ can be generated in state ν with non-zero probability. Property (2.3a) then guarantees that each link which fails to form in μ under M can also fail in ν under M^{-1} , and hence $\mathcal{F}^{(\nu)}$ such that $\mathcal{F}^{(\nu)} = \mathcal{F}^{(\mu)}$, can be generated with non-zero probability in state ν under M^{-1} . The fact that links are proposed to the same set of pairs in state μ under M , and in state ν under M^{-1} follows from the recursive rules (2.1a), (2.1b), and the fact that $\mathcal{L}^{(\nu)}$ has non-zero probability in state ν . It is clear that recursive rules (2.1a), (2.1b) imply that $\mathcal{B}^{(\nu)} = \mathcal{B}^{(\mu)}$. We have thus shown that the recursive rules (2.1a), (2.1b), and linking properties (2.3) imply that if $R_C^{(\mu)}$ in the free cluster selection has non-zero probability in μ , then $R_C^{(\nu)}$ such that $R_C^{(\nu)} = R_C^{(\mu)}$ has non-zero probability in ν .

Conditions (2.1a) and (2.1b) are general rules for the recursive cluster selection, and will be used for the validation of the algorithm. From the practical point of view, considering all pairs satisfying (2.1a) and (2.1b) is not efficient. To speed up the cluster selection, the pairs to which a link is proposed to select the cluster are further restricted by the optional condition (2.1c). A consequence of this restriction will be that the set of pairs to which a link is proposed is not generally the same in states μ and ν . More explicitly, there might be pairs in state ν to which a link is proposed under M^{-1} , but cannot be proposed in state μ under M ; and *vice versa*, there might be pairs to which a link is proposed in state μ under M , but not in state ν under M^{-1} . In Sec. 2.5, the recursive selection of the cluster considers pairs




Figure 2.1: The difference between the original and alternative formulations of the symmetrized VMMC is illustrated on a 2D system \mathcal{S} , consisting of circular particles (disks) in a plane, in two states μ and ν , interacting via a short-range attraction. State ν is created from μ by applying a rotational move map M to the particles in \mathcal{C} (dark disks). Black full lines represent formed links \mathcal{L} . Black dotted lines are failed links internal to the cluster \mathcal{F} , and failed links defining the boundary of \mathcal{C} . (a) Asymmetric boundary. Links are proposed only to pairs interacting in state μ under the forward move M , and to pairs interacting in ν under the reverse move M^{-1} . Grey lines connect the pairs interacting in μ and ν in the respective states μ and ν . (b) Symmetric boundary. Links are proposed to pairs interacting in state μ or in virtual state μ_i under the forward move M , and to pairs interacting in ν or in ν_i under the reverse move M^{-1} . Grey lines connect the pairs interacting in μ or in μ_i and ν or in ν_i , in the respective states μ and ν . Grey disks in the background are particles in their virtual position μ_i in state μ , and ν_i in state ν . From a physical perspective, the root particle is randomly displaced, and the cluster represents a possible region to which this fluctuation propagates.

with properties (2.1a), (2.1b), restricted by (2.1c) such that (i, j) is interacting in μ or in state μ modified by applying M to particle i . It is shown in Appendix A, that the difference between pairs to which a link is proposed does not affect the probability of selecting or accepting the cluster, which will serve as an argument that the acceptance probability of Sec. 2.4 derived for the free cluster selection can also be used in Sec. 2.5 for the cluster selected more efficiently via the restriction (2.1c). Sec. 2.8 considers an even stronger condition (2.1c) restricting the pairs to which a link is proposed, to just pairs that interact in state μ . The difference between the pairs considered in the cluster selection in μ and ν can no longer be ignored: it must be detected after the cluster selection, and included in the bias of the acceptance probability to preserve the exact sampling.

The link formation probability $p_{ij}^{(\mu)}$ will depend on various pairwise energy changes due to the movement by M . It is thus useful to have a compact notation for the energies of pairs of particles before and after one or the other has been moved by M . The energy of (i, j) in state μ is denoted by $\epsilon_{ij}^{(\mu)}$. Clearly, if (i, j) is not an interacting pair in μ then $\epsilon_{ij}^{(\mu)} = 0$. Given a pair (i, j) , the move map M is said to be virtual, if it is applied to a particle only in order to determine the pairwise energy change associated with that move or to determine the set of nearest neighbours of the moved particle. If the system \mathcal{S} is in state μ , and M is applied to particle i , \mathcal{S} is then said to be in the *virtual state* μ_i . Virtual state μ_j will denote a state μ after applying M to particle j ; ν_i will be \mathcal{S} in ν after applying M^{-1} to i ; and ν_j will be \mathcal{S} in ν after applying M^{-1} to j . $\epsilon_{i'j}^{(\mu)}$ will denote the energy of (i, j) in μ_i ; $\epsilon_{ij'}^{(\mu)}$ will be the energy of (i, j) in μ_j ; $\epsilon_{i'j}^{(\nu)}$ will be the energy of (i, j) in ν_i ; and $\epsilon_{ij'}^{(\nu)}$ will be the energy of (i, j) in ν_j . We note that the Boltzmann distribution of states implies that the pairwise energy $\epsilon_{ij}^{(\mu)}$ is strictly different from $+\infty$, and with the above defined properties of the pairwise potential $V(r)$, we can write $-\infty < \epsilon_{ij}^{(\mu)} < +\infty$. The virtual energy, on the other hand, can be assumed $+\infty$. For example, the virtual state μ_i generally has $-\infty < \epsilon_{i'j}^{(\mu)} \leq +\infty$. These inequalities will be used later. We also emphasize that the pairwise energy change is the energy change only between particles i and j , not the energy change of the system. In practice, the energy $\epsilon_{i'j}^{(\mu)}$ is obtained in three steps. Given a pair (i, j) in state μ : firstly, move particle i by M ; secondly, measure pairwise energy of (i, j) , denote it as $\epsilon_{i'j}^{(\mu)}$; thirdly, move particle i back (restore its original position). Energies $\epsilon_{ij'}^{(\mu)}$, $\epsilon_{i'j}^{(\nu)}$, $\epsilon_{ij'}^{(\nu)}$ can be obtained analogously.

Let us define the *complete boundary* as

$$\mathcal{B} = \mathcal{B}^{(\mu)} \cup \mathcal{B}^{(\nu)}, \quad (2.4)$$

Clearly, for the free cluster selection case

$$\mathcal{B} = \mathcal{B}^{(\mu)} = \mathcal{B}^{(\nu)}. \quad (2.5)$$

The following energy identities apply to the classes of pairs identified above. For any pair outside the complete boundary of \mathcal{C} in state μ and ν , it holds that

$$\left. \begin{aligned} \epsilon_{ij}^{(\mu)} &= \epsilon_{ij}^{(\nu)} \\ \epsilon_{i'j}^{(\mu)} &= \epsilon_{i'j}^{(\nu)} \\ \epsilon_{ij'}^{(\mu)} &= \epsilon_{ij'}^{(\nu)} \end{aligned} \right\} \quad \forall (i, j) \notin \mathcal{B}. \quad (2.6)$$

For boundary pairs:

$$\left. \begin{aligned} \epsilon_{i'j}^{(\mu)} &= \epsilon_{ij}^{(\nu)} \\ \epsilon_{ij'}^{(\nu)} &= \epsilon_{ij}^{(\mu)} \end{aligned} \right\} \quad \forall (i, j) \in \mathcal{B}. \quad (2.7)$$

It is useful to denote

$$\begin{aligned} \Delta\epsilon_{i'j}^{(\mu)} &= \epsilon_{i'j}^{(\mu)} - \epsilon_{ij}^{(\mu)}, & \Delta\epsilon_{i'j}^{(\nu)} &= \epsilon_{i'j}^{(\nu)} - \epsilon_{ij}^{(\nu)}, \\ \Delta\epsilon_{ij'}^{(\mu)} &= \epsilon_{ij'}^{(\mu)} - \epsilon_{ij}^{(\mu)}, & \Delta\epsilon_{ij'}^{(\nu)} &= \epsilon_{ij'}^{(\nu)} - \epsilon_{ij}^{(\nu)}, \end{aligned} \quad (2.8)$$

which by using Eqs. (2.7) and (2.6) leads to the identity

$$\Delta\epsilon_{i'j}^{(\mu)} = -\Delta\epsilon_{i'j}^{(\nu)} \quad \forall (i, j) \in \mathcal{B}, \quad (2.9)$$

and identities

$$\left. \begin{aligned} \Delta\epsilon_{i'j}^{(\mu)} &= \Delta\epsilon_{ij'}^{(\nu)} \\ \Delta\epsilon_{ij'}^{(\nu)} &= \Delta\epsilon_{i'j}^{(\mu)} \end{aligned} \right\} \quad \forall (i, j) \notin \mathcal{B}. \quad (2.10)$$

We will now use the above notation in order to express a trivial fact, that by moving a cluster of particles, the energy change of the system is given by the energy change of the boundary of the cluster. Relation (2.2) implies that the total energy of the system in state μ can then be expressed as

$$E_\mu = \sum_{(i,j) \in \mathcal{L}^{(\mu)}} \epsilon_{ij}^{(\mu)} + \sum_{(i,j) \in \mathcal{F}^{(\mu)}} \epsilon_{ij}^{(\mu)} + \sum_{(i,j) \in \mathcal{B}^{(\mu)}} \epsilon_{ij}^{(\mu)} + \sum_{(i,j) \in \mathcal{X}^{(\mu)}} \epsilon_{ij}^{(\mu)}. \quad (2.11)$$

Similarly, the total energy of \mathcal{S} in state ν is

$$E_\nu = \sum_{(i,j) \in \mathcal{L}^{(\nu)}} \epsilon_{ij}^{(\nu)} + \sum_{(i,j) \in \mathcal{F}^{(\nu)}} \epsilon_{ij}^{(\nu)} + \sum_{(i,j) \in \mathcal{B}^{(\nu)}} \epsilon_{ij}^{(\nu)} + \sum_{(i,j) \in \mathcal{X}^{(\nu)}} \epsilon_{ij}^{(\nu)}. \quad (2.12)$$

Since $R_{\mathcal{C}}^{(\mu)} = R_{\mathcal{C}}^{(\nu)}$, and $\epsilon_{ij}^{(\mu)} = \epsilon_{ij}^{(\nu)}$, provided (i, j) is in any of the sets $\mathcal{L}^{(\mu)} = \mathcal{L}^{(\nu)}$,

$\mathcal{F}^{(\mu)} = \mathcal{F}^{(\nu)}$, or $\mathcal{X}^{(\mu)} = \mathcal{X}^{(\nu)}$, by using Eq. (2.8) and Eq. (2.5) we can express

$$E_\nu - E_\mu = \sum_{(i,j) \in \mathcal{B}} \Delta \epsilon_{ij}^{(\mu)}. \quad (2.13)$$

2.3 Link formation probabilities

We have briefly explained the way in which particles, or rather pairs, are selected for consideration in the construction of the cluster, but we have not given any expressions for the acceptance or rejection of the link proposals, and this is the topic of the current section. After the cluster is constructed, we anticipate accepting or rejecting the Monte Carlo cluster move with a certain probability. The expression for this can be significantly simplified, if the link formation probability is the same for the forward and reverse moves, for pairs outside the boundary of the cluster. Let us thus express condition (2.3b) imposed on the link formation probability in terms of the complete boundary as

$$p_{ij}^{(\mu)} = p_{ij}^{(\nu)} \quad \forall (i,j) \notin \mathcal{B}. \quad (2.14)$$

Clearly, if (2.14) is satisfied, then (2.3b) is satisfied. The requirement (2.14) can be easily achieved in conventional (static) cluster algorithms, because $p_{ij}^{(\mu)}$ only depends on the energy of (i,j) in state μ , and $p_{ij}^{(\nu)}$ only depends on the energy of (i,j) in state ν , and these energies are clearly the same for pairs outside the boundary [see the first identity in Eq. (2.6) which applies for $(i,j) \notin \mathcal{B}$]. In contrast, $p_{ij}^{(\mu)}$ in (dynamic) VMMC also depends on the energy of (i,j) in state ν . To satisfy relation (2.14) in the dynamical linking process is not so straightforward, because it is not known whether (i,j) belongs to the boundary of \mathcal{C} , during the construction of \mathcal{C} . Hence, it is assumed that during the cluster selection (i,j) does not belong to the boundary, and this assumption is corrected by a bias in the cluster move acceptance probability for pairs which end up in the boundary. Moreover, in order to apply the early rejection scheme [Frenkel and Smit, 2002], which simplifies and accelerates the algorithm, the link is formed in two link formation tests as follows.

Given a pair (i,j) in state μ , we consider the following two probabilities defined as

$$p_{ij}^{(\mu)} = \max \left\{ 0, 1 - \exp \left(-\beta \Delta \epsilon_{ij}^{(\mu)} \right) \right\}, \quad (2.15a)$$

$$p_{ij'}^{(\mu)} = \max \left\{ 0, 1 - \exp \left(-\beta \Delta \epsilon_{ij'}^{(\mu)} \right) \right\}. \quad (2.15b)$$

The former represents the pre-link formation probability between (i, j) in state μ under the map M . The latter represents the pre-link formation probability in ν under M^{-1} ; note, we assume that particles are linked in the forward move, and this interpretation formally uses the equality (2.10). The link between (i, j) then forms in two random tests. The success probability of the first test is merely given by $p_{ij}^{(\mu)}$, and that of the second is given by $\min\{1, p_{ij'}^{(\mu)}/p_{i'j}^{(\mu)}\}$. The link between (i, j) forms only if both tests are successful. Hence, the link formation probability can be defined as the joint probability

$$p_{ij}^{(\mu)} = \text{Prob} \left\{ X_1 < p_{ij}^{(\mu)}; X_2 < \min(1, p_{ij'}^{(\mu)}/p_{i'j}^{(\mu)}) \right\}, \quad (2.16)$$

where (X_1, X_2) are two independent random numbers drawn from a uniform distribution $\mathcal{U}(0, 1)$. If the first test fails, then (i, j) is marked as an *outright failed* link (or pair) independently of the outcome of the second test. If the first test succeeds and the second test fails, then (i, j) is marked as a *frustrated* link (or pair). Note that, if $p_{i'j}^{(\mu)} = 0$, we do not need to consider the ill-defined fraction $p_{ij'}^{(\mu)}/p_{i'j}^{(\mu)}$, as (i, j) is always outright failed.

The link failure probability is defined as

$$q_{ij}^{(\mu)} = 1 - p_{ij}^{(\mu)}. \quad (2.17)$$

It is useful to define the pre-link failure probability as

$$q_{i'j}^{(\mu)} = 1 - p_{i'j}^{(\mu)}, \quad (2.18)$$

and express it as

$$q_{i'j}^{(\mu)} = \exp \left(-\beta \max\{0, \Delta\epsilon_{i'j}^{(\mu)}\} \right). \quad (2.19)$$

We point out that, in the wider context of recursive cluster selection, the link formation probability defined in this way is conditional upon μ , and upon a set of auxiliary variables $\{b\}$ comprising M , and (optionally) a limit specifying the maximum size of the cluster. We will later discuss the case where $p_{i'j}^{(\mu)}$ explicitly depends upon the current number of particles in the cluster. Note also that definitions (2.15a), (2.15b), and (2.18), use a notation which is slightly different (but more accurate) than that used in the original papers [Whitelam et al., 2009a]. It shows that the pre-link probability is not only conditional upon M (through the appearance of i' and j' in the above expressions) but, more importantly, it also shows that the probability is conditional upon the state μ itself.

Finally, let us show that the link formation probability defined in Eq. (2.16)

does satisfy the desired equality (2.14). Firstly, the equalities (2.10) imply that

$$\left. \begin{aligned} p_{i'j}^{(\mu)} &= p_{ij'}^{(\nu)} \\ p_{ij'}^{(\mu)} &= p_{i'j}^{(\nu)} \end{aligned} \right\} \quad \forall (i, j) \notin \mathcal{B}. \quad (2.20)$$

Hence, if both $p_{i'j}^{(\mu)} > 0$ and $p_{ij'}^{(\mu)} > 0$, then both $p_{i'j}^{(\nu)} > 0$ and $p_{ij'}^{(\nu)} > 0$, and the link formation probability (2.16) can be expressed as

$$\begin{aligned} p_{ij}^{(\mu)} &= p_{i'j}^{(\mu)} \min \left\{ 1, \frac{p_{ij'}^{(\mu)}}{p_{i'j}^{(\mu)}} \right\} = \min \left\{ p_{ij'}^{(\mu)}, p_{i'j}^{(\mu)} \right\} \\ &= \min \left\{ p_{i'j}^{(\nu)}, p_{ij'}^{(\nu)} \right\} = p_{i'j}^{(\nu)} \min \left\{ 1, \frac{p_{ij'}^{(\nu)}}{p_{i'j}^{(\nu)}} \right\} = p_{ij}^{(\nu)} \end{aligned} \quad (2.21)$$

for $\forall (i, j) \notin \mathcal{B}$, where equality (2.20) was used again. Similarly, if $p_{i'j}^{(\mu)} = 0$ or $p_{ij'}^{(\mu)} = 0$, then $p_{i'j}^{(\nu)} = 0$ or $p_{ij'}^{(\nu)} = 0$, and $p_{ij}^{(\mu)} = p_{ij}^{(\nu)} = 0$. This completes the proof that the definition (2.16) of $p_{ij}^{(\mu)}$ satisfies Eq. (2.14) or Eq. (2.3b).

For the sake of clarity, let us express the probability that a link between (i, j) does not form. This probability is given by the probability that (i, j) is outright failed or that (i, j) is frustrated which reads as

$$\begin{aligned} q_{i'j}^{(\mu)} + p_{i'j}^{(\mu)} \left(1 - \min \left\{ 1, p_{ij'}^{(\mu)} / p_{i'j}^{(\mu)} \right\} \right) &= q_{i'j}^{(\mu)} + p_{i'j}^{(\mu)} - \min \left\{ p_{i'j}^{(\mu)}, p_{ij'}^{(\mu)} \right\} \\ &= 1 - \min \left\{ p_{i'j}^{(\mu)}, p_{ij'}^{(\mu)} \right\} = 1 - p_{ij}^{(\mu)} = q_{ij}^{(\mu)}. \end{aligned}$$

2.4 Free cluster selection

In this section we derive a cluster move acceptance probability for the algorithm, which proceeds through free cluster selection. Selecting the cluster in this general way is inefficient, but the corresponding acceptance probability can be easily derived, because the set of pairs to which a link is proposed during the cluster selection does not differ in the old and in the new states. We will thus refer to the realization of the cluster to have both a *symmetric core* ($\mathcal{L}^{(\mu)} = \mathcal{L}^{(\nu)}$, $\mathcal{F}^{(\mu)} = \mathcal{F}^{(\nu)}$) and a *symmetric boundary* ($\mathcal{B}^{(\mu)} = \mathcal{B}^{(\nu)}$). The other significantly more efficient ways of selecting the cluster through asymmetric realizations are presented later. In what follows, we will assume that the cluster is selected freely, i.e. by considering all pairs (i, j) satisfying Eqs. (2.1a) and (2.1b).

We will use the two-step link formation probability defined in the previous

section to link particles to the cluster, which allows us to define an early rejection scheme simplifying the cluster move acceptance probability.

Let us only consider those $R_{\mathcal{C}}^{(\mu)}$ such that $\mathcal{B}^{(\mu)}$ consists entirely of outright failed links. We anticipate imposing the condition that all realizations of \mathcal{C} with at least one frustrated link in the boundary are rejected, and it is shown in Appendix C.1 that the scheme samples from the Boltzmann distribution, subject to this condition.

Since the linking probability $p_{ij}^{(\mu)}$ defined in the previous section satisfies conditions (2.3), for each $\mathcal{L}^{(\mu)}$ and $\mathcal{F}^{(\mu)}$, there is a non-zero chance to recursively obtain $\mathcal{L}^{(\nu)}$ and $\mathcal{F}^{(\nu)}$, with $\mathcal{L}^{(\mu)} = \mathcal{L}^{(\nu)}$ and $\mathcal{F}^{(\mu)} = \mathcal{F}^{(\nu)}$. Let us now show that if $\mathcal{B}^{(\mu)}$ is a corresponding boundary with outright failed links in state μ under M , then $\mathcal{B}^{(\nu)}$ can also be a boundary with outright failed pairs in state ν under M^{-1} . Indeed, $-\infty < \epsilon_{ij}^{(\mu)} = \epsilon_{i'j}^{(\nu)} < +\infty$, and the fact that $p_{ij}^{(\mu)} < 1$, implies that $p_{i'j}^{(\nu)} < 1$, and an outright failed pair in μ under M can become an outright failed boundary pair in ν under M^{-1} . We can thus expect that if $R_{\mathcal{C}}^{(\mu)}$ is such that $\mathcal{B}^{(\mu)}$ only has outright failed pairs, a realization $R_{\mathcal{C}}^{(\nu)}$ with $\mathcal{B}^{(\nu)}$ having outright failed pairs can be constructed with non-zero probability in ν under M^{-1} . In total, we have that for each $R_{\mathcal{C}}^{(\mu)}$ with outright failed boundary pairs, there is a non-zero chance to construct an $R_{\mathcal{C}}^{(\nu)}$ with outright failed boundary pairs such that $R_{\mathcal{C}}^{(\mu)} = R_{\mathcal{C}}^{(\nu)}$. This will be used to simplify the following.

Let us impose the microscopic reversibility condition on generating the realizations $R_{\mathcal{C}}^{(\mu)}$ and $R_{\mathcal{C}}^{(\nu)}$ with outright failed boundary pairs. This condition can also be seen as the superdetailed balance (SDB) condition [Frenkel and Smit, 2002; Frenkel, 2004] imposed on the realization of the cluster in states μ and ν . The condition reads as

$$\begin{aligned} & \exp(-\beta E_{\mu}) p^{(\mu)}(\{b\}) \prod_{(i,j) \in \mathcal{L}^{(\mu)}} p_{ij}^{(\mu)} \prod_{(i,j) \in \mathcal{F}^{(\mu)}} q_{ij}^{(\mu)} \prod_{(i,j) \in \mathcal{B}^{(\mu)}} q_{i'j}^{(\mu)} W_{\text{acc}}^{(\mu \rightarrow \nu | R)} \\ &= \exp(-\beta E_{\nu}) p^{(\nu)}(\{b\}) \prod_{(i,j) \in \mathcal{L}^{(\nu)}} p_{ij}^{(\nu)} \prod_{(i,j) \in \mathcal{F}^{(\nu)}} q_{ij}^{(\nu)} \prod_{(i,j) \in \mathcal{B}^{(\nu)}} q_{i'j}^{(\nu)} W_{\text{acc}}^{(\nu \rightarrow \mu | R)}. \end{aligned} \quad (2.22)$$

The leftmost terms on each side are the Boltzmann weights of states μ and ν . The term $p^{(\mu)}(\{b\})$ is the probability of generating auxiliary variables $\{b\}$ in state μ , and it is a product of three components

$$p^{(\mu)}(\{b\}) = p^{(\mu)}(M) p^{(\mu)}(i) p^{(\mu)}(N_{\mathcal{C}}), \quad (2.23)$$

specifically the probability of selecting the move map M , the root particle i , and the maximum cluster size N_C in state μ . The term $p^{(\nu)}(\{b\})$ on the right represents a similar product in state ν . The explicit products over links on the left of Eq. (2.22) combine to give the probability of constructing the specific realization $R_C^{(\mu)}$, which is conditional upon $\{b\}$. Analogous terms for state ν appear on the right. Finally, the terms $W_{\text{acc}}^{(\mu \rightarrow \nu | R)}$ and $W_{\text{acc}}^{(\nu \rightarrow \mu | R)}$ are the cluster move acceptance probabilities that we seek.

We note that the set of all pairs satisfying conditions (2.1a) and (2.1b) generally changes during the cluster selection, and that even if we proceed through the free cluster selection, the SDB condition in Eq. (2.22) may be seen not to express exactly what the algorithm is doing. In Sec. 2.10 we will take the change of set with (2.1a) and (2.1b) into account, and we will distinguish between the probability with which a pair (i, j) is selected, and the probability with which a link between (i, j) is formed. This will allow us to formulate the SDB condition in a slightly more explicit way than in Eq. (2.22).

Let us simplify the SDB condition (2.22). We suppose being able to ensure that $p^{(\mu)}(\{b\}) = p^{(\nu)}(\{b\}) > 0$, so that these terms cancel out. The products over linked pairs $\mathcal{L}^{(\mu)}$, $\mathcal{L}^{(\nu)}$, and unlinked pairs $\mathcal{F}^{(\mu)}$, $\mathcal{F}^{(\nu)}$ also cancel because of Eq. (2.14), the fact that these sets are the same in states μ and ν , and the fact that the probability of generating $R_C^{(\nu)}$ such that $R_C^{(\nu)} = R_C^{(\mu)}$ is strictly positive. Hence, Eq. (2.22) thus involves only products over boundary pairs $\prod_{\mathcal{B}^{(\mu)}}$ and $\prod_{\mathcal{B}^{(\nu)}}$ on each side. The reduced form of the SDB condition then reads as

$$\exp(-\beta E_\mu) \prod_{(i,j) \in \mathcal{B}^{(\mu)}} q_{i'j}^{(\mu)} W_{\text{acc}}^{(\mu \rightarrow \nu | R)} = \exp(-\beta E_\nu) \prod_{(i,j) \in \mathcal{B}^{(\nu)}} q_{i'j}^{(\nu)} W_{\text{acc}}^{(\nu \rightarrow \mu | R)}, \quad (2.24)$$

with a Metropolis (Rosenbluth factor) type of solution

$$W_{\text{acc}}^{(\mu \rightarrow \nu | R)} = \min \left\{ 1, \exp(-\beta(E_\nu - E_\mu)) \frac{\prod_{(i,j) \in \mathcal{B}^{(\nu)}} q_{i'j}^{(\nu)}}{\prod_{(i,j) \in \mathcal{B}^{(\mu)}} q_{i'j}^{(\mu)}} \right\}. \quad (2.25)$$

This expression can be simplified as follows. Let us separate the products over pairs into those for which $\Delta\epsilon_{ij}^{(\mu)} > 0$, denoted by a superscript $+$, and those for which $\Delta\epsilon_{ij}^{(\mu)} \leq 0$, denoted by $-$. Using expression (2.19), and relations (2.5), (2.9)

and (2.13), the ratio of products in Eq. (2.25) can be expressed as

$$\begin{aligned} \frac{\prod_{(i,j) \in \mathcal{B}^{(\nu)}} q_{i'j}^{(\nu)}}{\prod_{(i,j) \in \mathcal{B}^{(\mu)}} q_{i'j}^{(\mu)}} &= \frac{\prod_{(i,j) \in \mathcal{B}}^+ \exp(-\beta \Delta \epsilon_{i'j}^{(\nu)})}{\prod_{(i,j) \in \mathcal{B}}^+ \exp(-\beta \Delta \epsilon_{i'j}^{(\mu)})} = \prod_{(i,j) \in \mathcal{B}}^- \exp(\beta \Delta \epsilon_{i'j}^{(\mu)}) \prod_{(i,j) \in \mathcal{B}}^+ \exp(\beta \Delta \epsilon_{i'j}^{(\mu)}) \\ &= \prod_{(i,j) \in \mathcal{B}} \exp(\beta \Delta \epsilon_{i'j}^{(\mu)}) = \exp(\beta(E_\nu - E_\mu)), \quad (2.26) \end{aligned}$$

and the cluster move acceptance probability (2.25) can be simplified to

$$W_{\text{acc}}^{(\mu \rightarrow \nu | R)} = 1, \quad (2.27a)$$

provided all pairs in the boundary \mathcal{B} are outright failed, and as

$$W_{\text{acc}}^{(\mu \rightarrow \nu | R)} = 0, \quad (2.27b)$$

if there is a frustrated pair in \mathcal{B} . This means that the cluster is accepted whenever the boundary only contains outright failed links and is rejected otherwise. This form of the acceptance probability is similar to the original rejection-free cluster algorithms [Wolff, 1989; Liu and Luijten, 2004], and represents an example where the early rejection scheme is useful when applied to many-particle moves [Frenkel and Smit, 2002].

2.5 Selection of the cluster under high acceptance probability

This section formulates the version of symmetrized VMMC, which has a very simple form of the cluster acceptance probability, and which is efficient, in particular, if we expect the cluster acceptance probability to be high. Sec. 2.6 will show that this simpler formulation forms a basis for a further alteration of the algorithm which allows us to control properties of proposed clusters, without rejecting them *in situ*. Compared to the free cluster selection presented in Sec. 2.4, the speed is gained by restricting the pairs to which a link is proposed, from the set of pairs satisfying conditions (2.1a) and (2.1b) to the set of pairs satisfying conditions (2.1a), (2.1b), and (2.1c), where (2.1c) requires that (i, j) is interacting in μ or in μ_i . This is a natural restriction, because if (i, j) is not interacting in μ and in μ_i , then $\epsilon_{ij}^{(\nu)} =$

$\epsilon_{i'j}^{(\mu)} = 0$, and the definition of linking probability (2.15a) implies that a link proposed to (i, j) is outright failed in states μ under M . If (i, j) ends up outside the boundary, Eqs. (2.20) and (2.21) imply that (i, j) fails to form in state ν under M^{-1} . If (i, j) ends up in the boundary, Eqs. (2.7) imply that $p_{i'j}^{(\nu)} = 0$, and that (i, j) is outright failed in state ν . There is thus no reason to propose a link to that pair in the free cluster selection.

As discussed earlier for the SDB condition, the set of pairs to which a link is proposed in states μ and ν must be guaranteed to be the same to ensure the existence of $R_{\mathcal{C}}^{(\nu)}$ such that $R_{\mathcal{C}}^{(\nu)} = R_{\mathcal{C}}^{(\mu)}$. It is shown in Appendix A that although these sets are not the same under the restriction (2.1c) of this section, nothing changes in the algorithm if we assume that the sets are identical, and we can thus use the simple form of the cluster acceptance probability (2.27) for the free cluster selection. Since generally $\mathcal{F}^{(\mu)} \neq \mathcal{F}^{(\nu)}$ and $\mathcal{B}^{(\mu)} = \mathcal{B}^{(\nu)}$, we say that the realization of the cluster has *asymmetric core* and *symmetric boundary*.




Figure 2.2: Illustration of the asymmetric core on a system \mathcal{S} composed of attractive disks with range of interaction r_c determined by the radius of the disk. The figure shows a moving cluster composed of three particles selected via the translational map M in state μ , and by its inverse M^{-1} in state ν . The root particle is denoted as i_r . Particles in the virtual states corresponding to M in μ and M^{-1} in ν are green (with short-dashed margin). Full lines between particles represent formed links \mathcal{L} . Dashed line is the failed link in $\mathcal{F}^{(\mu)}$. The formed and failed links are enumerated according to the order they were proposed. The link between the first two top particles was not proposed in state ν , because the two particles do not interact both in states ν and ν_i . Thus $\mathcal{F}^{(\nu)} = \emptyset$, $\mathcal{F}^{(\mu)} \neq \mathcal{F}^{(\nu)}$, and the core is asymmetric.

The cluster move derived in Sec. 2.4 and accelerated by condition (2.1c) defined in this section can be summarized as follows:

1. Pick a random particle, and use it as the first (root) particle of the cluster \mathcal{C} .
2. Perform the recursive loop selecting all other particles to \mathcal{C} .
 - (a) Pick randomly a pair (i, j) , $i \in \mathcal{C}$, $j \notin \mathcal{C}$, which interacts in state μ or in virtual state μ_i , and to which a link has not yet been proposed. If no such pair exists, finish the cluster selection by exiting the recursive loop.

(b) Attempt to create a link between (i, j) as follows.

i. Form a pre-link with probability

$$p_{ij}^{(\mu)} = \max \left\{ 0, 1 - \exp(-\beta(\epsilon_{ij}^{(\mu)} - \epsilon_{ij'}^{(\mu)})) \right\} .$$

ii. If the pre-link does not form, label (i, j) as outright failed, go to (a).

iii. If the pre-link forms, calculate the reverse link formation probability

$$p_{ij'}^{(\mu)} = \max \left\{ 0, 1 - \exp(-\beta(\epsilon_{ij'}^{(\mu)} - \epsilon_{ij}^{(\mu)})) \right\} ,$$

where $\epsilon_{ij'}^{(\mu)}$ denotes the energy of (i, j) in μ after applying the move map M only to j , and form the link with probability

$$\min \left\{ 1, p_{ij'}^{(\mu)} / p_{ij}^{(\mu)} \right\} .$$

(c) If the link does not form, label (i, j) as frustrated, go to (a).

(d) If the link forms, include j into \mathcal{C} , go to (a).

3. Identify the boundary \mathcal{B} of cluster \mathcal{C} as those pairs (i, j) , $i \in \mathcal{C}$, $j \notin \mathcal{C}$, to which a link was proposed but failed to form, i.e. is either outright failed or frustrated.
4. Accept the cluster move, provided the boundary \mathcal{B} only contains outright failed links. If \mathcal{B} contains a frustrated link, the move of \mathcal{C} is rejected.

2.6 Controlling the cluster size

Cluster size control can be useful when generating small rotational Monte Carlo cluster moves, intended to preserve kinetic evolution of the system, or to represent the dynamics of Stokesian or Brownian particle motion [Whitelam, 2011], or in conditions of high particle density, where the moving cluster might otherwise span the entire system [Liu and Luijten, 2005]. The original symmetrized or unsymmetrized VMMC schemes control the size distribution of the moving clusters simply by rejecting any move reaching a certain limit (number of particles in the cluster is an example of such a limit). In the symmetrized version of VMMC this leads to a combination of two early rejection schemes; one by frustrated links in the boundary, the other by a maximum cluster size. Here, we present a way of accepting the clusters which would otherwise have been rejected by the latter scheme. This makes this ap-

proach potentially more efficient than the original symmetrized VMMC [Whitelam et al., 2009a] summarized in Sec. 2.8.

Let us assume that the auxiliary conditions $\{b\}$, chosen at the beginning of the cluster construction, include a number N_C determining the maximum number of particles in the chosen cluster. We now redefine the pre-link formation probabilities in state μ as

$$\left. \begin{aligned} p_{i'j}^{(\mu)} &= \max \left[0, 1 - \exp(-\beta \Delta \epsilon_{i'j}^{(\mu)}) \right] \\ p_{ij'}^{(\mu)} &= \max \left[0, 1 - \exp(-\beta \Delta \epsilon_{ij'}^{(\mu)}) \right] \end{aligned} \right\} \quad \text{if } n_C < N_C, \quad (2.28a)$$

where n_C denotes the number of particles in \mathcal{C} , and

$$p_{i'j}^{(\mu)} = 0 \quad \text{if } n_C = N_C. \quad (2.28b)$$

The link again forms as a result of two random tests, and the realization $R_C^{(\mu)}$ is constructed recursively from the root as described above. Hence, when the N_C -th particle is linked to the cluster, the selection is finished, and all pairs (i, j) , $i \in \mathcal{C}$, $j \notin \mathcal{C}$ interacting in one of the states μ or μ_i , and to which a link has not yet been proposed, will become boundary pairs with probability $q_{ij}^{(\mu)} = 1$.

To derive the acceptance probability, let N_C be a random integer, $1 \leq N_C \leq N$, and let $R_C^{(\mu)} = (\mathcal{L}^{(\mu)}, \mathcal{F}^{(\mu)}, \mathcal{B}^{(\mu)})$ be a realization after selecting \mathcal{C} via the free cluster selection following the recursive rules (2.1a) and (2.1b), and using the probability defined in Eq. (2.28). The boundary can be uniquely decomposed as

$$\mathcal{B}^{(\mu)} = \mathcal{B}^{*(\mu)} \cup \mathcal{B}^{\dagger(\mu)}, \quad (2.29)$$

where $\mathcal{B}^{*(\mu)}$ denotes pairs in $\mathcal{B}^{(\mu)}$ which failed by definition (2.28a), i.e. failed probabilistically, and $\mathcal{B}^{\dagger(\mu)}$ denotes pairs which failed by definition (2.28b), i.e. were forced to fail because the cluster reached its maximum size. Since $p_{ij}^{(\mu)}$ defined in Eqs. (2.28) satisfies conditions (2.3a) and (2.3b), one can construct $R_C^{(\nu)}$ such that $R_C^{(\nu)} = R_C^{(\mu)}$. Moreover, if the notation (2.29) is used in state ν , it is clear $R_C^{(\nu)}$ can be such that

$$\mathcal{B}^* = \mathcal{B}^{*(\mu)} = \mathcal{B}^{*(\nu)}, \quad (2.30a)$$

$$\mathcal{B}^\dagger = \mathcal{B}^{\dagger(\mu)} = \mathcal{B}^{\dagger(\nu)}, \quad (2.30b)$$

where sets \mathcal{B}^* and \mathcal{B}^\dagger have been defined.

We again use the SDB condition (2.22) to express the acceptance probability

as in Eq. (2.27). Eq. (2.27) can then be further simplified by considering relations (2.13), (2.29) which imply

$$E_\nu - E_\mu = \sum_{(i,j) \in \mathcal{B}} \Delta\epsilon_{i'j}^{(\mu)} = \sum_{(i,j) \in \mathcal{B}^*} \Delta\epsilon_{i'j}^{(\mu)} + \sum_{(i,j) \in \mathcal{B}^\dagger} \Delta\epsilon_{i'j}^{(\mu)}. \quad (2.31)$$

In a similar way to section 2.4, we further decompose the sums and products over pairs in the boundary \mathcal{B}^* into those for which $\Delta\epsilon_{i'j}^{(\mu)} > 0$ (denoted by $+$) and those for which $\Delta\epsilon_{i'j}^{(\mu)} \leq 0$ (denoted by $-$). The Boltzmann factor of the energy difference in Eq. (2.25) can be expressed as

$$\exp(-\beta(E_\nu - E_\mu)) = \frac{\prod_{(i,j) \in \mathcal{B}^*}^+ \exp(\beta\Delta\epsilon_{i'j}^{(\nu)})}{\prod_{(i,j) \in \mathcal{B}^*}^+ \exp(\beta\Delta\epsilon_{i'j}^{(\mu)})} \prod_{(i,j) \in \mathcal{B}^\dagger} \exp(-\beta\Delta\epsilon_{i'j}^{(\mu)}). \quad (2.32)$$

Using definitions (2.29), (2.28), and relation (2.9), the product in Eq. (2.25) can then be expressed as

$$\prod_{(i,j) \in \mathcal{B}} \frac{q_{i'j}^{(\nu)}}{q_{i'j}^{(\mu)}} = \prod_{(i,j) \in \mathcal{B}^*} \frac{q_{i'j}^{(\nu)}}{q_{i'j}^{(\mu)}} = \frac{\prod_{(i,j) \in \mathcal{B}^*}^+ \exp(-\beta\Delta\epsilon_{i'j}^{(\nu)})}{\prod_{(i,j) \in \mathcal{B}^*}^+ \exp(-\beta\Delta\epsilon_{i'j}^{(\mu)})}. \quad (2.33)$$

Combining Eqs. (2.33) and (2.32) gives the cluster move acceptance probability

$$W_{\text{acc}}^{(\mu \rightarrow \nu | R)} = \min \left\{ 1, \prod_{(i,j) \in \mathcal{B}^\dagger} \exp(-\beta\Delta\epsilon_{i'j}^{(\mu)}) \right\}, \quad (2.34a)$$

provided \mathcal{B}^* only contains outright failed links, and

$$W_{\text{acc}}^{(\mu \rightarrow \nu | R)} = 0, \quad (2.34b)$$

if \mathcal{B}^* contains a frustrated link. Note that, if the cluster is not restricted by N_C or other conditions, then $\mathcal{B}^* = \mathcal{B}$, $\mathcal{B}^\dagger = \emptyset$, and the acceptance probability reduces to Eqs. (2.27). Also, in the derivation of the acceptance probability, we assumed that \mathcal{C} was selected by free cluster selection, proposing links to all pairs satisfying (2.1a) and (2.1b). By using the results of Sec. 2.5, it is easily seen that, in practice, it is

enough to propose links to pairs satisfying (2.1a) and (2.1b), and interacting in one of the states μ or μ_i while still using Eqs. (2.34) to accept the cluster. The early rejection scheme is validated in Appendix C.2.

The conditions under which the algorithm of this section becomes efficient may arise for example in systems at high densities, where each move leads to the selection of a large cluster, and is likely to meet a condition restricting its size. If the system and conditions are such that the cluster selected in this way has a non-negligible chance of being accepted with Eq. (2.34), the algorithm may become orders of magnitudes faster than the version which would simply reject the proposed clusters. This is due to the fact that the computer time, spent on selecting the moving cluster is not wasted, and leads to a successful move of the cluster. A large gain in efficiency can be expected especially if the moving clusters are large.

2.7 Collective rotational Monte Carlo move

The move map for a rotation operation is specified by the position of the centre of rotation, the direction of the axis about which the rotation is performed, and the angle (magnitude) of the rotation. Although it is generally not problematic to choose these variables such that their probabilities are the same in the original and in the final state, collective rotation operations may involve several hidden issues. Firstly, rotating the cluster in periodic boundaries makes sense only if the cluster is smaller than half the minimum dimension L of the simulation box. A possible way of restricting the cluster size, to avoid this problem, is to extend the auxiliary conditions $\{b\}$ such that if (i, j) , $i \in \mathcal{C}$, $j \notin \mathcal{C}$ is a pair to which a link is to be proposed, with the distance of j from the root particle larger than $L/4$, then the link is forced to fail, i.e. included in \mathcal{B}^\dagger . Secondly, since the size of the particle displacement increases with its distance from the centre of rotation, any physically reasonable dynamics is approximated only if the displacement is restricted by a maximum distance specified again in $\{b\}$. Pairs (i, j) , $i \in \mathcal{C}$, $j \notin \mathcal{C}$, interacting in μ or in μ_i , which have a single-particle displacement j larger than some maximum (typically σ), are then again included in \mathcal{B}^\dagger . (Instead of using \mathcal{B}^\dagger , one can also reject the cluster without any chance of being accepted as was originally proposed by Whitlam and Geissler [2008].) Thirdly, we observed that collective rotational moves require double floating point precision, especially in systems with discontinuous potentials. Although single floating point precision is usually sufficient for MC simulation, random collective rotation operations with single floating point precision can systematically decrease or increase inter-particle distances, and bring pairs of

particles in \mathcal{C} from non-overlapping to overlapping positions. In fact, overlaps can occasionally occur even when using double floating point precision, and translational moves. For hard core systems or generally for systems with discontinuous potentials, we thus advise to check for the overlaps between particles of the moved cluster.

Let us now summarize the rotational move while applying the bias via forced failed links described in the previous section. The rotational move can be implemented as follows.

1. Pick a random particle, and use it as the first (root) particle of the cluster \mathcal{C} . Let \vec{r}_o denote the position vector (the centre of mass) of the root particle.
2. Pick the centre of rotation as a point with position vector \vec{r}_c defined as $\vec{r}_c = \vec{r}_o + a\vec{u}$, where \vec{u} is a random unit vector, $a \in \mathcal{U}(0, a_{\max})$ is a random number selected from a uniform distribution, and a_{\max} is a constant (typically $a_{\max} \approx \sigma$).
3. Pick another random unit vector \vec{u}_o defining the orientation of rotation, and a random angle with size $\theta \in \mathcal{U}(0, \theta_{\max})$, where θ_{\max} is a suitable constant. (Typically, $\theta_{\max} \approx \delta/\sigma$, with δ being the maximum size of the translational displacement.) Define the rotation matrix A corresponding to \vec{u}_o and θ . Let \vec{r}_i denote the position vector of particle i . The position vector of particle i after applying the Monte-Carlo move is expressed as $M\vec{r}_i = A(\vec{r}_i - \vec{r}_c) + \vec{r}_c$.
4. Perform the recursive loop selecting all other particles to \mathcal{C} .
 - (a) Pick randomly a pair (i, j) , $i \in \mathcal{C}$, $j \notin \mathcal{C}$, which interacts in state μ or in virtual state μ_i , and to which a link has not yet been proposed. If no such pair exists, finish the cluster selection by exiting the recursive loop.
 - (b) Test the maximum displacement size: If the size of $M\vec{r}_j - \vec{r}_j$ is larger than σ , label (i, j) as forced failed, and go to (a). Carry on, otherwise. (For small θ , this is equivalent to the test that $\vec{r}_j - \vec{r}_c$ is larger than σ/θ .)
 - (c) Test periodic boundaries: If the size of $\vec{r}_j - \vec{r}_c$ is larger than $L/4$, label (i, j) as forced failed, and go to (a). Carry on, otherwise.
 - (d) Attempt to create a link between (i, j) as follows.
 - i. Form a pre-link with probability
$$p_{i'j}^{(\mu)} = \max \left\{ 0, 1 - \exp(-\beta(\epsilon_{i'j}^{(\mu)} - \epsilon_{ij}^{(\mu)})) \right\} .$$
 - ii. If the pre-link does not form, label (i, j) as outright failed, go to (a).

iii. If the pre-link forms, calculate the reverse link formation probability

$$p_{ij'}^{(\mu)} = \max \left\{ 0, 1 - \exp(-\beta(\epsilon_{ij'}^{(\mu)} - \epsilon_{ij}^{(\mu)})) \right\} ,$$

where $\epsilon_{ij'}^{(\mu)}$ denotes the energy of (i, j) in μ after applying the move map M only to j , and form the link with probability

$$\min \left\{ 1, p_{ij'}^{(\mu)} / p_{i'j}^{(\mu)} \right\} .$$

(e) If the link does not form, label (i, j) as frustrated, go to (a).

(f) If the link forms, include j into \mathcal{C} , go to (a).

5. Identify the boundary \mathcal{B} of cluster \mathcal{C} as those pairs (i, j) , $i \in \mathcal{C}$, $j \notin \mathcal{C}$, to which a link was proposed but failed to form, i.e. is outright failed, forced failed, or frustrated.

6. Divide \mathcal{B} into forced failed pairs \mathcal{B}^\dagger , and all other pairs \mathcal{B}^* .

7. Accept the cluster move with probability

$$W_{\text{acc}}^{(\mu \rightarrow \nu | R)} = \min \left\{ 1, \prod_{(i,j) \in \mathcal{B}^\dagger} \exp \left(-\beta \Delta \epsilon_{ij}^{(\mu)} \right) \right\} , \quad (2.35)$$

provided \mathcal{B}^* only contains outright failed links. If \mathcal{B}^* contains a frustrated link, the move of \mathcal{C} is rejected.

2.8 Selection of the cluster under low acceptance probability

In this section we present a version of the symmetrized VMMC algorithm, which is suitable for situations where the cluster acceptance probability is expected to be low. Low acceptance can be caused naturally by, for example, high density and large moves, or artificially by accepting only clusters with some specific properties. This version of the algorithm has been described before by Whitlam et al. [2009a]; Whitlam [2011], but the following description is more detailed. Compared to Secs 2.4 and 2.5, the algorithm works with a cluster selection, which is further speeded up by proposing links only to pairs satisfying (2.1a), (2.1b), and (2.1c), where (2.1c) requires that (i, j) is interacting in state μ . This restriction again leads to an asymmetry of realizations of \mathcal{C} , because the boundary of the clusters in state

μ and ν cannot be required to be the same. In contrast to Sec. 2.5, the difference between the realizations cannot be ignored, and needs to be corrected with a bias in the acceptance probability (2.36). The bias is derived in Appendix B, and the corresponding early rejection scheme validated in Appendix C.3. Since $\mathcal{L}^{(\mu)} = \mathcal{L}^{(\nu)}$, $\mathcal{F}^{(\mu)} = \mathcal{F}^{(\nu)}$, but generally $\mathcal{B}^{(\mu)} \neq \mathcal{B}^{(\nu)}$, the realizations of clusters are said to have *symmetric core* and *asymmetric boundary*.

The algorithm can be summarized as follows:

1. Pick a random particle, and use it as the first (root) particle of the cluster \mathcal{C} .
2. Perform the recursive loop selecting all other particles to \mathcal{C} .
 - (a) Pick a pair (i, j) , $i \in \mathcal{C}$, $j \notin \mathcal{C}$, which interacts in state μ , and to which a link has not yet been proposed. If no such pair exists, finish the cluster selection by exiting the recursive loop.
 - (b) Attempt to create a link between (i, j) as follows.
 - i. Form a pre-link with probability

$$p_{ij}^{(\mu)} = \max \left\{ 0, 1 - \exp(-\beta(\epsilon_{ij}^{(\mu)} - \epsilon_{ij'}^{(\mu)})) \right\}.$$

- ii. If the pre-link does not form, label (i, j) as outright failed, go to (a).
- iii. If the pre-link forms, calculate the reverse link formation probability

$$p_{ij'}^{(\mu)} = \max \left\{ 0, 1 - \exp(-\beta(\epsilon_{ij'}^{(\mu)} - \epsilon_{ij}^{(\mu)})) \right\},$$

and form the link with probability $\min \left\{ 1, p_{ij'}^{(\mu)} / p_{ij}^{(\mu)} \right\}$.

- (c) If the link does not form, label (i, j) as frustrated, go to (a).
- (d) If the link forms, include j into \mathcal{C} , go to (a).
3. Identify the boundary $\mathcal{B}^{(\mu)}$ of cluster \mathcal{C} as those pairs (i, j) , $i \in \mathcal{C}$, $j \notin \mathcal{C}$ to which a link was proposed, but not formed, i.e. is either outright failed or frustrated.
4. If $\mathcal{B}^{(\mu)}$ contains a frustrated link, reject the move of \mathcal{C} .
5. If $\mathcal{B}^{(\mu)}$ only contains outright failed links, proceed as follows.
 - (a) Identify those pairs in $\mathcal{B}^{(\mu)}$ that are interacting in the original state μ but not interacting in the final state ν . Denote these pairs $\mathcal{B}^{(\mu\nu)}$.

- (b) Identify pairs (i, j) , $i \in \mathcal{C}$, $j \notin \mathcal{C}$, that are outside the interaction region in the original state μ , but end up in the interaction region in the final state ν . Denote these pairs $\mathcal{B}^{(\bar{\mu}\nu)}$.
- (c) Accept the cluster move with probability

$$\min \left\{ 1, \frac{\prod_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}}^+ \exp(-\beta \epsilon_{ij}^{(\nu)})}{\prod_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}}^+ \exp(-\beta \epsilon_{ij}^{(\mu)})} \right\}, \quad (2.36)$$

where the plus sign above the product in the numerator says that the product runs only over those pairs in $\mathcal{B}^{(\bar{\mu}\nu)}$ such that $\epsilon_{ij}^{(\nu)} > 0$. Similarly, the product in the denominator runs over pairs in $\mathcal{B}^{(\mu\bar{\nu})}$ such that $\epsilon_{ij}^{(\mu)} > 0$.

2.9 Recursive selection of the cluster

Cluster algorithms generally select the cluster under the implicit assumption, that pairs to which a link is proposed are taken randomly from all pairs satisfying conditions (2.1). Since we use the SDB condition to validate the algorithm, we discuss the rules of the cluster selection more carefully, and we distinguish between a random and an isotropic way of selecting the cluster. This will be used in the next section to formulate the SDB condition such that it takes the pair selection probabilities into account.

The general selection of the cluster (a random tree in graph theory) starts with a random choice of auxiliary variables, usually denoted by $\{b\}$, specifying the properties of \mathcal{C} , and the selection rules. We will consider a case where the definition of auxiliary variables $\{b\}$ consists of

1. random selection of the first (root) particle of \mathcal{C} ;
2. random selection of the maximum number of particles in the cluster, denoted by $N_{\mathcal{C}}$; and
3. random selection of the move map, M .

The cluster $\mathcal{C} \subset \mathcal{S}$ will again be selected in a recursive loop by linking particles of \mathcal{S} to \mathcal{C} . Let us define the queue as the set of pairs satisfying conditions (2.1). A link is proposed to a randomly chosen pair in the queue. If the link between (i, j) forms, j becomes a new member of \mathcal{C} , all pairs (k, l) such that $l = j$ are erased from

the queue, and all pairs (j, l) satisfying (2.1) that are not yet part of the queue are added to the queue. If the link does not form, the pair is erased from the queue.

Selecting pairs from the queue randomly like this is simple, intuitive, and common to most papers treating single cluster Monte Carlo algorithms [Wolff, 1989; Wu et al., 1992; Liu and Luijten, 2004; Troisi et al., 2005]. We call this rule *random ordering*; the resulting sequence of proposed links is called *random link order*; and the outcome is a *randomly ordered* cluster. An example of a randomly ordered cluster, and the corresponding random link order of \mathcal{C} , are illustrated in Fig. 2.3(a).

Random ordering may not be the best choice, depending on the objectives of the simulation. There may be physical reasons, which make it desirable to produce a cluster which is, for example, spatially distributed in a more isotropic fashion about the root, than a typical randomly ordered cluster. This could mimic a physical fluctuation, propagating isotropically in all directions from the root, out to some typical distance, which resulted in the coherent motion of a group of particles. Also, in the Brownian / Stokesian interpretation, it might be more justifiable to relate diffusion coefficients to cluster size, if the clusters are generated in this way. Alternatively, more isotropic or more compact clusters might be found (empirically) to generate more efficient simulations, in some circumstances.

Let us now illustrate how such a different rule for cluster selection might be formulated. We rely on the unique definition of *distance* between any two particles (vertices) in a tree, as the number of links (edges) joining them. Each particle in the growing cluster \mathcal{C} has a unique distance from the root, determined (once and for all) when it is added to the cluster. It is convenient to define a *generation* of particles in \mathcal{C} as the set of particles at the same distance from the root, measured in this way; the terminology is intentionally reminiscent of a family tree. A *living generation* is one such that there exists at least one pair (k, l) in the queue with k belonging to that generation; in other words, one which still has a chance to connect further particles l to the cluster. The *oldest* living generation is then the living generation with the shortest distance from the root (as defined above). The cluster selection then proceeds as above, but the links are proposed only to randomly selected pairs in the oldest living generation. A consequence of this selection rule is that there can be no more than two living generations at any stage. By analogy with the random case, we use the terminology *isotropic ordering*, *isotropic link order*, and *isotropically ordered* cluster (although this does not imply that the above rule is the only way of achieving similar results). An example is illustrated in Fig. 2.3(b). It might be possible to derive another kind of deterministic selection sequence by considering spiral ordering [Biyikoglu and Leydold, 2007].

It is easy to see that every isotropic link order is also a possible random link order, but not *vice versa*. This implies that the number of random link orders is higher than the number of isotropic link orders for a given \mathcal{C} , and hence that the probability of selecting an isotropic link order or isotropically ordered cluster is lower in random ordering than in isotropic ordering. If the cluster size is then restricted by an upper limit, say five particles in the cluster of Fig. 2.3, then the cluster has a higher chance of isotropic topology if isotropic ordering is used, than if random ordering is used. Isotropic ordering is thus an example of a cluster selection rule which affects the distribution of the moving cluster topologies.

2.10 Explicit formulation of the superdetailed balance condition

The SDB condition presented in Sec. 2.4 is implicit in that it does not take into account that the link formation probability $p_{ij}^{(\mu)}$ is conditional upon the probability of selecting a pair to which the link is proposed. In what follows, we will take the pair selection probabilities into account, and we will formulate the SDB condition explicitly.

We will again consider the free cluster selection, and we will denote the sequence of pairs to which a link was proposed during the recursive selection of $R_{\mathcal{C}}^{(\mu)}$ as

$$S_R^{(\mu)} = \left\{ \left(i, j | r^{(\mu)} \right) \right\}, \quad (2.37)$$

where $\left(i, j | r^{(\mu)} \right)$ denotes the $r^{(\mu)}$ -th pair to which a link was proposed. This sequence can be uniquely decomposed into subsequences

$$\left\{ \left(i, j | r_l^{(\mu)} \right) \right\}; \quad \left\{ \left(i, j | r_f^{(\mu)} \right) \right\}; \quad \left\{ \left(i, j | r_b^{(\mu)} \right) \right\}, \quad (2.38)$$

running over linked pairs $\mathcal{L}^{(\mu)}$, pairs in $\mathcal{F}^{(\nu)}$, and boundary pairs $\mathcal{B}^{(\mu)}$, respectively. Clearly, there is, in general, more than one recursive sequence $S_R^{(\mu)}$ leading to the realization $R_{\mathcal{C}}^{(\mu)}$. We shall denote the set of all possible sequences leading to $R_{\mathcal{C}}^{(\mu)}$ as $\left\{ S_R^{(\mu)} \right\}$. Also, there is generally, more than one realization $R_{\mathcal{C}}^{(\mu)}$ selecting particles to \mathcal{C} . The set of all possible realizations of \mathcal{C} in state μ will be denoted as $\left\{ R_{\mathcal{C}}^{(\mu)} \right\}$.

Let $L(r^{(\mu)})$ denote the number of pairs in the queue before the $r^{(\mu)}$ -th iterative step is applied. In random ordering, the probability of selecting a specific pair in the r -th iterative step is thus given by $1/L(r^{(\mu)})$.

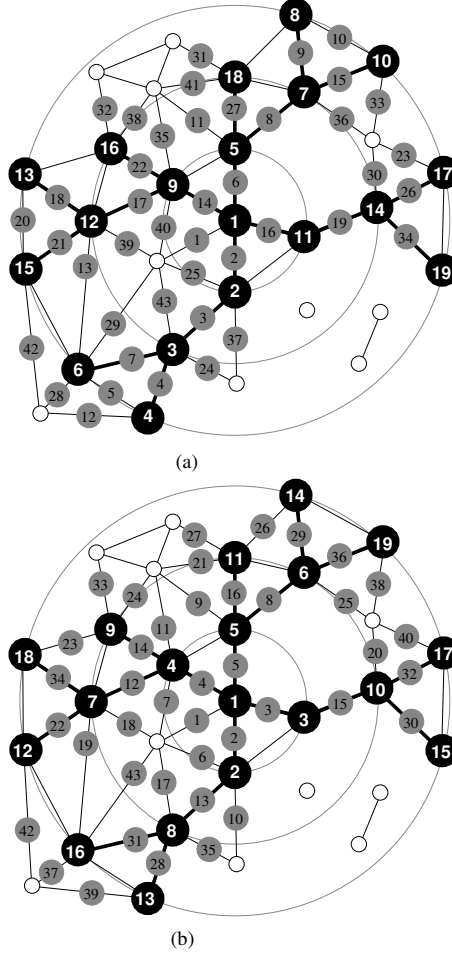


Figure 2.3: Diagrams showing two different orders in which the particles become members of the cluster, along with the order in which the links are proposed in a restricted cluster selection. Black numbered disks are particles in the cluster; small white disks are particles outside the cluster; thick black lines are formed links; thin black lines connect pairs interacting in μ (or in μ_i depending on the choice of the algorithm); numbered grey circles determine the link order. (In the next section, the link order is denoted as $S_R^{(\mu)}$ for the free cluster selection.) Large concentric circles identify particles within the same generation. (a) Random ordering. (b) Isotropic ordering.

For the free cluster selection, it is easily seen that

$$\left\{ S_R^{(\mu)} \right\} = \left\{ S_R^{(\nu)} \right\} \text{ and } \left\{ R_C^{(\mu)} \right\} = \left\{ R_C^{(\nu)} \right\}, \quad (2.39)$$

and also that for each $S_R^{(\mu)}$ and $S_R^{(\nu)}$ such that $S_R^{(\mu)} = S_R^{(\nu)}$, it holds that

$$\begin{aligned} L(r_l^{(\mu)}) &= L(r_l^{(\nu)}), \\ L(r_f^{(\mu)}) &= L(r_f^{(\nu)}), \\ L(r_b^{(\mu)}) &= L(r_b^{(\nu)}), \end{aligned} \quad (2.40)$$

which means that pair selection probabilities $1/L(r^{(\mu)})$ are the same in the old and in the new states at each step of the recursive cluster selection. We shall also denote

$$\begin{aligned} p(i, j | r^{(\mu)}) &= p_{ij}^{(\mu)}, \\ q(i, j | r^{(\mu)}) &= q_{ij}^{(\mu)}, \\ q(i, 'j | r^{(\mu)}) &= q_{i'j}^{(\mu)}, \end{aligned}$$

for a pair (i, j) , which is the r -th member of sequence $S_R^{(\mu)}$. The probability of selecting $S_R^{(\mu)}$ can then be expressed as

$$\begin{aligned} W_{\text{sel}}(S_R^{(\mu)}) &= \prod_{\left\{ (i, j | r_l^{(\mu)}) \right\}} \frac{1}{L(r_l^{(\mu)})} p(i, j | r_l^{(\mu)}) \\ &\quad \prod_{\left\{ (i, j | r_f^{(\mu)}) \right\}} \frac{1}{L(r_f^{(\mu)})} q(i, j | r_f^{(\mu)}) \\ &\quad \prod_{\left\{ (i, j | r_b^{(\mu)}) \right\}} \frac{1}{L(r_b^{(\mu)})} q(i, 'j | r_b^{(\mu)}), \end{aligned} \quad (2.41)$$

and the SDB condition imposed on $\left\{ S_R^{(\mu)} \right\}$ and $\left\{ S_R^{(\nu)} \right\}$ such that $\left\{ S_R^{(\nu)} \right\} = \left\{ S_R^{(\mu)} \right\}$, can be written as

$$\begin{aligned} \exp(-\beta E_\mu) p^{(\mu)}(\{b\}) W_{\text{sel}}^{(\mu)}(S_R^{(\mu)}) W_{\text{acc}}^{(\mu \rightarrow \nu | S)} \\ = \exp(-\beta E_\nu) p^{(\nu)}(\{b\}) W_{\text{sel}}^{(\nu)}(S_R^{(\nu)}) W_{\text{acc}}^{(\nu \rightarrow \mu | S)}, \end{aligned} \quad (2.42)$$

with terms having a similar meaning in Eq. (2.22). By using a reasoning analogous to

that in Sec. 2.4, and by using relations (2.39) and (2.40), the explicit SDB condition (2.42) can be simplified into the form (2.25), and further into the form (2.27). We note that auxiliary conditions in the SDB condition were originally defined such that $\{b\}$ should also include the probability of generating $S_R^{(\mu)}$. Since $1/L(r^{(\mu)})$ is conditional upon the set of already formed links, we do not adhere to this definition, and express $1/L(r^{(\mu)})$ separately from $p^{(\mu)}(\{b\})$.

One can also proceed slightly differently, and express the probability of selecting a realization as

$$W_{\text{sel}}(R_C^{(\mu)}) = \sum_{\{S_R^{(\mu)}\}} W_{\text{sel}}(S_R^{(\mu)}) . \quad (2.43)$$

The SDB condition imposed on $R_C^{(\mu)}$ and $R_C^{(\nu)}$ such that $R_C^{(\mu)} = R_C^{(\nu)}$, then reads as

$$\begin{aligned} \exp(-\beta E_\mu) p^{(\mu)}(\{b\}) W_{\text{sel}}^{(\mu)}(R_C^{(\mu)}) W_{\text{acc}}^{(\mu \rightarrow \nu | R)} \\ = \exp(-\beta E_\nu) p^{(\nu)}(\{b\}) W_{\text{sel}}^{(\nu)}(R_C^{(\nu)}) W_{\text{acc}}^{(\nu \rightarrow \mu | R)} , \end{aligned} \quad (2.44)$$

and can again be simplified into Eq. (2.25) or Eq. (2.27). Eq. (2.44) is, in fact, equivalent to the implicit formulation of the SDB condition in Eq. (2.22), or to the previous formulations of the detailed balance conditions [Niedermayer, 1988; Wu et al., 1992; Troisi et al., 2005] validating the cluster algorithms, where the realization of the cluster is understood as a “static pattern” of pairs rather than a “dynamic sequence” of pairs. The SDB condition applied to sequences of pairs offers an alternative insight into the validity of the algorithms described above, both for random and isotropic ordering. Note that there are alternative ways [Whitelam] to see that the selection of a specific pattern of pairs in the old and in the new states has the same probability.

2.11 Results

We firstly test the VMMC method on a system of $N = 2000$ particles interacting via a short-ranged attractive generalized Lennard-Jones [Vliegenthart et al., 1999] potential given by

$$V(r) = 4\epsilon \left[\left(\frac{\sigma}{r} \right)^{2\alpha} - \left(\frac{\sigma}{r} \right)^\alpha \right] - V_c, \quad (2.45)$$

with $\alpha = 18$, $\sigma = 1.0$, $\epsilon = 1.0$, in reduced units. The potential was truncated and shifted [Frenkel and Smit, 2002] at a cutoff distance $r_c = 1.8$, with the constant V_c chosen such that $V(r_c) = 0$. The simulations were carried out in the constant-

NVT ensemble. We took the critical packing fraction $\phi = \pi/6 \cdot N/V = 0.225$, and temperature $T = 0.6$ corresponding to the fluid phase [Vliegenthart et al., 1999]. We use the version of VMMC described in Sec. 2.5, and the isotropic ordering. A Monte Carlo cycle consists of N virtual translational or rotational moves. The rotational moves were performed in the same way as described in Sec. 2.7. The decision about performing a translation or a rotation is random, with on average 50% translations and 50% rotations. The size of the translational displacement was chosen randomly from the interval $(-\delta, \delta)$. The maximum sizes of translational and rotational displacements were taken as $\delta/\sigma = \theta_{\max} = 0.10$, $a_{\max} = \sigma$. The size of δ can significantly affect the efficiency of the algorithm, and our optimum choice will be discussed in Sec. 3.3. We compare the radial distribution functions $g(r)$ generated both with SPMC and VMMC only. For each case, the $g(r)$ is averaged over 10 independent simulations. The results are depicted in Fig. 2.4(a), and show that $g(r)$ are within the error range of each other. Similar tests were performed for fluid phases of similar short-ranged attractive systems at different temperatures, densities, and also in the constant- NpT ensemble. The VMMC method described in Sec. 2.8 was tested elsewhere [Whitelam].

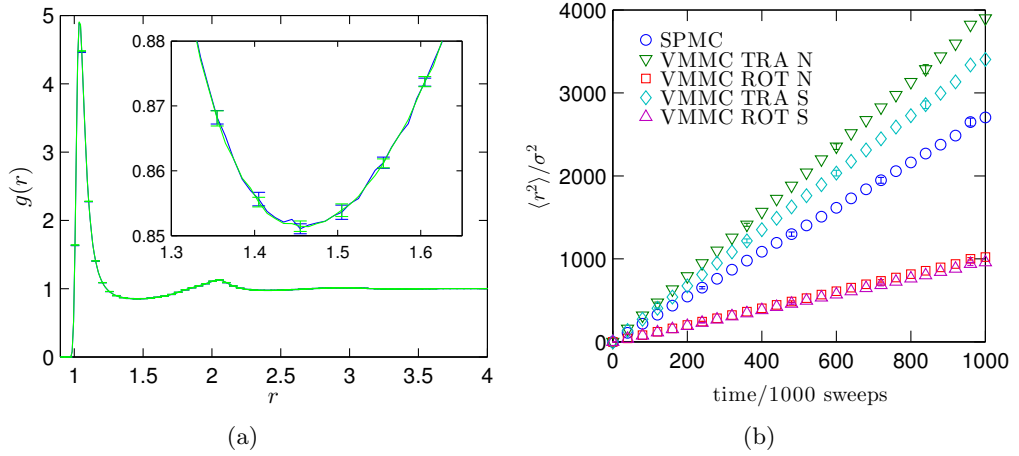


Figure 2.4: Simulations comparing single particle MC (blue) and collective VMMC moves (green). (a) Radial distribution functions in a fluid state of the system at $T = 0.6$ and $\phi = 0.225$. Inset: Magnification of the first minimum in $g(r)$. (b) Mean square displacement versus time measured in 1000 MC sweeps for the fluid at $T = 0.6$ and $\phi = 0.225$. VMMC either contains only rotations (ROT) or only translations (TRA), and simulations are either not scaled (N) or scaled (S) by $N_C = 1/x$. All error bars are estimated from 10 independent simulations.

Fig. 2.4(b) compares the mean square displacement (MSD) of SPMC with the MSD of VMMC. The SPMC has the same maximum displacement as the VMMC,

equal to 0.1σ . The VMMC is tested with 4 different setups. A MC sweep either includes 100% of rotational or 100% of translational moves, and the cluster selection is or is not scaled by condition (2.28). By scaled or not scaled dynamics we understand that the limit on the maximum cluster size in Eq. (2.28) is, or is not, applied. We use

$$N_C = 1/x, \quad (2.46)$$

with x being a random number from the uniform distribution $\mathcal{U}(0, 1)$. The scaling of the cluster selection can be used to approximate Brownian dynamics by collective MC moves [Whitelam et al., 2009a; Jack et al., 2007]. Although the time is non-physical, the gradient of the MSD is constant, defining a diffusion coefficient for each simulation in Fig. 2.4(b). Diffusion by the non-scaled and scaled translational VMMC is higher than in SPMC by 44% and 26%, respectively. This corresponds to the fact that more particles are moved in an average VMMC cycle. Indeed, the number of displaced particles in the respective VMMC simulations was higher, by 28% and 16%, than in SPMC. Diffusion by rotations is lower than by translations, despite the fact that the total number of displaced particles in the rotational VMMC is higher, by 22%, than in the translational VMMC. The rotational geometry of the displacements and the potentials is a plausible reason for this result.

Fig. 2.5(a) shows the evolution of the MSD in the Lennard-Jones system of Eq. (2.45), with a low packing fraction $\phi = 0.01$, quenched to the temperature $T = 0.25$, corresponding to the fluid-solid metastable region where clustering occurs [Perez et al., 2011]. The VMMC is again scaled by Eq. (2.46), and contains 50% of rotations and 50% of translations. The maximum displacements were chosen as $\delta/\sigma = \theta_{\max} = 0.30$. The initial configuration is a high temperature fluid conformation. The MSD is measured immediately from the start of the simulation, with no averaging over time origins. It can be seen that in the SPMC simulation of this system, particles get trapped by their neighbours in isolated aggregates, which do not move, meaning that the growth of the mean square displacement is extremely slow at later times. VMMC, on the other hand, moves both single particles and particles with their neighbours, thus producing a gel-like structure spanning most particles in the system, which itself moves as a single large cluster. The MSD in VMMC is thus several orders of magnitude larger than in SPMC, and the situation displayed in Fig. 2.4(b) is an example where VMMC is significantly more efficient than SPMC (at least on a certain timescale). The competition between local and macroscopic structure evolution in this system will be investigated in Sec. 3. Fig. 2.5(c) then shows that the efficiency of VMMC outperforms the efficiency of SPMC at later stages of the simulation when measured by the total MSD displacement. It must be

pointed out the gain in MSD is mainly caused by the motion of the whole aggregate, and that the local order evolves at about the same pace as in SPMC (see Sec. 3).

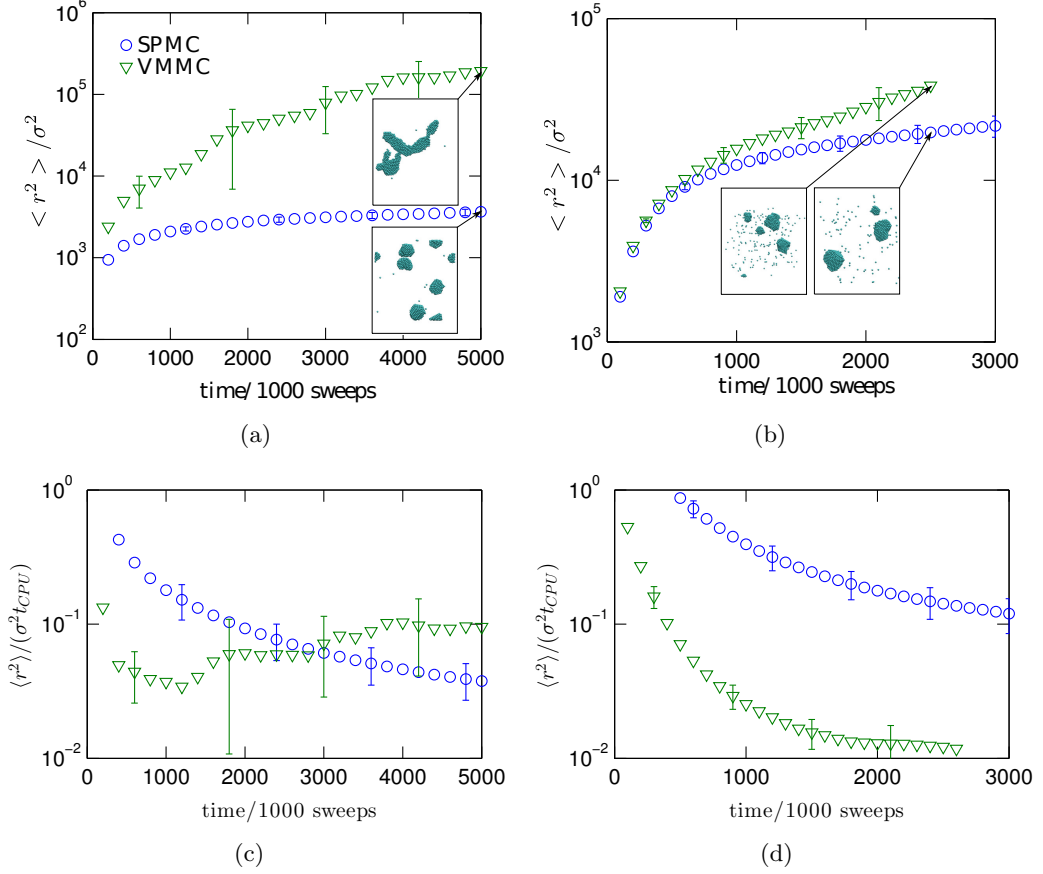


Figure 2.5: Simulations comparing single particle MC (blue circles) and collective VMMC moves (green triangles). (a) MSD for the fluid at $\phi = 0.01$ measured from the quench to $T = 0.25$. (b) MSD for short-range attractive and long-range repulsive fluid at $\phi = 0.01$ measured from the same quench. The insets are snapshots of typical configurations at times given by the arrows. (c), (d) Corresponding efficiencies, measured as MSD versus total CPU time in seconds. All error bars are estimated from ten independent simulations.

Fig. 2.5(b) shows the results from quenches with parameters identical to those presented in Fig. 2.5(a), except that $\delta/\sigma = \theta_{\max} = 0.20$ and the potential (2.45) has an additional long-range repulsive term of the Yukawa form $A \exp(r/\xi)/(r/\xi)$, where $A = 0.08$ and $\xi = 2.0$. The potential is cut off and shifted at $r_c = 3.0$. It will be seen in Sec. 3.5, that the long-range repulsion stabilizes the droplets [Sciortino et al., 2004], and that the system phase separates after the quench into 2–3 large and long-living aggregates. Fig. 2.5(b) shows that the MSD in VMMC evolves at about

the same pace as in the attractive system without a long-range repulsion, whereas the MSD in SPMC grows about an order of magnitude faster. This is likely to be caused by a single-particle exchange between the isolated aggregates. From the MSD point of view, the VMMC would thus not significantly outperform the SPMC. This can be better seen from Fig. 2.5(d), where the efficiency is measured with respect to the CPU time. The fact that the clusters can exchange single particles, together with the fact that the selection of moving clusters takes into account all particles with the large cutoff range $r_c + \sigma = 4\sigma$ (an extra σ is due to rotations), makes the VMMC much less efficient. (For SALR potentials, the VMMC can be speeded up by considering fictitious potentials to select the moving cluster [Whitelam, 2011]). The advantage of VMMC is, however, that it moves with the isolated structures, preserving their integrity, and can confirm their temporal stability against aggregation into a single domain. A detailed account of the interplay between the single-particle and collective motion will be addressed in Chap. 3.

2.12 Discussion

We have derived an alternative formulation of the symmetrized VMMC. Let us now discuss its efficiency. Proposing links to pairs interacting not only in μ but also in μ_i may seem to impact the speed, compared to the VMMC selecting clusters via realizations with symmetric core and asymmetric boundary, where links are proposed only to pairs interacting in state μ . However, one must take into account that determining the boundary of the cluster in the new state, in the latter version of VMMC, requires detection of the interacting neighbours of every particle in \mathcal{C} , in the new state. A similar reasoning applies to the static cluster algorithm [Troisi et al., 2005]. Without testing, we anticipate that the speeds of both versions of VMMC are comparable in the limit of high acceptance probability. If the cluster acceptance probability is low, selecting clusters via realizations with symmetric core and asymmetric boundary should be faster, because nearest neighbours in the new state do not need to be detected, provided the boundary contains a frustrated failed link. An advantage of the alternative version of the algorithm is that it may be simpler to implement, and allows us to simply accept clusters, which would have otherwise been rejected, through the scheme derived in Sec. 2.6. The efficiency generally depends on system properties and simulation parameters. For example, if the maximum size of displacement in M is smaller than the attractive range, and (i, j) is not interacting in μ but does so in μ_i , then also $p_{i'j}^{(\mu)} = 0$ and the proposed link is outright failed. This means that links can only be proposed to the

pairs interacting in μ , while still using the simplified acceptance probability (2.27). (Under the assumption of small displacements, proposing links to pairs interacting in μ_i does not affect the validity of the algorithm, but would slow it down.) Another example is a one-dimensional (1D) system with short-range attractions, where the boundary of the cluster is given by at most two pairs (corresponding to two border particles of \mathcal{C}). The set $\mathcal{B}^{(\bar{\mu}\nu)}$ thus only includes neighbours of these two particles; neighbours of other particles in \mathcal{C} can be omitted, and we can therefore expect the original version being faster than the reformulated version.

An alternative, and the original, versions of the symmetrized VMMC algorithm, were again proved via the use of the SDB condition, but with realization of the cluster defined as the set of all pairs to which a link is proposed during the cluster selection. It was shown that translational and rotational cluster moves, with clusters selected by the recursive stochastic linking proposed to nearest neighbour pairs, cannot generally possess a unique realization, because the set of pairs to which a link is proposed under the forward move is not the same as the set of pairs to which a link is proposed under the reverse move. This problem, which can be related to the unsolved task of enumerating the number of independent graphs in Monte Carlo cluster algorithms [Krauth, 2006], was described here in terms of the asymmetric core or boundary of the cluster (Sec. 2.1), and can be briefly summarized as follows: If links are proposed only to pairs interacting in the original state, the set of proposed links internal to the cluster is the same in the old and in the new state, but the sets of pairs to which a link is proposed and which end up in the boundary are generally different. The realization has a symmetric core but an asymmetric boundary in this case. If links are proposed to pairs interacting in the original or in the virtual state, the set of pairs in the boundary of the cluster is the same, however, the sets of pairs internal to the cluster are generally different. The realization then has a symmetric boundary but an asymmetric core. These iterative selection rules thus do not generally yield realizations which would have symmetric core and boundary at the same time. The SDB condition, imposed on the realization of the cluster, was thus formulated differently in terms of the free cluster selection, where links were proposed to all pairs satisfying conditions (2.1a) and (2.1b). Realizations are then guaranteed to have both symmetric boundary and symmetric core, and the SDB condition can be formulated exactly. It is then shown that under certain assumptions, not proposing pairs to all (2.1a) and (2.1b), but only to certain pairs, does not affect the validity of the algorithm.

Let us relate the notion of symmetric boundary and symmetric core to other cluster algorithms [Wolff, 1989; Wu et al., 1992; Troisi et al., 2005]. In the Wolff

cluster algorithm the realization of the cluster has both symmetric core and symmetric boundary. The SDB condition can thus be imposed on the realization of \mathcal{C} without any ambiguity. The symmetry of the core and boundary in the Wolff cluster algorithm follows from the fact that spins are flipped and stay at one lattice position, rather than being moved from one lattice position to another. The symmetry of the realization is also preserved in the translational cluster algorithm for charged lattice spin systems [Wu et al., 1992]. This is because the spins are linked with probability one, if they are close enough, and fail to form otherwise. A consequence of this deterministic linking is that there are no such spins (particles) j outside the cluster to which a link would be proposed in one state and not in another state, and the requirement about equal chances of selecting the pairs at each step of the recursive selection is naturally satisfied by the definition of this linking probability. We note that the same authors proposed a possible extension of their theory to stochastic linking, without further specifying the exact rules of the cluster selection. If the selection of the cluster was recursive following random or isotropic ordering, the stochastic linking would lead to realizations of clusters having asymmetric boundaries, and considerations similar to those presented in this chapter would be needed to validate the algorithm. In the algorithm of Troisi and Bhattacharyay [Troisi et al., 2005; Bhattacharyay and Troisi, 2008] the links are not proposed in a recursive way, but to all pairs of the system. For each state μ this leads to the definition of a square matrix with elements given by $p_{ij}^{(\mu)}$, $i, j = 1, \dots, N$. The cluster is then detected as those particles which are at least singly joined by links. This selection leads to a cluster acceptance probability which is, indeed, biased exactly by the energy change on the boundary of the cluster. In practice, the authors do not consider the entire matrix $p_{ij}^{(\mu)}$, but only a part of it by a recursive process described above in this chapter. In fact, restriction of the whole matrix to only a part of it, used to validate their algorithm, is very similar to the restriction of pair selection by condition (2.1c) used to derive the acceptance probability in the VMMC.

We will now point out a weak point in our tests performed to numerically validate the VMMC. Fig. 2.4(a) shows that the radial distribution function of a generalized Lennard-Jones fluid is within the error range of the radial distribution function obtained by a SPMC. We have thus verified that the VMMC samples the fluid phase within the accuracy of the SPMC. Although we performed these tests for square-well and 36-18 LJ fluids at different conditions, we emphasize that the probability of generating or accepting a cluster decreases exponentially with cluster size, and that the dominant moving cluster size is one or two particles. The contri-

bution of larger clusters to the total number of accepted clusters is thus negligible, and large cluster moves are likely to have only a very little effect on generating an equilibrium MC chain. It would thus be appropriate to perform tests on an equilibrium (ergodic) system where the dominant moving cluster size is significantly higher than one particle. The simplest example of such systems might be particles with short-range attractions and long-range repulsions at a suitable $\epsilon/k_B T$. Those systems were proved to form stable cluster phases [Sciortino et al., 2004; Charbonneau and Reichman, 2007b], and we have observed that a suitably chosen displacement can indeed lead to the dominant moving cluster size larger than 1 (and smaller than N). We have not compared the radial distribution function with a radial distribution function of a SPMC, because we know that the latter equilibrates those systems too slowly. Another non-local algorithm such as aggregation volume bias MC [Chen and Siepmann, 2000] would be needed in order to sample the configurational space of an equilibrium cluster phase. At sufficiently low densities such that structural arrest does not take place [Charbonneau and Reichman, 2007b], this might then be used in a very long simulation, to verify that the cluster algorithm samples from the Boltzmann distribution even when the dominant clusters are large. Another possibility for a reference simulation is to compare with a standard single particle Monte Carlo simulation, combined with simple static cluster moves [Troisi et al., 2005]. These tests were not performed here. In addition to the 3D tests, we did 1D continuum model tests similar to the 1D lattice model tests of Troisi et al. [2005]. Our continuum model tests matched the analytical result of Walmsley [1988] exactly. We emphasize that one-dimensional tests are necessary but not sufficient to validate the SDB condition, because there is only one realization or one sequence of proposed links selecting the one dimensional moving cluster. We thus conclude that although our tests confirm the validity of VMMC in the limit of small clusters, they do not provide strong numerical evidence that the algorithm is valid for large moving clusters, and in that limit we can only rely on analytical considerations deriving the VMMC. We stress that other unpublished tests [Whitelam] have been performed by Romano et al. [2012] comparing VMMC with Molecular Dynamics simulations in systems where collective motion is important.

We will now address a more general aspect of the algorithm. The form of the linking probability in Eq. (2.16) implies that no links and hence no clusters are formed as a result of a repulsive interaction ($\epsilon_{ij}^{(\mu)} = 0$, $\epsilon_{i'j}^{(\mu)} > 0$). In Chap. 5, we will show a general way of linking the particles and accepting the relevant collective translational and rotational MC moves, which is applicable to any form of pairwise interaction, including hard-core repulsion. Although this can also be done by

fictitious potentials [Whitelam, 2011], the general linking is different in that it proposes links to pairs interacting in the virtual state corresponding to $\epsilon_{ij'}^{(\mu)}$, and defines general properties of the linking function. We also mention that new event-based rejection-free MC algorithms have recently been formulated for general potentials. Firstly, event-chain Monte Carlo [Bernard and Krauth, 2012; Bernard et al., 2009] displaces chains of particles in an efficient way. Secondly, event-driven Monte Carlo [Peters et al., 2012] displaces single particles, but selects and accepts them in a way which is dynamic. It would be interesting to see how these algorithms complement the VMMC methods in capturing the kinetic and thermodynamic crossover in glassy systems [Jack et al., 2007; Klotz and Jack, 2011].

This chapter aimed to clarify the way of creating collective translational and rotational Monte Carlo moves, based on local pairwise energy changes, and to shed more light on the technical details, as well as to provide a clear validation of the algorithm. Apart from the tests of the MSD in Fig. 2.4, Stokesian or Brownian scaling was not considered here, because it requires one to approximate a MC sweep composed of collective rotations and translations with a timestep. Presentation and tests of these approximations are outside the scope of this chapter. Nevertheless, it is worth mentioning that one of the main reasons for the asymmetric boundary in the original version of the symmetrized VMMC was to speed up the generation of proposed clusters, because a large fraction of them are rejected immediately without any chance of being accepted. This is done for clusters exceeding a certain size in order to approximate the real dynamics. Sec. 2.6 provides a scheme controlling the cluster size, the advantage of which is that its acceptance rate is potentially higher than in the case of purely rejecting the clusters. Approximation of the dynamics by this scheme is an example situation where the alternative formulation of the VMMC algorithm (Secs 2.5 and 2.6), might naturally outperform the original scheme (Sec. 2.8), because more generated clusters can be accepted. Higher acceptance rate is, indeed, desirable in MC schemes approximating Brownian or other dynamics [Sanz and Marenduzzo, 2010]. Moreover, the size distribution of proposed clusters is strongly correlated with the maximum size of the Monte Carlo map M . We can speculate that the distribution of accepted clusters might be additionally optimized, for example, by drawing the size of the Monte Carlo map M not from a uniform but from another (possibly Gaussian) distribution such that the distribution of proposed and accepted cluster sizes ends up being closer to the distribution in a real dynamics.

2.13 Conclusion

We have formally described and analysed the symmetric version of the VMMC algorithm, and made additional numerical tests. The early-rejection scheme is explicitly used to show that the algorithm samples from the Boltzmann distribution, while treating the superdetailed balance condition in another way. To clarify the theory behind the algorithm, we have presented an alternative formulation of the VMMC algorithm which has a simple acceptance form. The main advantage of this reformulated scheme is that it provides an easy way of accepting clusters which would otherwise be rejected. This may find its use in controlling the cluster size distribution, but also in the implementation of kinetically realistic rotational moves. We anticipate that the original formulation of the symmetrized VMMC is still the fastest, but system- and simulation-dependent exceptions might exist.

Similarly to what was originally reported by Whitlam et al. [2009a], this thesis shows that the VMMC algorithm is particularly useful in predicting the gelation boundary of low- and intermediate-density short-ranged attractive systems below the critical point [Perez et al., 2011], where the density fluctuations [ten Wolde and Frenkel, 1997], phase separation [Foffi et al., 2005b], or spinodal decomposition [Lu et al., 2008], lead to the formation of meta-stable droplets, fractal aggregates, or even stable clusters [Sciortino et al., 2004; Charbonneau and Reichman, 2007b]. The cluster moves enhance the mobility of those structures which can kinetically slow down (arrest) and form glasses or gels. The results of this thesis or the results of Whitlam et al. [2009a] which report the effect of collective motion on gelation are restricted to a narrow range of system parameters $(\rho, \epsilon/k_B T, \alpha)$, determining the onset of gelation. Other simulation parameters including the optimum displacement size, the simulation ensemble, particle shape, or the quench or compression (crunch) rate [Royall and Malins, 2012] may also affect the structure of the gel. One can thus imagine a whole class of possible studies examining the affect of collective motion on gel formation, or crystallization, in attractive colloidal or molecular systems, making use of cluster algorithms of this kind.

Chapter 3

Kinetically slowed down phase separation

In this chapter, we provide an insight into the behaviour of the Virtual Move Monte Carlo cluster algorithm. We show that the maximum size of Monte Carlo displacement can significantly affect the efficiency of the algorithm, and we investigate how different choices of the displacement affect the dynamical trajectories leading to arrested or kinetically slowed down phases. Firstly, we show that a short-ranged attractive equilibrium fluid has a universal exponential scaling of moving cluster size distribution in a region above the inflection point of the fluid-solid binodal. We then quench the same system into the temperature region where the extrapolated glass line meets the spinodal (Fig. 1.1). We reveal that the system exhibits universal linear growth regimes under single-particle MC correlating the growth rate to the local structure. These regimes are suppressed under the collective motion where droplets aggregate into a single large disordered domain. Finally, it is shown that addition of long-range repulsion to the short-range attraction leads to the formation of long-living spherical aggregates after those quenches. The long-range repulsion seems to recover the linear growth regime even under the collective motion, and the dynamics scales similarly as in the single-particle Monte Carlo simulations of attractive systems without a long-range repulsion.

3.1 Introduction

In the previous chapter, we have presented the symmetrized VMMC algorithm which makes collective translational and rotational moves, and which can efficiently explore long-time evolution of metastable colloidal droplets created by quenches of systems

with short-range attractions. This chapter presents systems, conditions, and optimum parameters, under which VMMC is efficient. In particular, it focuses on the optimum maximum displacement size δ , a parameter governing the moving cluster size distribution. Three systems are examined. First, an equilibrium short-range attractive fluid at moderate densities shows how δ affects the equilibrium distribution of proposed and accepted clusters of a given size. Second, a low-density and short-range attractive fluid quenched to the metastable vapour-liquid transition region is used to show how metastable droplets developed after the quench move and crystallize both in VMMC and in SPMC simulation. Finally, the same system of short-range attractive particles, but equipped with long-range repulsions stabilizing the droplets, is quenched under the same conditions. It is shown that the evolution of the structure is similar to the structure evolution of an attractive system without long-range repulsion simulated with the SPMC.

We gather an ensemble of trajectories from simulations following the quenches to measure averages of observables quantifying local and macroscopic properties of the system. We show that a consequence of the stochastic collective motion based on local energy gradients is that the local order measured on the ensembles of trajectories reveals a scaling with respect to the time and the maximum δ , and that this scaling is similar in SPMC and VMMC simulations. The trajectory ensembles will also show that some observables are not, and other observables may be, invariant with respect to δ . For example, the maximum size of the second largest partition varies with δ , but the system seems to strive for a similar long-time non-equilibrium structure, characterized by a certain number of isolated crystalline domains independently of δ . Some of these results confirm previous observations that the kinetically slowed down evolution of the system is invariant with respect to the dynamics Foffi et al. [2005a].

3.2 Methods

3.2.1 Simulation model

We use a system of $N = 2000$ particles interacting via the short-range attractive generalized Lennard-Jones (LJ) [Vliegenthart et al., 1999] potential given by Eq. (2.45). To simulate the cluster phases we use the short-range attractive and long-range repulsive (SALR) potential

$$V(r) = A \frac{e^{-r/\xi}}{r/\xi} + 4\epsilon \left[\left(\frac{\sigma}{r} \right)^{2\alpha} - \left(\frac{\sigma}{r} \right)^\alpha \right] - V_c, \quad (3.1)$$

where A is the Yukawa amplitude and ξ is the screening length. We took $A = 0.08$, $\xi = 2.0$, $\alpha = 18$, $\sigma = 1.0$, $\epsilon = 1.0$. The potential is truncated and shifted at a cut-off distance $r_c = 3.0$, with V_c having meaning analogous to that of Eq. (2.45). The potentials (2.45) and (3.1) are displayed in Fig. 3.1.




Figure 3.1: Short-range attractive 36-16 Lennard-Jones potential (blue). Short-range attractive and long-range repulsive potential with $A = 0.08$ and $\xi = 2.0$ (green). Zero potential $V(r) = 0$ (black) is displayed as a reference potential.

The simulations were carried out in the constant- NVT ensemble. We used two different packing fractions. First, $\phi = 6/\pi \cdot N/V = 0.225$ corresponds to the critical packing fraction [Vliegenthart et al., 1999], and is used to examine the equilibrium Lennard-Jones liquid. Second, relatively low packing fraction $\phi = 6/\pi \cdot N/V = 0.01$ is used to simulate the kinetic slowing in attractive systems [Perez et al., 2011; Khan et al., 2009] and to examine the formation of stabilized partitions in systems with competing interactions [Sciortino et al., 2004]. The fluid phase is simulated at a high temperature $T = 2.0$, and a low temperature $T = 0.6$, which is slightly above the critical temperature $T_c = 0.43$ of 18-36 LJ system [Vliegenthart et al., 1999]. To simulate the kinetic slowing down and cluster phases we used $T = 0.25$ across all the simulations. The temperature $T = 0.25$ and the packing fraction $\phi = 0.01$ then seem to be close to the region of the kinetic phase diagram of Zaccarelli [2007], where the spinodal intersects the extrapolated glass line (Fig. 1.1), and where phase separation starts to be logarithmically slowed down to an extent that long-lived disordered aggregates can be observed [Perez et al., 2011]. Moreover, this temperature lies on the boundary of the optimum region for self-assembly of a similar short-range attractive colloidal system, gleaned from analysing the fluctuation-dissipation ratios [Klotsa and Jack, 2011; Grant and Jack, 2012]. We thus expect both the large scale aggregation of drops and the evolution of local order happen simultaneously on the timescales we shall consider. To study directly how repulsion affects the kinetics of

phase separation, the same temperature is used for the short-range attractive system with long-range repulsion. The inverse temperature is denoted as $\beta = 1/k_B T$, where k_B is the Boltzmann constant.

The simulations of kinetic arrests (or more precisely kinetically slowed down phase separations [Testard et al., 2011; Perez et al., 2011]) start from a randomly distributed set of particles without high energetic overlaps. These conformations can be thought to be random high temperature fluid states. *A priori* simulations equilibrating the system into a typical fluid state [Foffi et al., 2005a] at $T = 1.0$ were not performed here, because we only start at low packing fractions, where initial large inhomogeneities are unlikely. Since the simulation starts from a fluid state, but under conditions corresponding to an equilibrium state in which the crystal coexists with the gas, and since no time is given to the system to equilibrate at intermediate conditions between these two states, the simulations is referred to as the *rapid* (non-equilibrium or instantaneous) *quench*, as opposed to slow (equilibrium or finite) quenching [Cates et al., 2004; Toledano et al., 2009; Royall and Malins, 2012], where the system is allowed at least a short time to find a local equilibrium and locally phase-separate.

3.2.2 Analysis

To analyse the speed of aggregation and stability of cluster phases, we monitor the size of phase separated aggregates. We say that particles i and j are neighbours, if they are separated by a distance less than 1.4σ . This distance roughly corresponds to the local maximum in the interaction potential in Eq. (3.1), but is more universal, and can characterize clustering in systems of attractive particles with repulsive core, or even just hard-sphere spheres [Schilling et al., 2010]. A *partition* (also *domain* or *aggregate*) is a group of particles which can be connected through a set of neighbouring particles, i.e. particles k and l are members of the same partition if and only if there is a set of neighbours such that $(k, a), (a, b), \dots, (y, x), (x, l)$. The number of particles within a partition is defined as the size of a partition. We will monitor the time evolution of the ten largest partitions. If necessary, we restrict our attention only to the three largest partitions, or to the first, second, and tenth largest partition of the system.

The term cluster is normally used for what we term here the partition, but it must be stressed that this section analyses the VMMC algorithm, and that we reserve the words *moving cluster* or simply *cluster* for the set of particles that are moved within the Monte Carlo cluster move, and that a cluster and a partition are generally not the same sets. The moving cluster is generated randomly, and although

it often spans the entire partition (is equal to the partition of the root particle) it can just be only a subset of it, or it can extend across different partitions. Two different words, partition and cluster, thus need to be carefully distinguished in this section. With this definition it would make sense to refer to the cluster phase [Sciortino et al., 2004] as the partition phase. We believe that the term “partition phase” would not be appropriate, and we adhere to the original “cluster phase”.

The time evolution of local crystallinity will be analysed with q_6q_6 parameters [Steinhardt et al., 1983; ten Wolde et al., 1996; Schilling et al., 2010]. If $N_b(i)$ is the number of neighbours of particle i , the orientational structure can be characterized by

$$q_{lm}(i) = \frac{1}{N_b(i)} \sum_{j=1}^{N_b(i)} Y_{lm}(\vec{r}_{ij}), \quad (3.2)$$

where $Y_{lm}(\vec{r}_{ij})$ is the spherical harmonic corresponding to the orientation vector \vec{r}_{ij} pointing from particle i to particle j . It is known that the crystal structure of the 36-16 LJ system developed after a suitable quench is a mixture of hexagonal and face-centred cubic lattices, with a slight predominance of the hexagonal structure [Perez et al., 2011]. We have thus focused on $l = 6$. A pair (i, j) is said to be orientationally bonded if

$$\frac{\sum_{m=-6}^6 q_{6m}(i) \bar{q}_{6m}(j)}{(\sum_{m=-6}^6 |q_{6m}(i)|^2)^{1/2} (\sum_{m=-6}^6 |\bar{q}_{6m}(j)|^2)^{1/2}} > 0.7, \quad (3.3)$$

where $\bar{q}_{6m}(j)$ is the complex conjugate of $q_{6m}(j)$, and $|q_{6m}(i)|$ is the size of the thirteen dimensional vector $q_{6m}(i)$. The number of neighbours which are orientationally bonded to i is denoted as $n_b(i)$. A particle i is called *crystalline* if $n_b(i) = 12$, and *weakly bonded* if $n_b(i) = 2$. Sometimes, the context will suggest that a crystalline particle refers to a particle with $n_b(i) > 10$, and weakly bonded particle to a particle with $1 < n_b(i) \leq 5$. A group of at least two bonded crystalline particles, will be called a *crystal* or a *crystalline domain*. A group of at least two particles with $n_b(i) > 5$, will be called a Low Symmetry Cluster (LSC). We will be distinguishing between the LSC with and without a crystal inside. Note that it would also make sense to use the term “low symmetry partition”, instead of “low symmetry cluster”. To be consistent with Schilling et al. [2010] we use the term “low symmetry cluster”, even if the low symmetry cluster does not describe what we call here the cluster, but refers to a structure of what call here the partition, as above.

To characterize the products of a kinetically slowed down phase separation we will define the *disordered aggregate* as a state of the system such that it contains

a partition with at least two particles in that partition being separated by a distance of at least half of the box size. At our low density, this generally means that such an aggregate must be highly aspherical and concave. We typically assume that the aggregate does not disintegrate over a sufficiently long trajectory, and that it is a product of a kinetically slowed down phase separation [Perez et al., 2011; Testard et al., 2011, 2013]. It will be seen that these assumptions are satisfied from monitoring the structure of partitions and from the behaviour of VMMC. The definition of a disordered aggregate is similar to the definition of a gel in simulations with periodic boundary conditions [Toledano et al., 2009; Sciortino et al., 2005]. A gel is defined as a structure that, if repeated, percolates through the entire system in at least 50% of randomly chosen conformations taken over a sufficiently long trajectory. A gel defined in this way is a disordered aggregate, but not *vice versa*, a disordered aggregate is not necessarily a gel (see for instance Fig. 3.6). In fact, these definitions are technical and suitable for our small system. In a very large system the disordered aggregate may extend to a large portion of space, and even if it does not percolate through the box, the aggregate itself may have a low density and a fractal-like structure, and thus properties of a gel.

To analyse the fraction of moves which are accepted in the VMMC algorithm we need to distinguish between the probabilities of generating and accepting differently sized moving clusters. We then study how the distribution of cluster sizes evolves in time depending on the initial choice of δ . Let $N_g(n_C, \delta)$ be the average number of all clusters of size n_C (accepted or rejected) generated with a displacement taken from the range $(-\delta, \delta)$. The average is taken over several (typically ten) independent simulations, and over a time specified later. We define the probability density of generating a cluster of size n_C with a Monte Carlo move map of maximum size δ as

$$f_g(n_C, \delta) = \frac{N_g(n_C, \delta)}{\sum_{n_C=1}^N N_g(n_C, \delta)}. \quad (3.4)$$

We will call it the *generation density* for short. Similarly, we can define the probability density of accepting a proposed cluster of size n_C with M as

$$f_a(n_C, \delta) = \frac{N_a(n_C, \delta)}{\sum_{n_C=1}^N N_g(n_C, \delta)}, \quad (3.5)$$

where $N_a(n_C, \delta)$ has the same meaning as $N_g(n_C, \delta)$, except that it only counts accepted clusters. We will call it the *acceptance density* for short. The *acceptance probability* of a cluster of size n_C generated with a displacement of size δ will be

defined as

$$P_a(n_{\mathcal{C}}, \delta) = \frac{N_a(n_{\mathcal{C}}, \delta)}{N_g(n_{\mathcal{C}}, \delta)} = \frac{f_a(n_{\mathcal{C}}, \delta)}{f_g(n_{\mathcal{C}}, \delta)}. \quad (3.6)$$

3.3 Equilibrium fluids

Before we show how different choices of maximum displacement δ affect the non-equilibrium VMMC simulation of the kinetically slowed down structures, it is worth illustrating the optimum choice of δ in equilibrium fluids. We use the generalized LJ attractive system defined by Eq. (2.45) with $\phi = 0.225$, $T = 0.6$, and different values of δ . To study the dependence on δ directly, the move map M is a translational displacement which is *not* uniformly chosen from $(-\delta, \delta)$, but is *fixed* at δ in a specified (randomly-chosen) direction. Figure 3.2(a) shows that the probability density of proposed cluster sizes has an exponential dependence on $N_{\mathcal{C}}$, when the displacement is sufficiently small, $\delta \leq 0.10\sigma$. Figure 3.2(b) then shows that only clusters created by displacements $\delta \leq 0.10\sigma$ have a reasonable chance of being accepted. If $\delta \ll 0.10\sigma$, the simulation does not involve movements of large clusters. If $\delta \gg 0.10\sigma$, the simulation becomes inefficient, as most of the computer time is spent on generating large clusters, which are doomed to be rejected. The maximum size of the displacement should thus be $\delta \approx 0.10\sigma$ in this case.

We note that the cluster move acceptance probability in Fig. 3.2(b) has an exponential dependence $\propto \exp(-N_{\mathcal{C}}/N^*)$, with the same characteristic size $N^* \approx 3.0$, for all values of the maximum displacement studied. To verify whether this is a universal behaviour in equilibrium fluids, we also did a similar analysis for a higher temperature $T = 2.0$ ($\phi = 0.225$) with results shown in Figs. 3.2(c) and 3.2(d). It can be seen that the cluster move acceptance probability no longer has an exponential form. (The optimum displacement size increases to at least $\delta = 0.20\sigma$.) Interestingly, an exponential scaling similar to that in Fig. 3.2(b) is observed for square-wells and other short-range Lennard-Jones attractions $\alpha = 12$, $\alpha = 50$, (but not for wide range of attraction given by $\alpha = 6$) for fluids at packing fractions and temperatures in the vicinity of the inflection point of the fluid-solid binodal. It would be interesting to study the relation of this behaviour to the scaling of the dynamics [Foffi et al., 2005b] or to the extended law of corresponding states [Noro and Frenkel, 2000] observed in similar systems.

As for the rotations, the distributions of proposed clusters differ from the translations. We fixed the maximum distance of the centre of rotation from the root as $a_{\max} = \sigma$. We observed that large clusters are most likely to be generated with $\delta = 0.10\sigma$, i.e. the probability of generating a cluster of size $N_{\mathcal{C}}$ is smaller for any

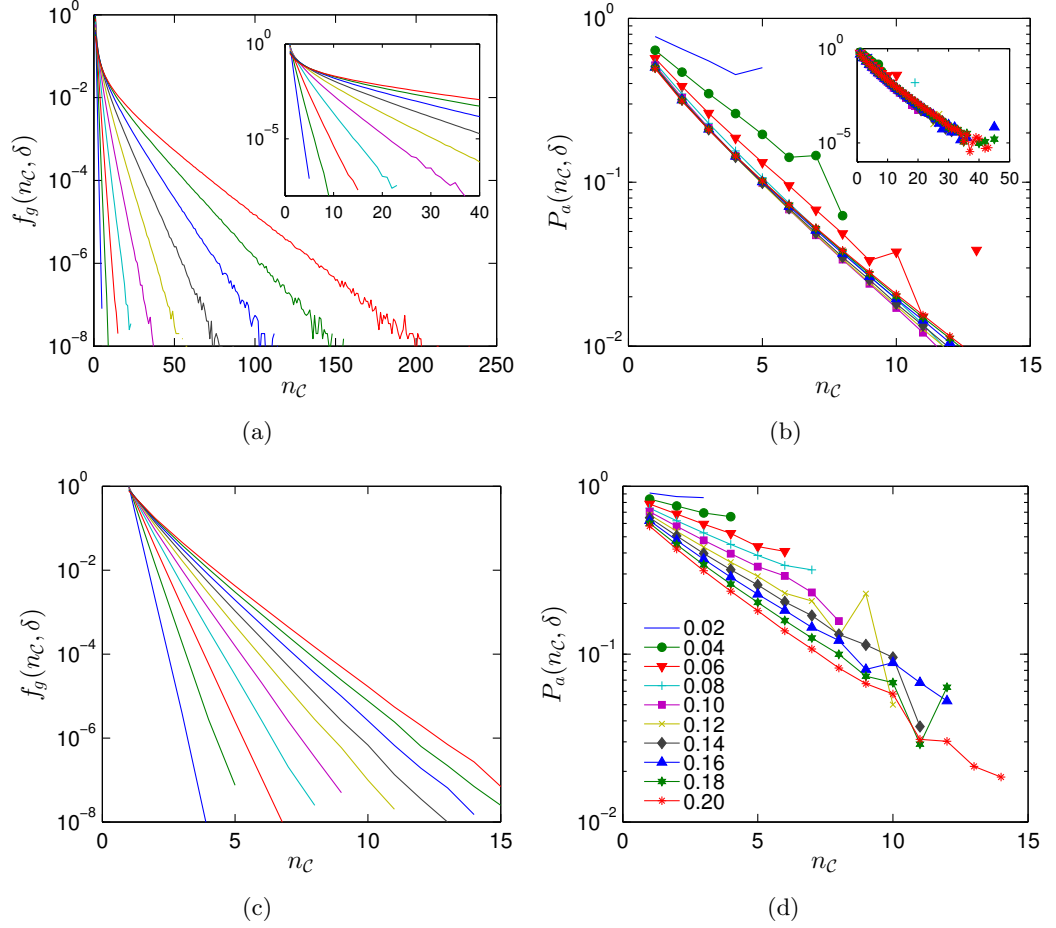


Figure 3.2: Probability density (log scale) of generating and probability of accepting a moving cluster of size n_c in a 18-36 Lennard-Jones fluid for different maximum displacements δ . Different colours (symbols) represent different sizes of translational displacement δ , see also legend in (d). For clarity, graphs of densities $f_g(n_c, \delta)$ do not contain markers, and correspond to $\delta = 0.02\sigma, 0.04\sigma, \dots, 0.20\sigma$ from left to right (colors are preserved). All runs were done at the critical packing fraction $\phi_c = 0.225$. Simulations in (a) and in (b) are at temperature $T = 0.60$, not far above the critical temperature $T_c = 0.43$. Simulations of (c) and in (d) are at $T = 2.0$ corresponding to a fluid phase. Inset of (a): Magnification in the region of small clusters and high probability densities. Inset of (b): Acceptance probabilities of large clusters.

δ which is larger or smaller than 0.10σ . The probability of accepting a rotational cluster of size N_c decreases with increasing δ_{\max} . We thus take $\delta_{\max} = 0.10\sigma$. An exponential scaling of acceptance probability for rotations was not observed.

3.4 Undercooled attractive systems

In this section, we examine the kinetic evolution of a low density short-range attractive fluid, after a quench to the temperature $T = 0.25$. At the same time, we study the properties of our Monte Carlo methods. In particular, we focus on the maximum Monte Carlo displacement size δ .

3.4.1 Coarsening in single-particle MC

Let us firstly investigate the coarsening which happens under the SPMC motion only. The dynamics of this model is largely unphysical, but can represent a realistic physical system, where single particles diffuse significantly faster than large clusters. The objective of the following analysis, is to mainly compare with collective motion, and to spotlight the coarsening mechanisms in the limiting case where only single particles are allowed to move. This will allow us to analyse the self-assembly and coarsening phenomena in more physically realistic systems as a superposition of processes and growth regimes which are consequences of either collective or single-particle motion.

We will investigate the coarsening of the attractive model with potentials in Eq. (2.45), quenched to the temperature $T = 0.25$. Figure 3.3 shows that the system under single-particle motion coarsens into several separated and long-lived domains. We will have a closer look at the kinetics of this phase separating process. Growth of the domains under single-particle motion can be characterized by three different growth regimes, with distinct growth rates.

The first stage of coarsening under single-particle motion, is characterized by a time window with a large number of partitions (Fig. 3.4(a)) and with the dominant number of particles having less than five orientational bonds (Fig. 3.4(c)). A high number of partitions implies a large surface, which means low stability from the thermodynamic point of view, and low repulsive barrier to the detachment from the kinetic point of view (as long as the surface is mostly convex). The reorganization is fast: weak attachment and large surface means that particles can easily detach, diffuse, and attach [Grant et al., 2011; Grant and Jack, 2012]. The time window is characterized by a fast and linear (not shown) increase in the number of particles in crystals and in LSCs with and without a crystal core inside. The total number of

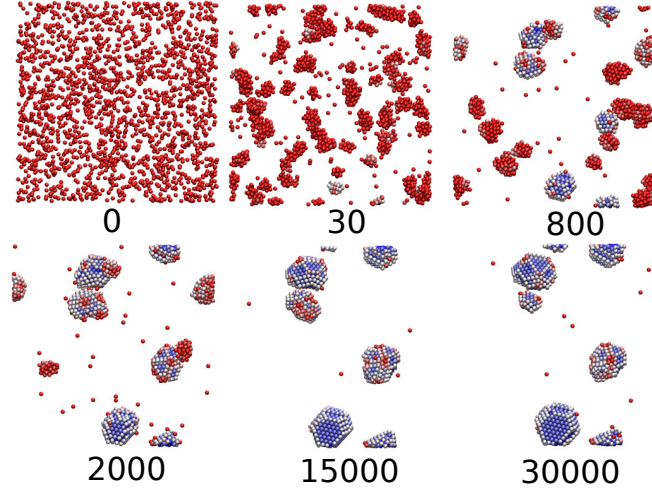


Figure 3.3: Time evolution of the system under single-particle motion. The particles are unlikely to escape the crystalline droplets which themselves hardly move. The time is measured by the number of MC sweeps (1 time unit corresponds to 1000 MC sweeps). Particles are coloured according to the number of orientational bonds n_b ascendingly from red via white to blue.

crystals and LSCs of both kinds increases rapidly and reaches a maximum, signalling the end of this time window (compare also with Figs. 3.4(a) and 3.4(b)). The end is also characterized by approximately equal concentrations of weakly bonded and crystalline particles (Fig. 3.4(c)). Figures 3.4(e) and 3.4(f) imply that LSCs can contain more than one such crystalline domain, and that new isolated crystalline domains generally nucleate only during this time window, and not at later stages of the simulation.

The second stage of coarsening is dominated by dissolution of non-crystalline partitions. Figure 3.4(c) shows that the concentration of weakly bonded particles is gradually exceeded by stronger bonded particles ($5 < n_b < 10$). Growth of the largest domains is slower than in the previous regime, because the number of partitions (and hence the total surface) is decreasing, and particles are attached to their partitions by an increasing number of bonds. Nevertheless, the growth is faster than in the upcoming time window, because the largest domains grow not only by adsorption of smaller crystals, but mainly because of the presence of low symmetry clusters without a crystal inside or even of lower ordered partitions. Indeed, Fig. 3.4(f) shows that the number of LSCs without a crystal core undergoes a rapid and distinct period of decrease, and that most partitions which survive this window have a crystal core inside. Figure 3.4(d) then implies that crystals within

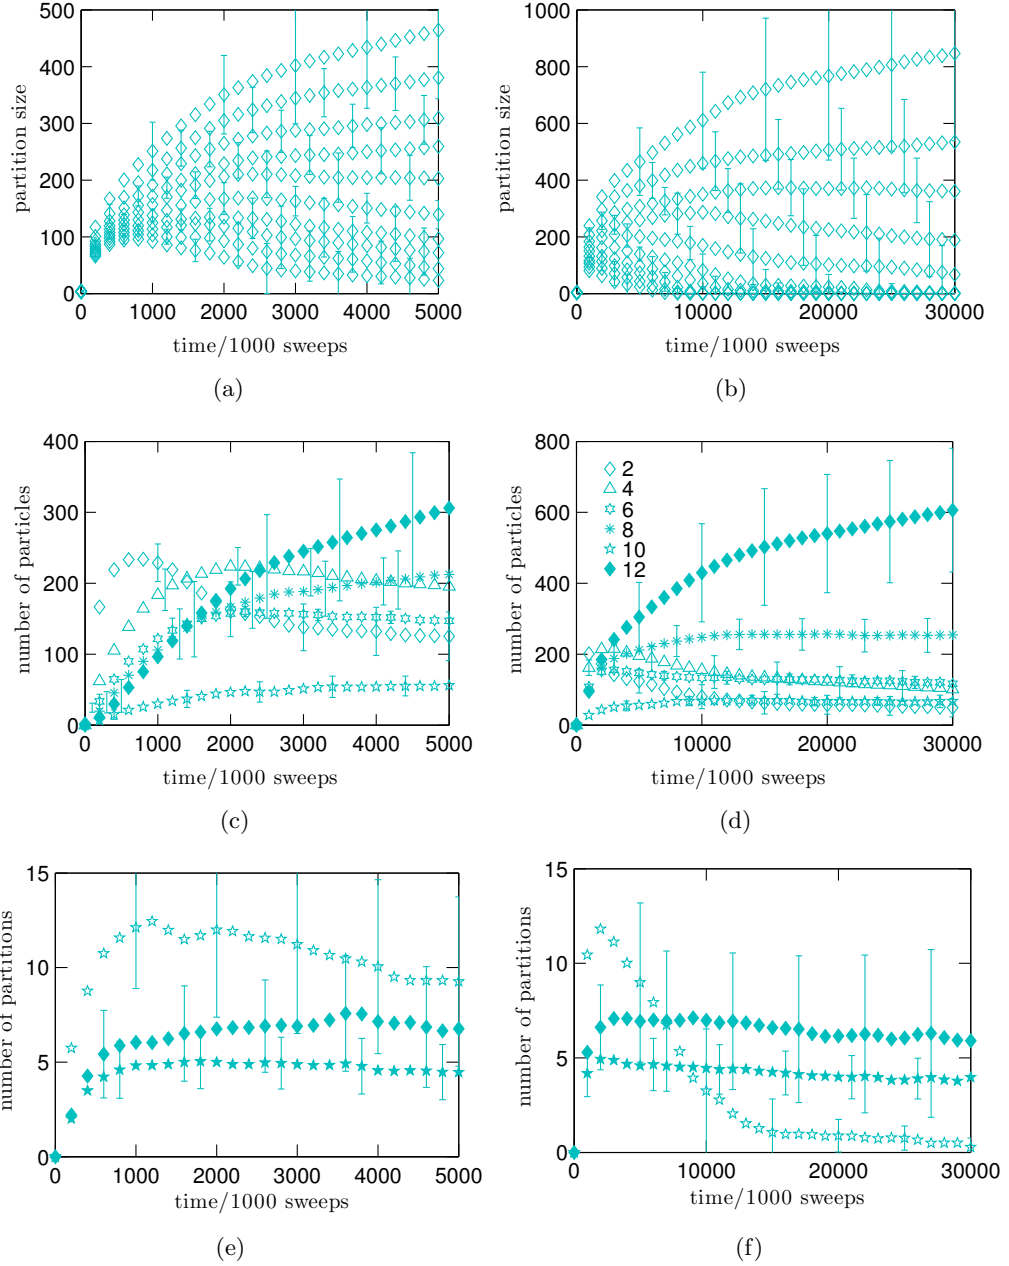


Figure 3.4: Time evolution of long- and short-range structure in a system of attractive particles without long-range repulsion which is quenched and evolves under a single-particle motion with maximum displacement $\delta = 0.30\sigma$. (a), (b) Size of the first ten largest partitions. (c), (d) Local order measured by the number of particles with a fixed number of orientational bonds given in the legend. (e), (f) Number of crystals (diamonds), low symmetry clusters (empty stars), and low symmetry clusters with crystals inside (full stars). Figures (a), (c), and (e) (on left) are short time magnifications of results displayed in Figures (b), (d), and (f) (on right). All error bars are averages from ten independent simulations.

LSCs grow within this intermediate time window. The window exists due to the fact that the timescale needed for crystallization is much longer than the timescale needed to dissolve the low-structured partitions.

The final stage of coarsening is a growth region with a low concentration of weakly bonded particles. The averages in Fig. 3.4(b) show that the system has several growing and several dissolving partitions. The number of particles with six and eight orientational bonds does not change, indicating that the crystalline surfaces and edges of the domains are well developed. The average number of particles in LSCs is also approximately constant, pointing to the fact that large partitions grow by adsorbing the smaller ones. The growth rate is again smaller than in the previous window, because the overall surface area of the partitions is smaller and because, to detach a particle from a well formed surface, more bonds need to be broken than to detach a particle from a low symmetry cluster or from the disordered (liquid) surface of a crystal. The growth rate of the largest domain is linear on the timescale of our simulation window, and crystals within LSCs keep growing linearly but at a slower pace than in the previous window (Fig. 3.4(d)).

3.4.2 Displacement size in single-particle MC

The first and second growth regime is also recovered for smaller displacements δ . The onset of the third growth regime was not achieved by MC simulations with $\delta \geq 0.20\sigma$ on the timescales of our simulations, but it is reasonable to expect a similar behaviour. Figure 3.5 shows that the growth rate of local structures and of the largest partitions discussed previously for $\delta = 0.30\sigma$ scales as follows. If t_1 is the time of crossover between the first and the second growth regime in a single-particle MC simulation with maximum displacement δ_1 , then the time t_2 of the same crossover for another δ_2 can be estimated as

$$t_1\delta_1^2 = t_2\delta_2^2. \quad (3.7)$$

The total number of crystalline or weakly bonded particles, and even the corresponding error bars follows the same scaling; recall that t_1 and t_2 are measured in MC sweeps. For example, given a total number of crystalline particles at t_1 and δ_1 , one should expect the same number of crystalline particles at t_2 and δ_2 given by Eq. (3.7).

Similarly to the time step dependent scaling in Molecular Dynamics simulations, δ dependent scaling in VMMC cannot be expected to be exact. For instance, the peaks in the number of weakly bonded particles have different heights for dif-

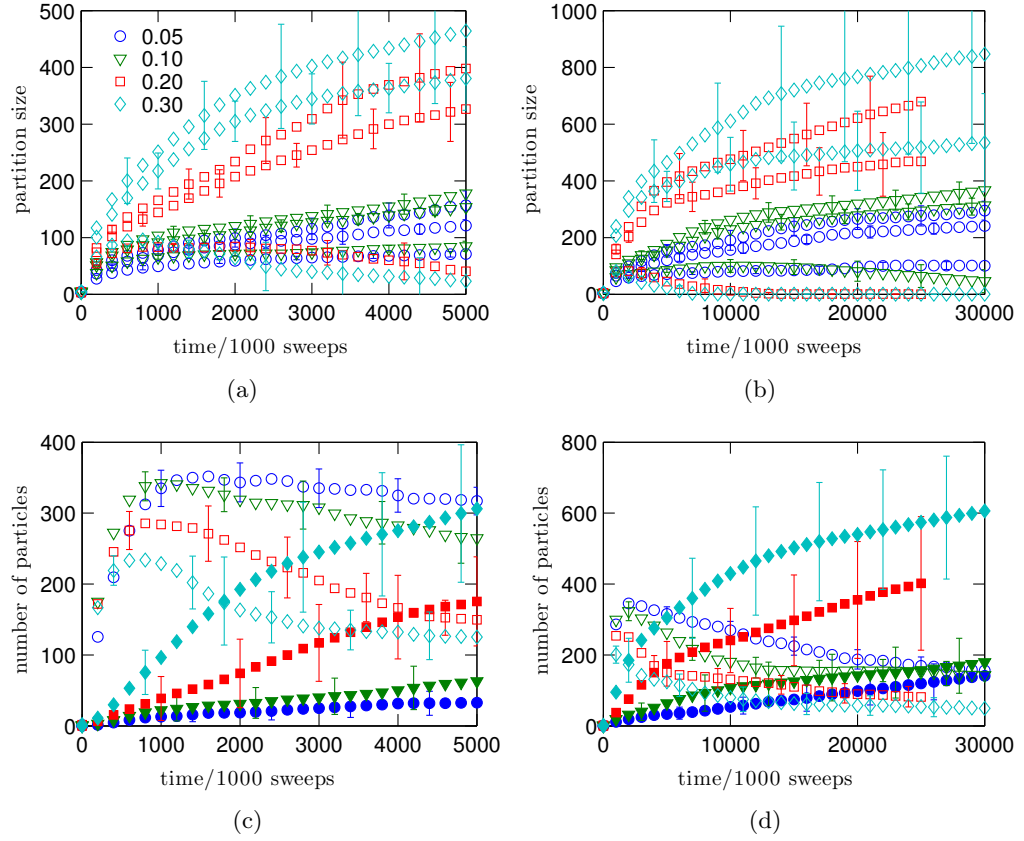


Figure 3.5: Time evolution of long- and short-range structure in a SPMC, and its dependence on the maximum displacement size δ given in the legend in a low-density and short-range attractive system following a quench. (a) (b) Size of the first, second and tenth largest partitions. (c) (d) Number of weakly bonded particles ($n_b = 2$) are represented by empty symbols and crystalline particles with full symbols.. Error bars are from ten independent simulations. The plots on the right are long-time data of simulations on the left.

ferent δ . In fact, we will also see for the VMMC, that well-ordered local structure generally evolves faster for larger δ , and that the less ordered local structure evolves faster for smaller δ . However, the scaling is reasonably approximate, and reinforces the possibility that a MC cycle can be associated with a rough timescale even in non-equilibrium simulations of quenched systems where phase separation and crystallization is underway.

We note that if we assume that SPMC simulation behaves like a diffusive process [Allen and Quigley, 2013], and that the MC acceptance probabilities for δ_1 and δ_2 are similar, we may write that

$$\tau = t_1\tau_1 = t_2\tau_2, \quad (3.8)$$

where τ_1 and τ_2 are physical timesteps, and τ is the physical time. The diffusion equation then implies

$$6D\tau = t_1\delta_1^2 = t_2\delta_2^2, \quad (3.9)$$

which is nothing else but Eq. (3.7), relating the number of sweeps to the physical time. Equation (3.9) allows one to plot the observables as a function of the physical time, and highlight the deviations from the diffusive scaling. This has not been done in this thesis.

3.4.3 Phase separation under collective motion

Let us now turn to the simulation of phase separation where collective motion is allowed. The probability of linking a particle to the moving cluster is proportional to the energy gradient [see Eqs. (2.21) and (2.15)]. For the system under consideration, if the maximum size δ of MC displacements is smaller than the width of the attraction, the pairwise energy gradient is small and links are unlikely to form. Clusters generated in the VMMC then have only one or a small number of particles, and the dynamics is similar to the dynamics under single-particle motion (compare Figs. 3.5 and 3.11). A more significant collective motion appears, only when δ is comparable to or larger than the attraction range. Figure 3.6 shows that the growth of the largest partition is then many orders of magnitude faster than in the SPMC simulation, but local order still evolves at about the same pace. This means that the crystallization pathway is not generated faster or in a more approximate or coarse-grained fashion than in the SPMC but is simply different.




Figure 3.6: The time evolution of a system where collective motion is allowed leads to the rapid formation of amorphous droplets, which coalesce into a disordered aggregate. Crystals within this aggregate develop later and are separated by a disordered (liquid) surface.

Aggregation in VMMC

Let us have a closer look at the average structure and the pathway leading to a single large partition obtained by VMMC. We first consider the timescale determined by the average time needed to merge most particles of the system into a single partition. Figures 3.7(a) and 3.4(b) show that this timescale is many orders of magnitude smaller under the collective than under the single-particle motion. In fact, the timescale of complete phase separation by SPMC was not even achieved. Figure 3.7(a) shows, that similarly to the SPMC, the first stage of the simulation in VMMC, leads to a rapid clustering within the system into a large number of small partitions. However, contrary to the SPMC, small partitions not only coarsen but also move easily, and the growth proceeds not only via single-particle exchange, but also via partition-partition attachment. In fact, the latter mechanism is so fast that the three linear growth regimes characteristic of SPMC simulation (Figs. 3.4(a) and 3.4(b)), and entirely suppressed. The main reason for this behaviour is that the timescale for detachment of a single particle from its partition and its diffusion to another partition is comparable to the timescale needed for the whole partition to diffuse that distance. Such dynamic equilibrium was defined as clustodal by Whitelam et al. [2009a]. Since small partitions diffuse faster than large partitions, we can generally say that large partitions grow by adsorbing the smaller ones.

After a very short time, the second largest partition reaches its peak size, and starts dissolving, at the expense of the largest one. The dissolution rate is significantly faster than in the SPMC. Figure 3.7(c) shows that this is again likely to be caused by a small number of bonds with which a particle is attached to its partition. The collective motion, allowing whole groups of particles to detach from their partitions and diffuse, and a low repulsive kinetic barrier [Zaccarelli and Poon, 2009] in locally dense disordered regions, may also accelerate the dissolution process. The dissolution then on average stops, and the partition size reaches a plateau before the second largest partition dissolves or before it merges with the largest partition. The sudden drop in the partition size suggests that the latter mechanism is more likely. Large partitions generally survive longest, because they diffuse slowly, taking longer for them to merge with another partition, and because they are more likely to contain a crystal or a LSC which take longer to dissolve. Indeed, Figs. 3.7(a) and 3.7(e) indicate that the largest partition contains on average two crystals at about the same time as the second largest partition merges with it.

Figures 3.7(a) and 3.7(e) show that after some time ($50 \cdot 10^3$ VMMC sweeps), there are on average only two partitions whereas there are more than six LSCs, and their number is growing even if the number of partitions remains the same.

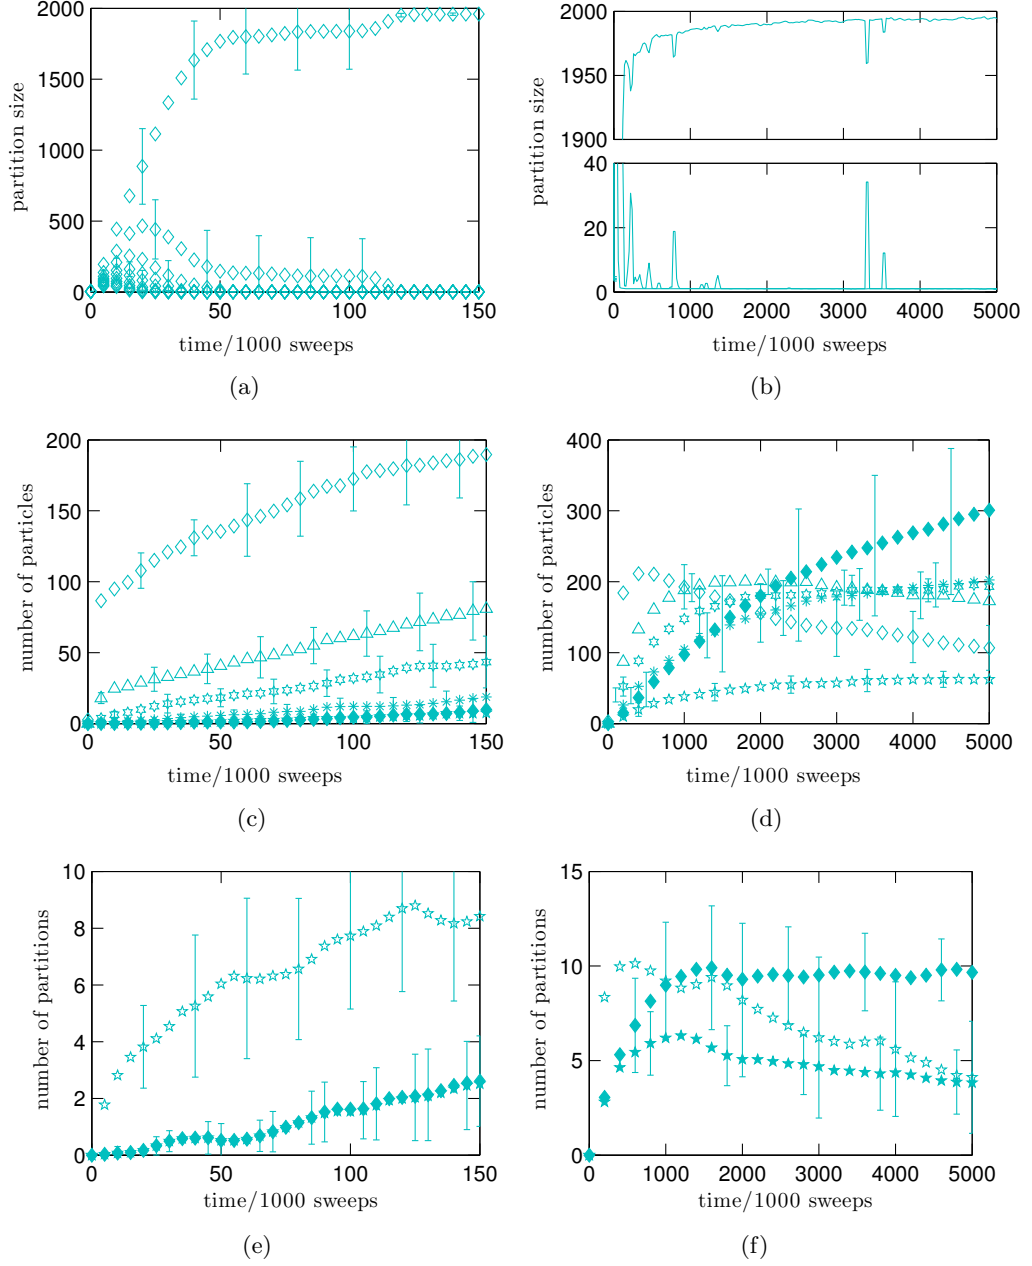


Figure 3.7: Structural evolution in a quenched short-range attractive fluid simulated with VMMC simulation where collective motion is allowed. Figures and symbols have meaning analogous of that in Fig. 3.4.

Figure 3.7(c) shows that this initial time window is dominated by particles embedded in partitions with weakly bonded particles. These results confirm that the local order evolves within the amorphous drops. Equal numbers of crystals and LSCs with a crystal in Fig. 3.7(e) means that such LSCs only contain one such crystal in the initial stages of the simulation. Figures 3.7(c) and 3.7(e) imply that the crystals are small with less than ten crystalline particles.

Figure 3.7(b) shows the long time evolution in the size of the first, second, and the third largest partition within the system. The fluctuations in the size of the first and second largest partitions indicate detachment and attachment of larger aggregates from the largest partition. The fluctuations decay with time as a consequence of an increasing crystallinity within the system making these detachments less likely. The size of the third largest partition is on average one particle, which points to the coexistence of the largest aggregate with gas. The size of the largest partition slowly increases. This is a kinetic phenomenon, resulting from the fact that the system is not in thermodynamic equilibrium, because the evolution of crystallinity depends on the interface between the gas and the aggregate, and because the detachment and attachment of single particles is becoming more difficult, due to an increasing attractive and repulsive barrier in the denser parts of the system.

Figure 3.7(f) implies that LSCs can eventually contain more than one crystal. The number of crystals, and the number of LSCs with crystals inside, stop growing, while the number of LSCs without a crystal starts decreasing. This indicates that no new crystals form within the LSCs, and that LSCs without a crystal core gradually dissolve at the cost of other growing LSCs with a crystal in them. Rearrangements at this late stage resemble coarsening rather than assembly, and happen without significantly distorting the large aggregate via single-particle or small cluster motion. Figure 3.7(b) shows that large collective rearrangements via partition division at this stage are less likely and rare, because they are associated with large and unlikely energy changes.

Finite size effects

It is shown in Figs 3.8 and 3.9 that the finite size effects at this scale are negligible, and the observables are linearly proportional to the system size. The larger system under consideration has $N = 8000$ particles, and was quenched to the same temperature and conditions as before. We run simulations of six independent systems for $N = 8000$ particles, compared to ten systems for $N = 2000$, which may explain why the error bars in the larger system are larger for some observables. The size of the maximum displacement is $\delta = 0.20$ for both system sizes. It is worth mentioning

that finite size effects may emerge at the different parameters of the simulation, which were not examined in this thesis.

Aggregation in VMMC after deeper quenches

Figure 3.10 shows kinetic evolution of the system under a deeper quench to temperature $T = 0.10$, and using maximum displacement $\delta = 0.20$. It can be seen that the system forms very rapidly a low density network (time 1), which collapses (time 2), and remains in this macroscopic state for the rest of the simulation, where the local structure of the aggregate evolves by very slow coarsening of the branches. In comparison with the structure shown in Figure 3.6 for higher temperature, the kinetically slowed down state has distinct large scale ‘holes’, and properties of a gel, although we cannot call it a gel, since it does not percolate the periodic box [Toledano et al., 2009]. Its ramified structure may be characterized by fractal dimension [Sciortino et al., 2005], which should be lower in Fig. 3.10 than in Fig. 3.6. The quench of the same system at this temperature was also examined by Perez et al. [2011] using Brownian Dynamics, with the resulting structure being visually comparable. This gives another evidence that the VMMC dynamics leads to physically realistic kinetically slowed down aggregates.

Aggregation in VMMC under different δ

Let us now have a closer look at the scaling and time evolution of global and local observables in a VMMC simulation for different δ . First, we examine the macroscopic aggregation. Figure 3.11 shows again the time evolution of the first, second, and the tenth largest partitions in the system, and that the system merges into a single partition on the timescale of our simulation as long as $\delta \geq 0.10\sigma$. Generally, the larger the δ , the less time it takes for the system to coalesce. This scaling can be easily seen from the peak size of the second largest partition and its position in time, which also shows that higher peaks correspond to larger values of δ . (A possible way of identifying an optimum δ might be to require that the peak is reached at $N/2$.) It is illustrative to point out the strength of this dependence. The time difference between the peaks of $\delta = 0.20\sigma$ and $\delta = 0.30\sigma$ is small relative to the difference between $\delta = 0.20\sigma$ and $\delta = 0.10\sigma$. If $\delta = 0.05\sigma$, the peak lies far beyond the timescale of our simulation. This is a consequence of our linking method, implying that clustering occurs only for δ larger than the width of the attraction range (at least for our choice of $\epsilon/k_B T$.) Long-time data for $\delta = 0.30\sigma$ are not displayed, because the simulations were too slow.

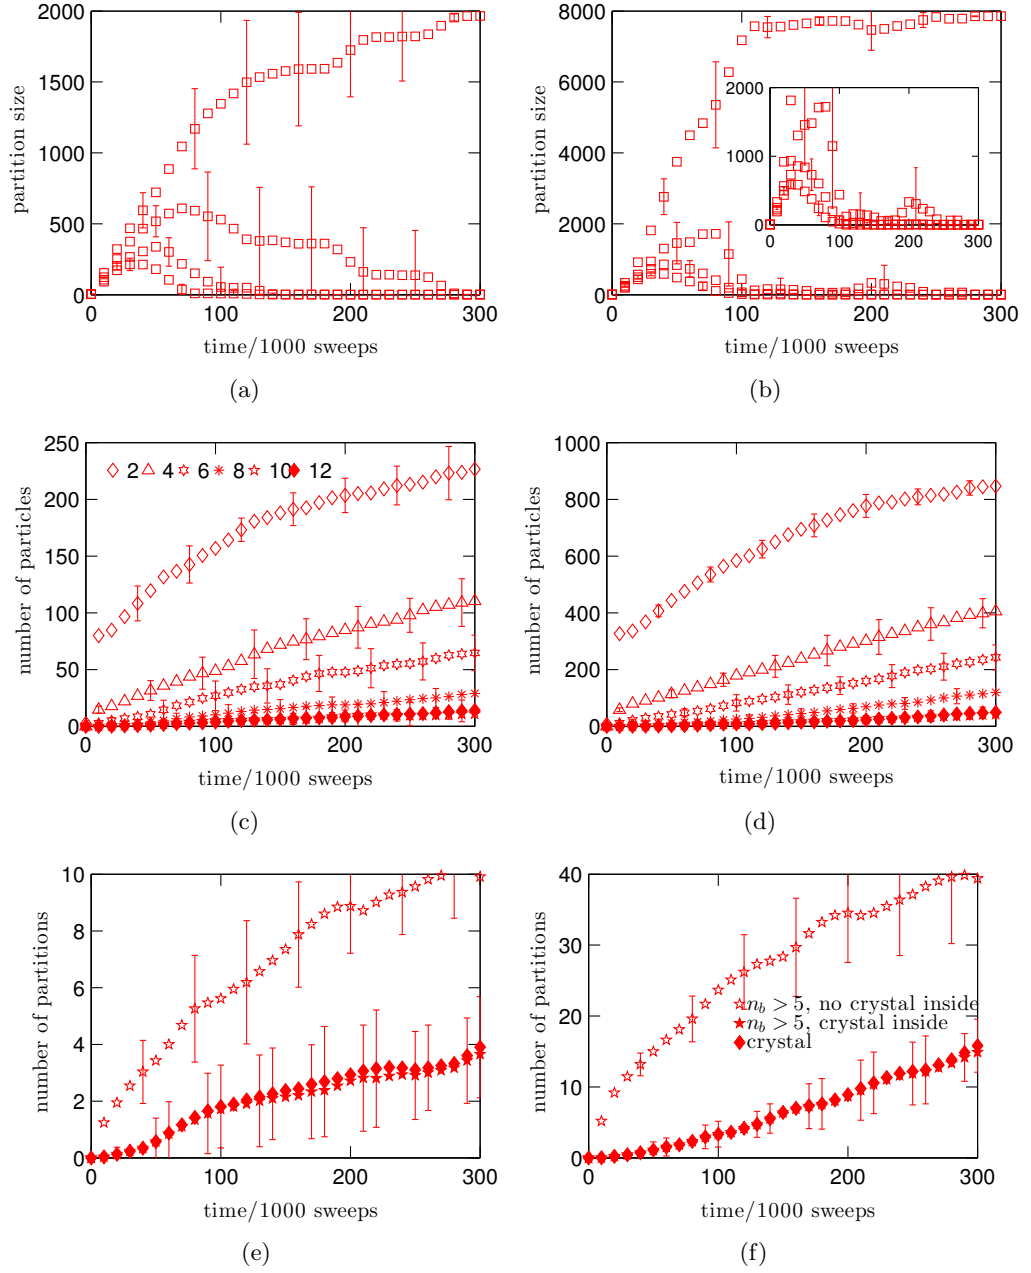


Figure 3.8: Finite size effects on a short timescale. Symbols have the same meaning as in Figs. 2.5 or 3.4, except that (b) shows the size of four instead of ten largest partitions. Plots on the left correspond to the system size $N = 2000$, and plots on the right to $N = 8000$. The size of the maximum displacement is $\delta = 0.20$ for both system sizes.

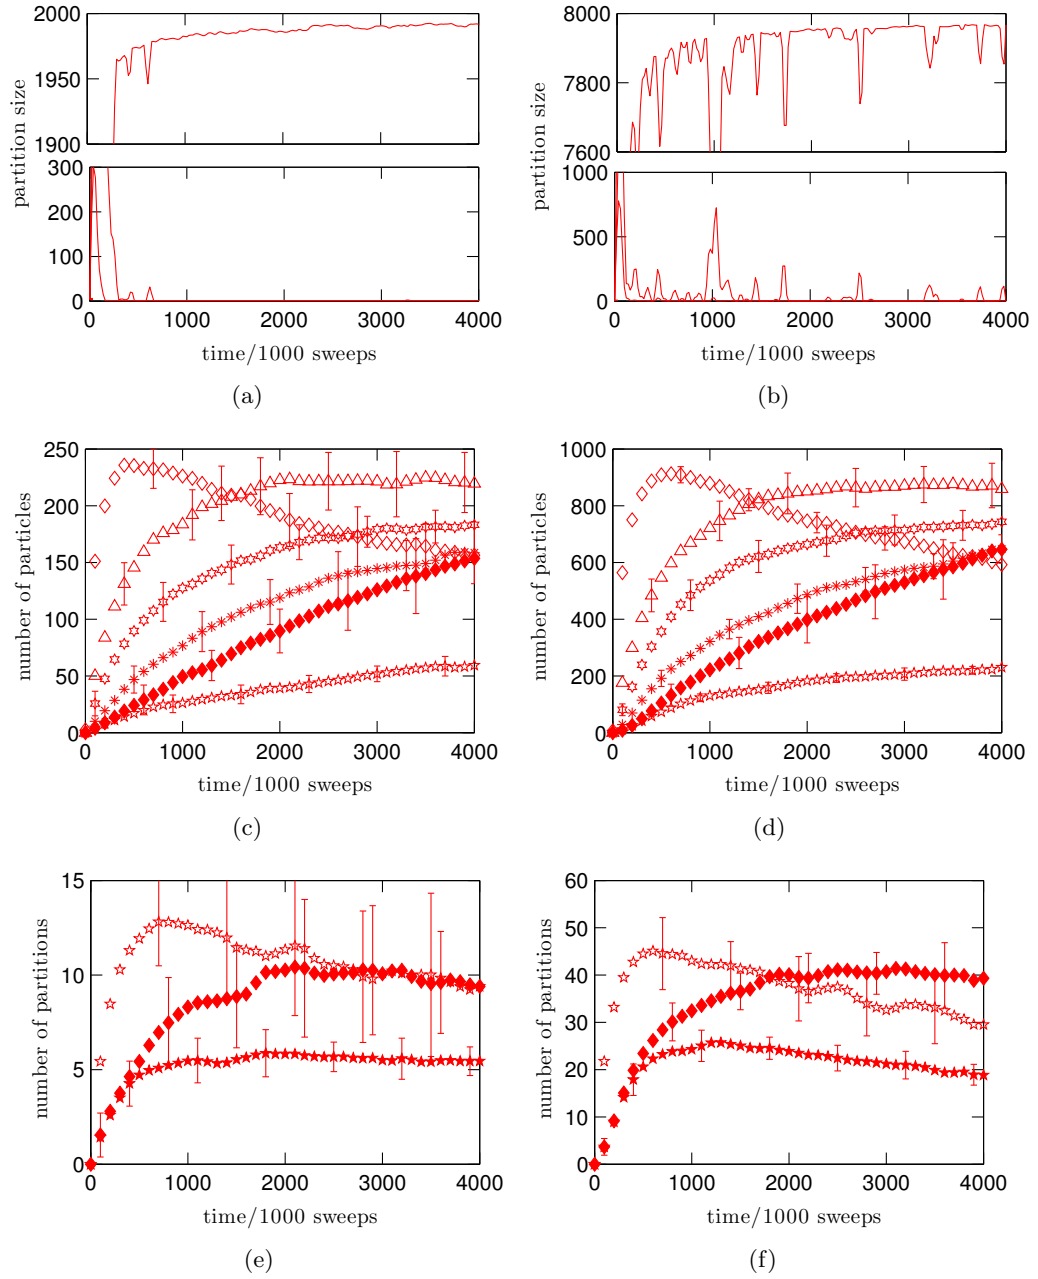


Figure 3.9: Finite size effects on a longer timescale. Meaning of the plots is the same as in Fig. 3.8 for shorter timescales.




Figure 3.10: Kinetic evolution of a system quenched deeply to the temperature $T = 0.1$. The colouring has the same meaning as in Fig. 3.6.

Single-particle rearrangements within the initial large and low-density aggregate are easy, because the moving clusters are selected through the local energy gradients, and because particles can travel along the large surface of the aggregate. The initial crystalline structure of the low-density aggregate can be quantified by the number of crystals and by the number of LSCs with and without a crystal core inside. The time evolution of these macroscopic observables is shown in Fig. 3.11. Figure 3.11(c) shows that, despite the large differences in the initial dynamics observed in Fig. 3.11(a), the evolution of structure under these global observables is independent of δ at the initial times. Figure 3.11(f) then shows that at late time the observables scale with δ . The scaling is again not exact but reasonably approximate. The number of LSCs without a crystal is shifted in time and decays at the same rate for different δ . It can be seen that the number of crystalline domains is independent of δ at long times, confirming that the final structure of the aggregate can be expected to be similar for different δ . Similarly to SPMC, the maximum number of crystals and LSCs with and without a crystal core inside is reached at about the same time as the time where the concentration of low- and high-crystalline particles meet (compare Figs. 3.11(d) and 3.11(f)). This peak is reached slightly earlier in VMMC than in SPMC, reflecting the fact that the time is measured in MC sweeps, and that a VMMC sweep generally moves with a larger number of particles than a SPMC sweep. The final number of isolated crystals is slightly higher in VMMC than in SPMC, suggesting that the presence of amorphous drops enhances crystal nucleation.

Interestingly, a similar number of LSCs with at least one crystal inside are observed at long times both in VMMC and in SPMC, even if one system is merged into a single partition and the other is not. This suggests that crystals grow from

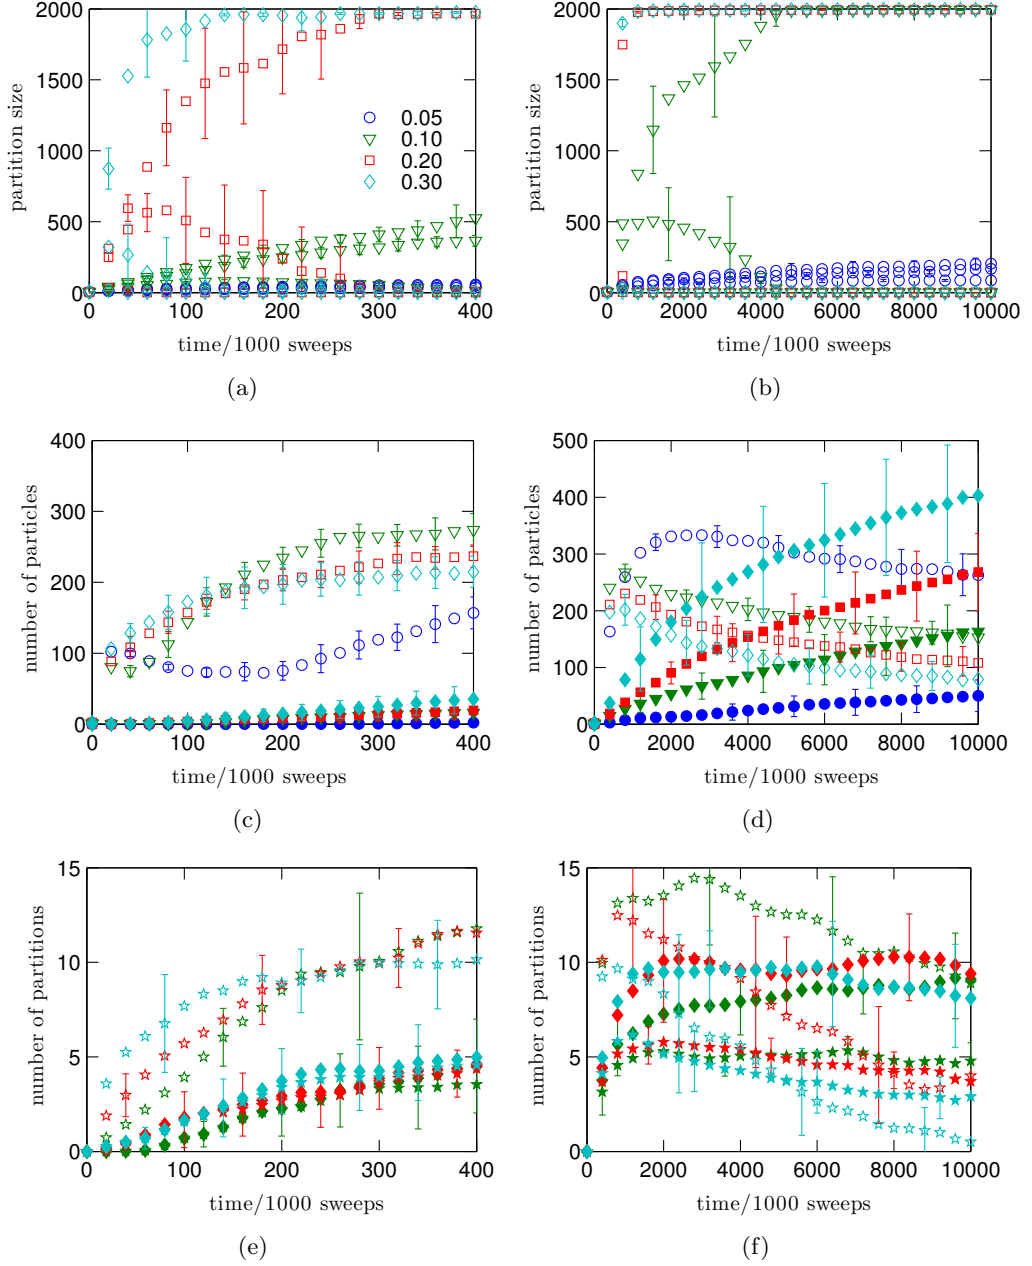


Figure 3.11: Dependence of structural evolution on the choice of maximum displacement in the VMMC simulation. Figures (a), (b), (c), and (d) have meaning analogous to that of the SPMC in Fig. 3.5. Figures (e) and (f) show the number of LSCs with no crystal inside (empty stars), the number of LSCs with a crystal inside (full stars), and the number of crystals (diamonds). Different colours correspond to different values of δ given in the legend of (a). All error bars are averages from ten independent simulations. Simulations with $\delta = 0.30\sigma$ were too slow, and did not reach the end of the time axis.

the amorphous drops independently of whether these drops are merged or isolated. Contrary to the SPMC, we did not reach timescales where most LSCs without a crystal inside dissolve. However, from Fig. 3.11(f) we can see that the number of LSCs with a crystal inside keeps decreasing at long times, while the number of crystals remains the same. This suggests that very weakly bonded particles, forming an interface between two separate LSCs with crystals inside, slowly become stronger bonded particles, and eventually produces a single LSC with multiple crystalline cores. Also note, that the length scale of clusters in Fig. 3.3 and the width of the branches (or subpartition sizes) in Fig. 3.6 may be seen to be similar at similar times. This may again be a kinetic (non-equilibrium) property, because we know that the thermodynamically stable state is the fluid-solid coexistence, and the reason for this similarity may be again a similar mechanism behind the coarsening, i.e. growth via single particle detachment and attachment at this early stage of the crystallization process. It would be interesting to compare the length scales of the crystals in Fig. 3.6 at later times, because, as mentioned before, the single particle motion may be facilitated by more energetically favourable motion along the surfaces of the aggregate.

Let us now look at the time evolution of local order quantities. Figure 3.11(d) shows that the scaling of local crystal order has a qualitative behaviour, which is similar to the one observed in the SPMC, and this is despite the fact that the trajectory is different when quantified with macroscopic observables such as the size of the largest partition. This result may be a consequence of cluster generation via local energy gradients. Qualitatively, concentrations of weakly bonded and crystalline particles are equal at earlier times than in SPMC. A plausible explanation for this behaviour is again that time is measured in MC cycles, rather than being proportional to the number of particles in proposed (or accepted) moves. Generally, the total number of moved particles in a VMMC cycle is greater than in a SPMC cycle. It might be possible to obtain a scaling relation similar to Eq. (3.7), by integrating over the total number of particles in proposed (moved) clusters. An insight into such relations may improve the association of a timestep with a VMMC cycle [Haxton and Whitlam, 2012]. A detailed treatment of this problem is beyond the scope of this work.

We no longer observe any connections between the growth of the largest aggregate and the local structure as is seen in SPMC simulations. We rather observe that the average dynamical trajectories are different for different δ to an extent where they cannot be considered comparable. This can be seen if we take $\delta \geq 0.20\sigma$ and $\delta \leq 0.10\sigma$. For $\delta \geq 0.20\sigma$, the number of weakly bonded particles peaks long

after the peak of the second largest partition. If $\delta \leq 0.10\sigma$, the weakly bonded particles and the second largest cluster reach their peak at about the same time. This simple example points to one of the central results of this section concerning the choice of δ in the VMMC. For smaller maximum displacements δ , isolated aggregates move slowly, and have more time to develop a local crystal structure, before they merge with other aggregates. For larger δ , aggregates move faster and coalesce well before a significant local order develops. The coalesced partition consists of weakly bonded particles. Figures 3.5 and 3.11(d) then show that the crystallinity within the partition evolves at later times at about the same pace as in SPMC. Figures 3.11(c) and 3.11(d) also show that, similarly to SPMC, the growth in the number of weakly bonded particles is faster for smaller δ ; however, the peak position in the number of weakly bonded particles for VMMC is nearly independent of $\delta \geq 0.10\sigma$, and the peak for $\delta = 0.05\sigma$ is not far from them. Whether the dynamics scales better for different smaller $\delta \leq 0.05\sigma$ still needs to be verified.

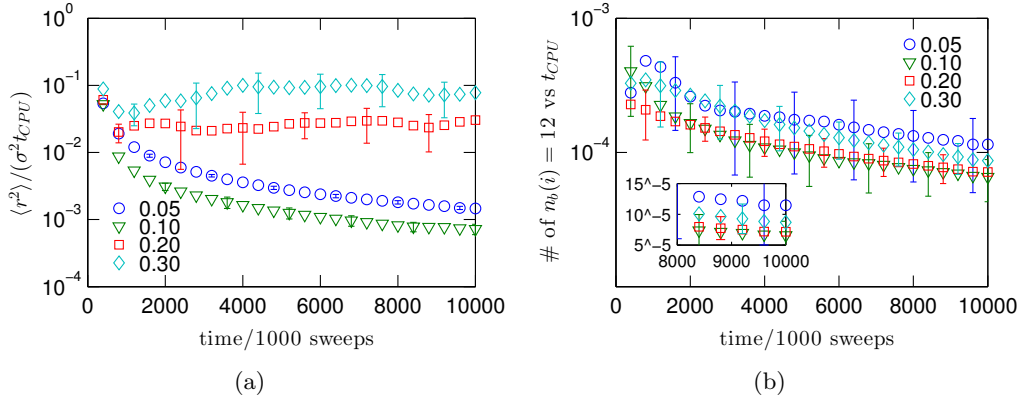


Figure 3.12: Efficiency of the VMMC with respect to the choice of maximum displacement δ (specified in the legend) in the simulation of kinetically slowed down phase separation. (a) efficiency measured as the mean square displacement vs total CPU time in seconds. (b) efficiency measured as the number of crystalline particles with $n_b(i) = 12$ vs total CPU time.

The choice of large δ may seem to be more efficient, because both local and global structure evolve faster (in a lower number of sweeps). However, the larger the δ , the longer it takes for the computer to complete a MC sweep. This is partly because of the larger number of particles to which a link is proposed, but mainly because MC sweeps contain many cluster moves of size larger than $N/2$, which are rejected. The construction of such moves results in a lot of waste, and slows down the simulation. Larger error bar in the number of crystals (Fig. 3.11(d)) indicate that choosing a large δ may lead to a more diverse and less realistic microscopic

dynamics. If the aim is to achieve the maximum local order while not neglecting the collective phenomena leading to the aggregation, there is thus an optimum δ . Fig. 3.12(b) shows that the CPU time needed to produce a given number of crystalline particles does not differ significantly for different δ . Small displacement size $\delta = 0.05$ performs best in this respect, however, Fig. 3.11(b) shows that the partitions remain isolated, and the motion resulting from the large scale collective aggregation is not sampled sufficiently fast (Fig. 3.12(a)). Again, $\delta = 0.30$ would perform best, because both the number of crystals and the local structure evolves fast even with respect to the CPU time. Nevertheless, the difference in efficiency between $\delta = 0.20$ and $\delta = 0.30$ is only small, and since we can expect that the microscopic dynamics for smaller δ would be more realistic, we choose $\delta = 0.20$ as the optimum displacement. For this choice of δ local order develops sufficiently fast, and the collective motion leads to the aggregation into a single domain. This choice corresponds to the width of the attraction. Indeed, Fig. 3.12(a) shows that there is a distinct jump in MSD for δ below ($\delta < 0.20$) and above ($\delta \geq 0.20$) the attractive range. A similar choice of δ has been made elsewhere [Whitelam et al., 2009a; Fusco and Charbonneau, 2013]. It can be expected from our previous considerations that the optimum δ generally also depends on other quantities determining the onset of gelation [Lu et al., 2008], which include $\epsilon/k_B T$ or volume fraction.

3.4.4 Generation and acceptance densities

Trajectories of the aggregation process studied in Fig. 3.11 were further analysed by means of the probability densities of generating and accepting a cluster. These probabilities evolve during the simulation. Results are shown in Fig. 3.13. For each δ we monitored the probability density of generating and accepting a cluster of a given size in different time windows of the simulation. For $\delta \geq 0.10\sigma$, the windows are determined by distinct times, characterizing the size of the second largest partition. These times correspond to the start of the simulation, the peak size of the second largest partition, and two other state points at later stages of the simulation, where the second largest partition is dissolved, and where the system is aggregated into a single large partition. Since the size of the second largest partition does not possess any such features for $\delta = 0.05$ on our simulation timescale, we distributed the time windows exponentially for that δ . We firstly discuss properties of f_g and f_a for values of δ which are not lower than the range of the attraction ($\delta \geq 0.10\sigma$). The case $\delta = 0.05\sigma$ is addressed separately. Densities f_g and f_a are different for rotations and translations, and we will distinguish between them.

We will analyse the dependence of f_g and f_a on time δ , and on n_C . Let us now

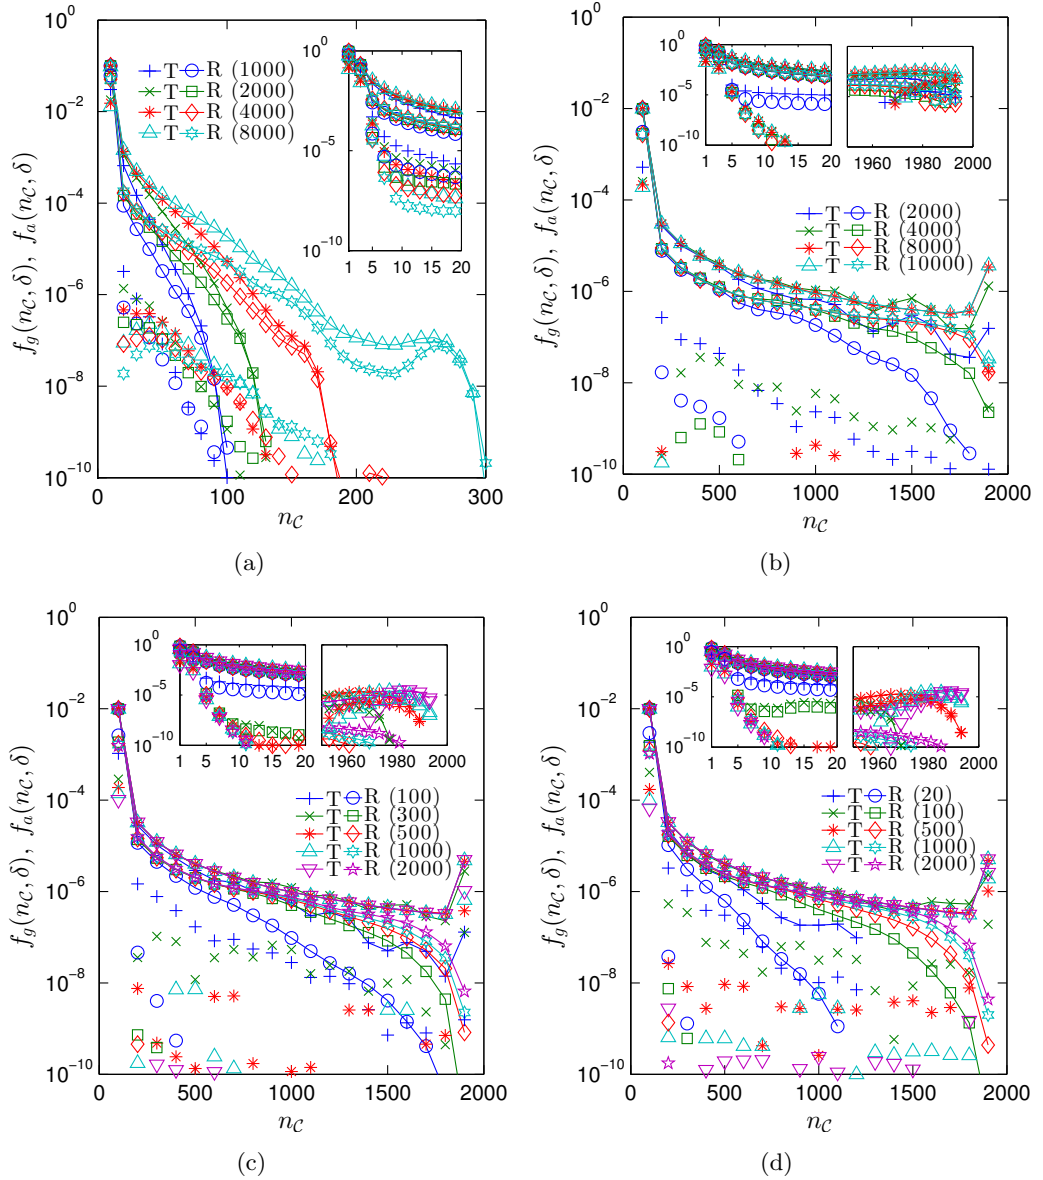


Figure 3.13: Probability densities of generating and accepting a moving cluster of size n_c at different time windows of the simulation. Connected symbols, and identical but unconnected symbols represent proposed and accepted moves, respectively. These can be easily distinguished because $f_a \leq f_g$. Symbols corresponding to translations (T) and rotations (R), together with the upper limit of the time window (in MC sweeps) are given in the legend. The same colour refers to the same time window. Data in (a), (b), (c), (d), correspond to δ equal to 0.05σ , 0.10σ , 0.20σ , 0.30σ , respectively. The symbols represent arithmetic averages over a cluster size window with radius equal to the distance between the symbols. All data are averages from ten independent simulations. Error bars can reach several orders of magnitude, and are not displayed. Insets are magnifications for small and large moving clusters.

discuss f_g for $\delta \geq 0.10\sigma$. If $n_C \ll N$, the density of proposed clusters f_g has similar behaviour both for translations and rotations, with f_g being slightly greater for translations, and with both being decreasing functions of n_C . If $n_C \rightarrow N$, f_g increases with n_C for translations and keeps decreasing for rotations. The reason for this will be discussed later. The density of proposed clusters f_g generally rapidly evolves in time towards larger sizes of the moving clusters, especially for $n_C \rightarrow N$, both for rotations and translations. This reflects that partitions in the system coarsen, and that the cluster size is proportional to the partition size. Rapid evolution of f_g slows down at later times, especially when the system coalesces into a single partition. The density f_g at those times has a similar behaviour for all δ .

For translational moves, proposed clusters can stop growing due to the condition in Eq. (2.46), to approximate the moving cluster size distribution in Brownian dynamics. This leads to

$$f_g(1, \delta) \approx \frac{1}{n_C} f_g(n_C, \delta), \quad (3.10)$$

for sufficiently large n_C , and a suitable δ . If $n_C \rightarrow N$, then $f_g(n_C, \delta) = 0$, which is a consequence of the coexistence of the gas with the aggregate.

We note that if the cluster size selection is not restricted by the condition in Eq. (2.46), then all clusters with no frustrated links in the boundary are accepted, and $f_g(n_C, \delta)$ is roughly symmetric about $N/2$ for translational moves, i.e.

$$f_g(n_C, \delta) \approx f_g(N - n_C, \delta), \quad (3.11)$$

for $1 \leq n_C \leq N/2$, and for a suitable δ . This can be seen in Fig. 3.14, although δ in this case should be chosen lower for the equality in Eq. (3.11) to be more precise. Note that Eq. (3.11) fits very well the average number of translational and rotational moves in the situation shown in Fig. 3.14, which is a desirable property to control the distribution of moving clusters and for the approximation of the dynamics.

Large clusters are more likely generated by translations than by rotations for all δ and for all studied time intervals, and the difference becomes larger for large $n_C < N$, where f_g increases for translations and decreases for rotations. The difference is mainly caused by the fact that the rotational clusters, as opposed to translational ones, stop growing if their size exceeds half of the simulation box size, or if particles to be included in \mathcal{C} are too far from the centre of rotation. As for the acceptance density, rotational moves are constructed such that particles that are far from the centre of rotation can be moved by a distance larger than δ (but smaller than σ). This leads to a larger pairwise energy gradient, and consequently to a larger linking probability. It can be seen from Fig. 3.13(a) that for small δ ,




Figure 3.14: Generation and acceptance densities of variously sized clusters in a VMMC simulation, where selection of clusters is not restricted by condition in Eq. (2.46). The maximum displacement here is $\delta = 0.20$. and the simulations are started from final configurations of simulations in Fig. 3.13(c), i.e. from ten configurations at times $2 \cdot 10^6$ MC sweeps after the quench. The length of the simulations was 10^4 MC sweeps, and plots are the averages.

where particles are mostly spherical during the simulation, the largest rotational moves are accepted more often than the largest translational moves.

For small clusters ($n_c \rightarrow 1$), the difference between f_g for translations and rotations is relatively small, and besides the reasons discussed for large clusters, a non-negligible contribution to the difference may also arise from the geometry of the potentials, and the geometry of the translational and rotational MC moves. Indeed, pairwise energy gradients produced by rotations are generally smaller than gradients produced by translations. For example, if the centre of rotation lies in the centre of its nearest neighbour, then a suitably defined rotation may lead to an energy gradient equal to zero, and the link has no chance to form. Such a zero gradient situation is not possible with translations of our spherical potentials, which is a possible reason why link formation is more likely for translations than for rotations, and translational clusters should be larger than rotational clusters.

We will now focus on the acceptance density. Clearly, $f_a < f_g$ both for translations and rotations. The statistics for f_a are poor, but allow us to spot the main trends. We firstly focus on f_a of translations. If $n_c \ll N$, f_a decreases with time both for translations and rotations and for all δ . If $n_c \rightarrow N$, f_a increases

in time. In fact, the acceptance rate in that case is large ($f_a \rightarrow f_g$), which is a consequence of an evolving crystallinity on the surface of the large aggregate. Indeed, surface particles are better bonded, and are more likely to become part of the moving cluster rather than forming a kinetic obstacle to the move. The acceptance density also seems to be symmetric about $N/2$. It is reasonable to expect that if the cluster size was not scaled by $1/n_C$, the relation for f_g in Eq. (3.11) might also hold for f_a . The difference $f_g - f_a$ seems to have its maximum at $n_C = N/2$, and minima for $n_C \rightarrow 1$, and $n_C \rightarrow N$. Also, $f_g - f_a$ increases in time for all $\delta > 0.05\sigma$ and all $n_C \rightarrow N/2$, because f_g increases and f_a decreases for these n_C .

The intermediate sized clusters ($n_C \rightarrow N/2$) are accepted mostly before the dissolution of smaller clusters, but their movements can occur even at later stages of the simulation, when most particles are merged into a single aggregate, suggesting that large collective rearrangements within the aggregate are possible. Figure 3.13 shows that these intrapartition cluster moves are more likely for large δ , suggesting again that the initial aggregate has large non-crystalline and low-density regions for large δ , which makes these collective rearrangements possible. Figure 3.11 implies that the peak size of the second largest partition increases with δ , which is why movements of partitions with $n_C \rightarrow N/2$ should happen for large δ at times before the system merges into a single partition.

The acceptance density f_a for $n_C > 1$ is larger for translations than for rotations, which is again caused by large pairwise energy gradients between pairs far from the centre of rotation. The acceptance density f_a for single-particle rotation ($n_C = 1$) is greater than that for translations, because of the geometry of the potentials and the MC map, as was already discussed for f_g .

Note that the plots in Fig. 3.13 must be interpreted carefully, because the averaging is crude. For example, point $n_C = 100$ is an arithmetic average taken over cluster sizes in the range $[1, 200]$, and the huge weight of single-particle accepted clusters is the reason why rotations are, on average, accepted more often than translations for the point $n_C = 100$. If the average was taken over a narrower window of the cluster sizes n_C , we would have obtained the opposite result consistent with that in the inset.

The acceptance density of rotations is non-zero only for $n_C < N/2$, even before small partitions dissolve. For larger n_C the acceptance is zero. Moreover, the rotational f_a is decreasing with time, and for $n_C \rightarrow N/2$, the density f_a tends to be non-zero only at early stages of the simulation. This may be again due to the low-density disordered aggregate which forms rapidly at the beginning of the simulation. Rotations of isolated clusters within that aggregate are unlikely.

Moreover, many rotations are rejected due to the finite size of the simulation box. Similarly to translations, we observe that relatively large rotational moves in the size interval $[100, 300]$, can be accepted even after the dissolution of small partitions. These rotations occur more often for larger δ , confirming again that the initial large partitions are locally disordered. It is possible that rotational moves with other parameters, such as the distance from the centre of rotation to the root particle, may lead, for example, to more frequent large rearrangements within the aggregate.

We now address f_g and f_a for small δ . In fact, it was discussed previously that the average pairwise gradient change is not large enough for the system to develop a significant collective motion. The system thus evolves similarly as in SPMC, into several long-living spherical partitions which coarsen in time, implying that the average size of large moving clusters also increases in time (compare Figs. 3.5(a), 3.5(b), and 3.13(a)). The rotational acceptance density for small δ has unique properties. Figure 3.13(a) shows that f_a develops a local maximum for $1 \ll n_c \ll N/2$ at later stages of the simulation, representing the average size of a large moving cluster. The maximum moves towards large cluster sizes in time, reflecting the coarsening of partitions. The acceptance density f_a decreases with time for clusters with a size below that maximum, and increases for larger clusters. This may again reflect the slow coarsening into spherical partitions with a high internal crystalline structure, and the fact that the movement of a subpartition is less likely than movement of the whole partition. This is because the group of particles which are embedded in such a partition and forming a moving cluster is more likely to cause a high energetic overlap in the new state or to have a boundary with a possible frustrated link in it. We note that f_a for $\delta = 0.05\sigma$ and the largest proposed clusters is larger for rotations than for translations. This is again a consequence of the moving clusters spanning the entire partition, and of the fact that a rotations of a spherical cluster is less likely to change its boundary than a translation. To know how VMMC behaves in the low δ limit may be useful in the approximation of the Brownian dynamics by MC simulation, where small δ is usually required. We can expect that using VMMC in that limit would make sense only for lower temperatures or deeper wells, where clustering is more likely even if δ is smaller than the well width. SPMC is otherwise expected to be sufficient and faster in the approximation of Brownian dynamics than VMMC. Another situation where VMMC would be useful even if δ is small, is a system where the range of the attraction is shorter than δ .

3.5 Undercooled charged attractive systems

We now examine the kinetic evolution of a low-density system with short-range attractive particles which are quenched to the same temperature as in the previous section, with the difference that the particles are charged, and repel each other weakly at long distances (Fig. 3.1). This repulsion is represented by the Yukawa potential in Eq. (3.1). Stable or metastable cluster phases may then form in such systems (Fig. 1.2). Our simulation will not be long enough to reach timescales, where structural observables reach an equilibrium value, and remain constant for the rest of the simulation. In fact, the pace at which the structure evolves will be similar to the pace at which the system without repulsion (Sec. 3.4) crystallizes under the dynamics where collective motion is not allowed (Sec. 3.4.1). This will suggest that the equilibrium state of the system is the fluid-solid coexistence (Sec. 1.5). Since VMMC is expected to achieve kinetically realistic states at long timescales [Whitelam and Geissler, 2007], the repulsion is used here to show, that it drastically slows down the phase separation. We thus talk about the stabilized aggregates rather than about the equilibrium cluster phase.

3.5.1 Stabilization of the aggregates

Figure 3.15 shows that after the quench of charged attractive spheres, the system phase separates rapidly into several large partitions, which move with respect to each other without merging into a single aggregate as is the case for the attractive systems shown in Fig. 3.3. Figure 3.16 then shows that such a simulation does not necessarily result in several separated crystalline domains (Fig. 3.16), can also produce a single large crystalline partition (Fig. 3.16(a)), or partitions with low internal crystal order (Figs. 3.16(b) or 3.16(c)).

From the kinetic point of view, addition of a long-range repulsion to the short-range attraction decreases the strength of the attraction (Fig. 3.1), thereby reducing the bond life time, and leading to more frequent bond breaking and bond making. Figures 3.11(d) and 3.17(d) show that local crystallinity evolves at least twice as fast in the SALR system than in the attractive system without the repulsion, and we seem to be closer to the optimum temperature for self-assembly [Klotsa and Jack, 2011]. From the thermodynamic point of view, addition of a long-range repulsion decreases the liquid-vapour critical point [Pini et al., 2000]. Our simulations made with different values of amplitude A (results of which are not presented here) indicate that the size of isolated droplets for a given quench temperature increases with decreasing attraction strength or with increasing repulsion.




Figure 3.15: Time evolution of the system composed of particles with short-range attraction and long-range repulsion. The simulation is started from a quenched fluid state. The particle colouring is the same as in Fig. 3.3.




Figure 3.16: Different possible products of a simulation with parameters of the simulation in Fig. 3.15, started from a quenched fluid and with the snapshots taken at time $2000 \cdot 10^3$ sweeps.

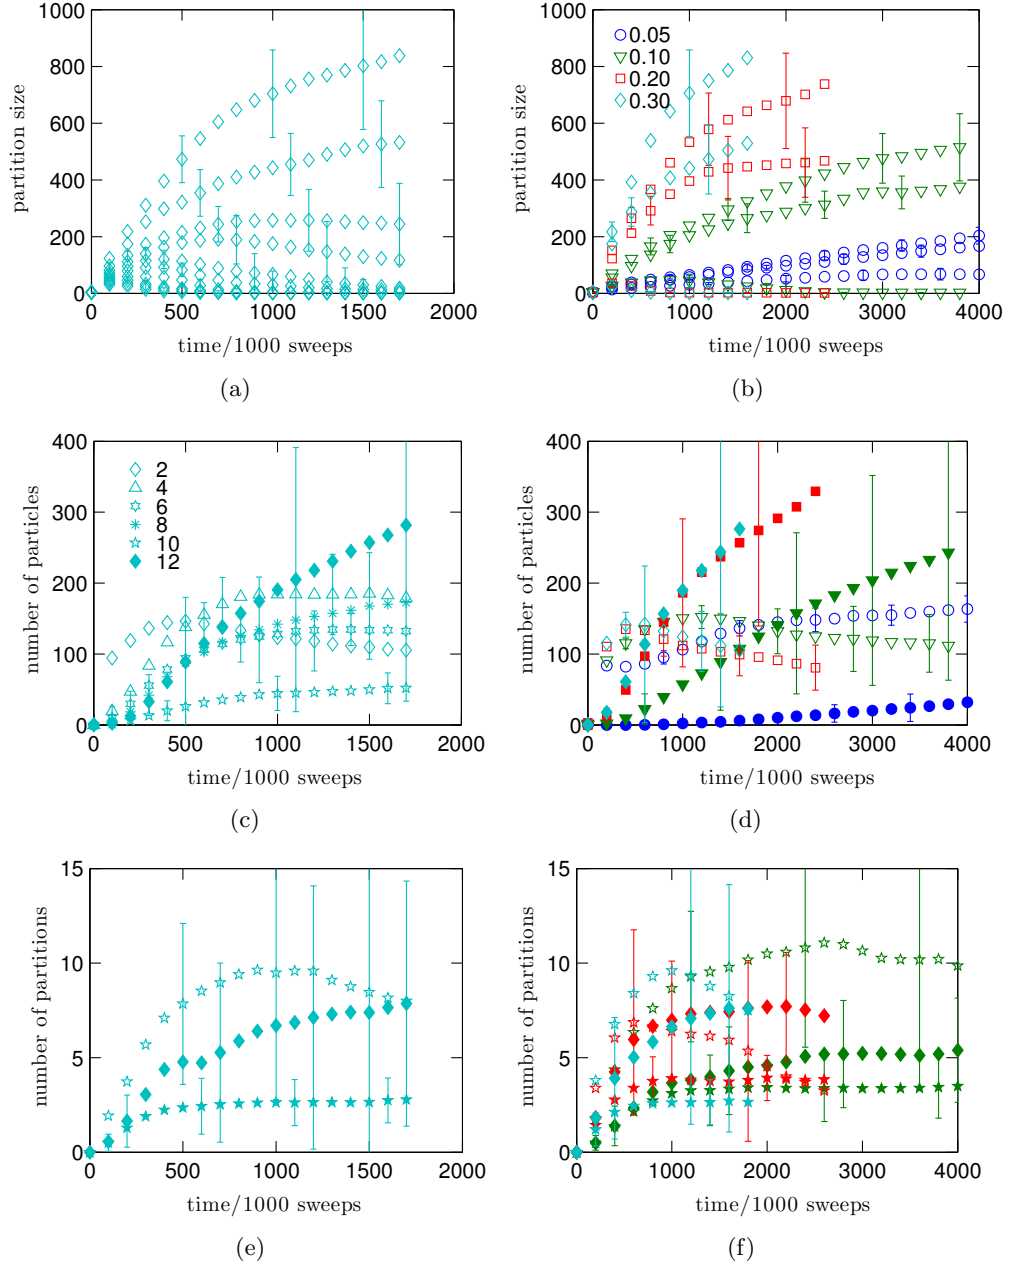


Figure 3.17: Structure evolution in a system of particles with short-range attraction and long-range repulsion, following a quench simulated by VMMC. (a) Size of the first ten largest partitions. (b) Size of the first, second and tenth largest partitions. (c) Particles with a given number of orientational bonds n_b given by the legend. (d) Very weakly bonded particles ($n_b = 2$) are represented by empty symbols and crystalline particles ($n_b = 12$) by full symbols. (e), (f) The number of LSC with no crystals inside (empty stars), LSC with crystals inside (full stars), and crystals (diamonds). Figures (a), (c), and (e) (on left) are results for simulations with $\delta = 0.30\sigma$. Figures (b), (d), and (f) (on right) are results for δ given by the legend of (b).

We observe similarities in structure evolution between the SALR systems under collective motion, and the short-range attractive systems under single-particle motion. First, large isolated partitions do not merge on our simulation timescale. Second, Fig. 3.17 shows first two linear growth regimes with similar properties: Figures 3.17(a) and 3.17(c) show that the growth regimes are again correlated with the evolution of local structure, suggesting that in the SALR system, growth of the largest partitions is governed by single-particle attachment and detachment. Figure 3.17(e) then shows that similarly to Fig. 3.4(e), isolated crystals nucleate mainly during the first growth regime, and also that the SALR system has a similar number of isolated crystalline domains. The main and obvious difference between the SALR system under VMMC motion and the short-range attractive system under SPMC motion is that in the former, partitions are moved and rotated with respect to each other. The SALR system also exhibits larger polydispersity of partitions with smaller partitions being dissolved or merged more rapidly even before the end of the first growth regime.

Figures 3.4 and 3.17 show that addition of the long-range repulsion leads to larger error bars, in particular, in the number of isolated crystals, and in the number of LSCs without a crystal core inside. Figure 3.17(e) shows that, despite these large fluctuations, the peak in the number of LSCs without a crystal inside still captures the start of the second linear stage of growth (Fig. 3.17(a)) also characterized by the equal concentration of particles with low and with high local order (Fig. 3.17(c)). The large error bars may not only be caused by a more complex energy landscape of SALR potentials [Mossa et al., 2004; Archer and Wilding, 2007], but may also be affected by kinetic and finite size effects. It can be seen from Fig. 3.15 that crystals grow mainly through addition of monomers to crystalline surfaces of the partitions. Figure 3.16(c) then shows a situation where the system is trapped in a state where exchange of particles between two amorphous surfaces might last for a long time without any growth of the crystals. This is a consequence of the finite size of our system. If the system were larger, there would be a larger probability of having a crystalline partition, which would attract particles from the amorphous structures.

The existence of trajectories with intermediate states like the one shown in Fig. 3.16(c) is a plausible reason for the large error bars associated with the number of crystalline particles in Fig. 3.17(c). To understand better the large fluctuations in the number of crystalline particles, we firstly point to the fact that the fluctuations occur even during the first growth regime. As discussed earlier for short-range attractive systems, LSCs without a crystal core which survive the first regime are more likely to dissolve at the expense of growing crystals than to recrystallize. If

there is no crystalline partition at the initial stages of the simulation, growth of LSCs without a crystal core inside continues and gives rise to a system containing LSCs only. Figure 3.15 shows that the initial partitions with no crystals in them are spherical even at early times (10^5 sweeps), and we can thus expect that these partitions are also spherical.

Let us present another argument which suggests that long-range repulsion may stabilize the LSC, and make the crystal formation within that partition more difficult. This can be illustrated if we consider an ideally spherical and sufficiently large amorphous partition, composed of particles with sufficiently long-range repulsion. (Repulsion range is equal to at least the diameter of the partition.) Contribution of the total long-range potential of the particles near the surface to the particle in the centre of the partition is large and such a particle is in the energy maximum given by the long-range forces (Fig. 3.18). In other words, energy per particle in the centre of the partition is higher than the energy per particle on the surface, i.e. to add a particle to the centre of the partition is less energetically favourable than to add a particle to the surface of the partition. Contribution of long-range repulsion to the particles in the surface from other particles of the partition may also be seen to be high, but is lower than that in the centre of the partition. It is thus reasonable to expect that particles near the centre of the partition may be destabilized by the long-range interactions to an extent, where it is more difficult to form a crystalline domain. Body density [Sweatman et al., 2014] of the partition may not be constant, and may be locally lower in the core than in the surface, and the core may even become fluid-like. Clearly, if there are no crystals, or more precisely no crystalline surfaces in the system, the partitions may remain in such a LSCs state for a long time, as is observed in Fig. 3.16(c). It is worth noting that the number of particles with $n_b = 8$ in Fig. 3.17(c), corresponding to a crystalline surface, has a much larger error bar than $n_b = 10$, corresponding to local (point) defects. This also confirms that the concentration of particles in amorphous surfaces of crystals, or in amorphous aggregates, varies significantly through the simulations, when compared to short-range attractive systems without long-range repulsion. We also note that the disordered and crystal like partitions observed in Figs. 3.15 and 3.16 may correspond to the transition between low and high density cluster phases as proposed by Sweatman et al. [2014]. It would be interesting to compare the cluster size in our system with the cluster size formula of Sweatman et al. [2014]. This may allow to determine whether the clusters are thermodynamically stable, and characterize that state with respect to the equilibrium phase diagram of Sweatman et al. [2014].

The similarity between the SPMC simulation of the short-range attractive




Figure 3.18: Strong affect of long-range repulsion of particles in the surface on the particles in the centre of the partition. The long-range potential is illustrated by thin circles, and the surface of the partition by the thick circle.

system and the VMMC simulation of the SALR system was the motivation to examine the behaviour of the SALR system under the SPMC motion. The results are presented in Fig. 3.19. The only difference of VMMC compared to SPMC is that VMMC leads to slightly higher error bars in all considered observables. This may be due to an insufficient initial mixing of partitions in SPMC. Also, the number of partitions with a crystal inside is slightly higher in the VMMC than in the SPMC, and the number of partitions with no crystal inside decreases faster in VMMC, suggesting that the dissolution of low structured aggregates via collective motion competes with the single-particle based mechanism.

3.5.2 Scaling of the dynamics

As far as scaling of trajectories with respect to δ is concerned, Figure 3.17(d) shows that equal concentrations of weakly bonded and crystalline particles for different δ is again observed for local observables at different times measured by MC sweeps, and that these times are correlated with different linear growth regimes of Fig. 3.17(b). Moreover, similarly to SPMC of SALR systems, the times where a peak occurs in the number of LSCs without a crystal in Fig. 3.17(d) correspond to the times of

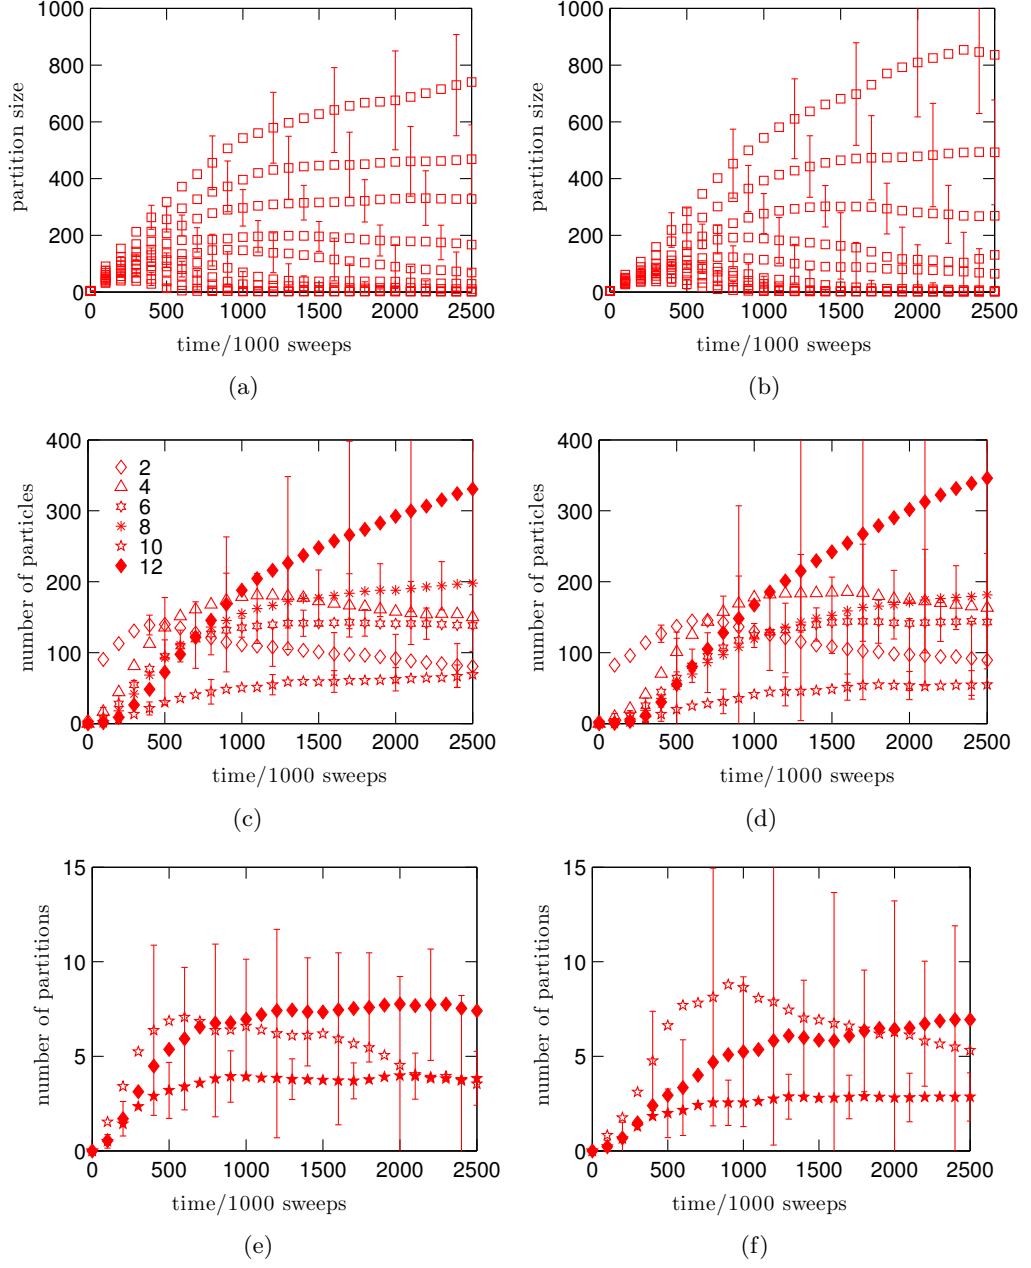


Figure 3.19: Simulations of SALR potentials performed with $\delta = 0.20\sigma$. Results in the left column are simulations from VMMC, and in the right column from SPMC simulations. The meaning of the symbols is analogous to that of previous figures.

equal numbers of crystalline and weakly bonded particles in Fig. 3.17(f) for all δ . These results suggest that the definition of a physical time, and the approximation of the dynamics is more feasible in SALR systems with several isolated long-living aggregates, than in attractive systems without long-range repulsion where a single disordered aggregate forms. Note that error bars increase again with δ , and that our results show only a little difference between the evolution of local structure for $\delta = 0.20\sigma$ and $\delta = 0.30\sigma$.

3.5.3 Generation and acceptance densities

The probabilities of generating and accepting the cluster of a given size with small displacements ($\delta = 0.05\sigma$) for SALR systems are similar in all respects to short-range attractive systems without long-range repulsion, with only a few differences. The generation density f_g is lower (slower) for SALR than for the short-range attractive system at the initial stages of the simulation but higher (faster) at later stages. The initial slowing down may be due to weaker bonds and hence weaker linking, or to the small energetic barrier caused by the long-range repulsive shoulder. The acceleration of f_g and f_a towards larger clusters may be a consequence of the larger size of metastable droplets (partitions). The acceptance density f_a is smaller for SALR than for short-range attractive systems throughout the simulation, however, it extends to larger clusters, especially for rotations. For clusters that do not span the entire partition, this could be caused by long-range interactions between the cluster and the surface of the partition, making the interactions and hence the linking weaker. The reason for large rotations being accepted more often than large translations was discussed previously for attractive systems.

We now present and discuss f_a and f_g for $\delta \geq 0.10\sigma$, i.e. for δ larger or comparable to the width of the attraction. Figure 3.20 only shows the plot for $\delta = 0.30\sigma$. Plots for $\delta = 0.10\sigma$ and $\delta = 0.20\sigma$ have similar qualitative features. It can be seen again that f_g decreases with time for small clusters and increases with time for larger clusters, due to a coarsening process similar to that discussed for attractive systems without long-range repulsion. The density f_g is decreasing with n_c with a drop around $N/2$ at later stages reflecting the fact that the system phase separates into, on average, two large long living (growing) partitions [Figs. 3.17(a) and 3.17(e)], and that the moving cluster spans these partitions. The acceptance density f_a for large translations increases rapidly with time, reflecting an increasing local order, and that $\delta \geq 0.10\sigma$ is needed for the clusters to span the entire partitions. The rotational f_a is very low compared to the translational f_a , as a consequence of the large partitions and of the fact that most rotational moves tend to include particles

in the moving cluster that would move by a distance larger than σ . Optimization of parameters in the rotational move may increase the acceptance ratio.

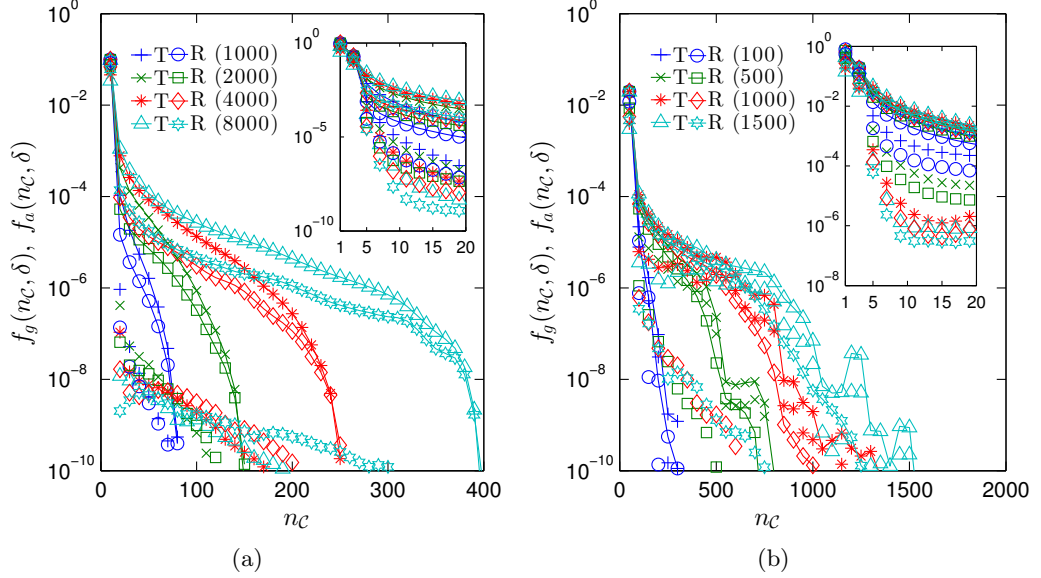


Figure 3.20: Generating and acceptance densities for clusters of size n_c in the system of particles with short-range attractions and long-range repulsions forming long-living isolated aggregates after a quench. Maximum MC displacement δ is equal to 0.05σ and 0.30σ in (a) and in (b), respectively. Meanings of symbols and legends is analogous to that of Fig. 3.13.

The acceptance density f_a for rotations is roughly constant at later stages of the simulations for all $\delta \leq 0.10\sigma$. The maximum size of an accepted rotational move is around six hundred particles for all δ (not shown for 0.10 and 0.20) at the timescales of our simulation. To some extent, this supports the scaling of the dynamics, and the fact that large rotations are constructed in a similar way, i.e. by a displacement which decreases with the distance from the centre of rotation, and the maximum size of this displacement is not larger than σ for all δ .

3.6 Discussion

The primary motivation for this chapter is to find the optimum parameters of VMMC, leading to an efficient and a kinetically realistic simulation, results of which are thermodynamically, or at least temporally, stable isolated crystalline structures. This was not straightforward, because the system under consideration is a non-equilibrium system. Equilibration in such systems is kinetically slowed down by strong attractive or excluded volume interactions. Moreover, there is no free energy

barrier for nucleation and aggregation is mainly driven by diffusion. Since δ is closely related to the diffusion, it turns out that different choices of δ lead to different time evolutions of the simulated system, which may result in different kinetically slowed down structures or different kinetic traps. We have reported results from simulations showing that although the dynamics may be different in some details, there are structural observables that are independent of δ , especially in the long time limits of such simulations. This provides some evidence that the kinetically slowed down structures observed in this thesis have a universal significance and may be physical. An *a priori* choice of δ , which remain constant during the course of the simulation is thus possible, while still predicting realistic products of a kinetically slowed down phase separation.

Optimum δ in simulations of quenches

Having verified that most long-time structural features of the system are independent of δ , we can return to the original objective of this work, and ask what size of δ most efficiently simulates the kinetically slowed down phase separation. However, before that we have to evaluate how to measure the efficiency. While the diffusion coefficient gleaned from the mean squared displacement provides a good insight into the efficiency of VMMC in equilibrium fluids (Sec. 2.9), the diffusion measurements are not adequate in the presence of kinetic slowing down or dynamic heterogeneity which emerges in the undercooled systems under consideration. To deal with such a situation, we fix a computationally achievable value of a quantified structural feature, and we determine the optimum δ such that the value is reached fastest by the simulation. An example of such a value is the average number of crystalline particles, or the average maximum size of the second largest partition.

Whilst in equilibrium fluids, an optimum δ exists for a given system and conditions, and the efficiency of VMMC under this choice remains constant during the simulation, the performance of VMMC in undercooled systems may change in time for a fixed δ as the large aggregates coarsen. Typically δ becomes too large at later stages of the VMMC simulation, becoming ineffective in local equilibration of the system, as most of the time is spent on moving the large partitions. A decrease of δ in such situations is desirable.

Polydispersity and size of δ

Polydispersity of stabilized partitions in Fig. 3.15 may be correlated with the choice of δ in the Monte Carlo simulation or analogously by the choice of timestep in

Molecular Dynamics simulation. Let us assume that the system is rapidly quenched from a fluid phase to a suitable temperature below the spinodal, and that particles in the initial low density region are uniformly distributed across the system, i.e. there are no density inhomogeneities. We can expect that if the simulation of the system starts with a very low δ , the polydispersity of phase separated droplets is lower than those in simulations with larger δ . This could not be verified from observing the largest partitions, because their size does not converge. However, it can be seen from the simulations shown in Fig. 3.17(b), that the onset of the second growth regime is associated with larger partitions for larger δ . We have also observed (details not shown), that the maximum size of the tenth largest partition is larger for smaller δ .

Hence, in fast quenches, very short timescales following the quench seem to affect the polydispersity of phase separated droplets, which may have a larger polydispersity only as a consequence of a larger choice of δ coarsening the dynamics. A low value of δ should thus be chosen when simulating droplet formation after quenches at least at the beginning of the simulation. In fact, if δ is small, phase separation into long-living drops with a low polydispersity may be observed in attractive systems without long-range repulsion.

Hierarchical self-assembly

As described in the Introduction, approximate analytical models of multiscale partition formation have been developed using the mode coupling theory [Cates et al., 2004], and the density functional theory for short-range attractive systems with long-range repulsion [Sweatman et al., 2014].

In what follows we will suggest that a slow variation of δ during the course of the simulation may result in a multiscale modelling protocol, which efficiently captures the collective motion on a scale, which is characteristic of a specific time window of the self-assembly. In a hierarchical self-organization, this scale increases with time, and the size of δ could, for example, be made to depend on the average size of crystals in the system. The early stages of this multiscale simulation protocol would be associated with the shortest δ , the shortest physical timescales, and with the self-assembly of the smallest building units. Later times of the simulation would be characterized by a larger δ , longer times, and self-assembly at a scale given by the size of the aggregates self-assembled earlier. Local motion at later stages would be less frequent, but not completely forbidden. This requires an assumption that local motion is important on one scale but not important later at a larger scale. Such a multiscale simulation requires an assumption that a systematic increase of δ leads to the simulation of very long timescales.

We can illustrate this on our attractive system without long-range repulsion, and speculate about the existence of an efficient simulation protocol which would vary δ during the simulation such that the aggregation happens via the motion of isolated crystalline units. Let us assume that we start the simulation of our rapidly quenched short-range attractive system with a sufficiently low δ ($\delta = 0.05\sigma$). Figures 3.11(b) and 3.11(d) show that the system eventually phase separates into a large number of partitions with a relatively low polydispersity and that the first crystals appear. As implied by our VMMC simulations for small δ , or by the Brownian Dynamics simulations of [Perez et al., 2011], those partitions remain isolated for a long enough time, and do not reach the timescale at which the amorphous partitions merge. It is reasonable to expect that a recrystallization of the partitions may happen on such a timescale.

For large systems and a suitable temperature [Klotsa and Jack, 2013], we may also expect that the droplets will be spherical and will have low polydispersity as discussed above. Such a system then develops a transient but long enough time window, where crystallinity within each of the partitions can develop, and partitions can be temporally described from another scale, as new and larger renormalized spherical particles with attractions that are shorter ranged relatively to the size of the spheres. Moreover, depending on the geometry of the particles and the size of the partitions, the attractive potential of the renormalized spheres may be stronger than that between the constituent particles. This is because a contact between the large spheres can be formed by more than one contact between the constituent spheres, which leads to a deeper well or equivalently to a larger ϵ . The phase diagram of such a renormalized system may change in two different and competing ways. The (relatively) shorter-ranged attraction moves the fluid-solid binodal and the critical point to lower temperatures [Vliegenthart et al., 1999]. The deeper well has the opposite effect, because of the equivalence between the temperature and the energy through $k_B T / \epsilon$. Our observations suggest that the latter effect is stronger, because a single aggregate of partitions is formed rather than a fluid of partitions. It is in the interest of hierarchical self-assembly to avoid the dynamic arrest at the scale of renormalized spheres, and to produce spherical partitions of partitions rather than a single partition from smaller partitions. It is possible that there is a simulation protocol, where T and δ are varied slowly and simultaneously, producing large long-living spherical aggregates from the smaller scale renormalized spheres. This protocol may or may not correspond to physical reality.

Let us stress that to preserve realistic kinetics, δ should never be larger than the size of the smallest building unit in the system, i.e. δ must be lower than σ to

preserve the microscopic dynamics. Nevertheless, one can still obtain a reasonable approximation of the dynamics, even if δ is larger than σ . To do so, all proposed clusters generated by a displacement larger than σ and having less than a certain number of particles need to be rejected. Should the small cluster motion produced by these displacements still be the dominant motion, the waste recycling technique [Frenkel, 2004] might be used to avoid the loss in efficiency by constructing large clusters from those which are not large enough instead of rejecting them *in situ*.

The maximum displacement δ may be varied according to the aims of the simulation. In the simulation of gelation or a rapid initial aggregation into a single low-density domain, we may want to decrease δ in the later stages in order to enhance sampling of the time evolution of local order. On the other hand, if a fluid of crystalline partitions is under consideration, one may want to gradually increase δ to enhance their mobility, and to validate their stability against the aggregation.

Alternative approaches to varying δ or linking strength

Slow variation of δ through the course of the simulation may not be the only way of achieving a realistic dynamics or efficient sampling. Instead of drawing δ from a uniform distribution, δ drawn from another, possibly Gaussian, distribution, may lead to a more realistic distribution of moving clusters approximating realistic dynamics. Properties or parameters of such a distribution may also be slowly varied in time. It is also worth noting that we can draw an analogy between the optimum size of the particle displacement in the VMMC and the optimum pseudotemperature regulating the linking strength in static cluster algorithms [Troisi et al., 2005; Bhattacharyay and Troisi, 2008]. A pseudotemperature in the VMMC algorithm can be introduced simply by using another temperature in the linking function of Eqs. (2.15) or (4.7), and correcting with a bias in the acceptance probability. This approach is equivalent to using fictitious potentials for linking, and real potentials for accepting, the moving clusters [Whitelam, 2011]. One could then consider a slow change of pseudotemperature instead of a slow change of δ .

Rapid periodic changes of δ are possible in single-particle MC simulations [Swendsen, 2011], but can generally bring the system out of equilibrium sampling. If δ is changed slowly in an undercooled (or even equilibrium) system, we expect that the violation of equilibrium sampling caused by these infinitesimal changes is negligible compared to the effects of coarse-graining or non-equilibrium effects of the quench. We note that the changes of δ or the pseudotemperature in the MC simulations during assembly may also be stepwise and follow a real-time feedback. A similar approach has been used by Klotz and Jack [2013], who proposed an efficient

simulation protocol where the temperature of the system is systematically changed to produce high-quality crystals.

Whether, and how, the maximum displacement δ in the VMMC simulation together with its change are represented in experiments is unclear, but time control of bond strength is experimentally achievable by exposure to light in repeated pulses [Jha et al., 2012], and a slow cooling process is used in hierarchical self-assembly of DNA origami [Rothmund, 2006]. The use of microwave heating [Godinho et al., 2008] to locally disturb the bond strength may also be an option.

Origins of long-range repulsion

The addition of long-range repulsion to the short-range attraction leads to thermodynamically stable aggregates, which may form fluid or solid cluster phases. In this thesis, we use similar parameters for the repulsion as in Sciortino et al. [2004], but our long-range cut-off is significantly shorter than that used by Sciortino et al. [2004] or by others Archer and Wilding [2007]; Sweatman et al. [2014]. The range of our repulsion is only comparable to the models of de Candia et al. [2006] and Xia et al. [2011] whose long-range repulsive shoulder is, however, stronger. Although it would be desirable to simulate an ideal model with a long cut-off and Ewald sums, as was done by Toledano et al. [2009], our choice makes the system computationally cheaper to simulate, and allows us to target longer timescales. This may be at the expense of having a less physically realistic model. Here, we assume that the long-range repulsion does not necessarily need to represent the Columb charge, but can also have other contributions such as counterions or salt ions within the solution or polymers accumulated on the surface of the droplets in the form of both steric [Likos et al., 1998] and accumulation repulsion [Louis et al., 2002]. We assume that for many practical situations, and for the purposes of this chapter, the exact form of the repulsion might not be important. To a certain approximation, the qualitative results observed in our model and elsewhere match well. Indeed, we clearly show that addition of even a weak long-range repulsion significantly increases the life-time of the phase-separated droplets, and suppresses their coalescence into a single aggregate [Lu et al., 2006]. We have also observed that for larger A , structures similar to Bernal spirals [Sciortino et al., 2005] are formed even if the range of repulsion is shorter. It is beyond the scope of this section to investigate whether these long-lived partitions are thermodynamically stable as was done by Mossa et al. [2004].

Coarse-graining and long-range repulsion

We have shown that structural features in the SALR system evolve similarly under VMMC and SPMC. This was mainly because single-particle attachment and detachment is the main mechanism for growth and coarsening for the given system and temperature. Since large partitions merge only rarely, expensive simulation of growth via a frequent collective mechanism can be neglected. Nevertheless, the omission of collective motion in an equilibrium or temporary cluster phase may alter the equilibrium partition-partition distance, and affect the average time for a single particle to move from one partition to another. The equilibration of inter-partition distance may be achieved by only a small fraction of collective moves. Indeed, collective moves commonly account for only a fraction of total moves in studies simulating structural arrest or self-assembly [Kraft et al., 2012; Sweatman et al., 2014].

We have even observed common characteristics in structure evolution with attractive systems under SPMC motion, and that avoiding the collective motion in fact helps to mimic the long-range repulsion. This observation opens up a way towards coarsening of both the dynamics and the potentials. These ad-hoc simulations would approximate a dynamical evolution of a SALR as a short-range attractive system under SPMC motion with a small fraction of collective moves taking the long-range repulsion into account to ensure a realistic inter partition-partition distance. We stress that these approximations would only be reliable for a temperature range where single-particle motion is the main mechanism for structure evolution, and where contribution of collective modes of motion can be integrated out. Also, evolution towards an equilibrium cluster phase cannot be modelled by this approach, since the cluster size is expected to be given by the long-range repulsive amplitude [Sweatman et al., 2014].

System size and temperature choice

The choice of the quench temperature T , and the size of the system N , are closely related. They are carefully chosen such that the number of isolated spherical partitions, which emerge after the quench, is large enough, and such that the simulation time is long enough to allow the partitions to crystallize or dissolve. If the temperature were lower than $T = 0.25$, partitions would be smaller, but local order would develop more slowly, and partitions would be less spherical. If, instead of $T = 0.25$, the quench temperature were equal to the ideal self-assembly value $T = 0.38$ [Jack et al., 2007; Whitlam et al., 2009a; Klotz and Jack, 2011; Grant and Jack, 2012], our system of $N = 2000$ particles would phase separate into only a few (1–3) large

partitions, and a study of the interaction between those partitions would be limited or impossible. If the system size were larger, it would not be possible to reach the timescales at which LSCs without a crystal core dissolve.

The size of the system is also related to the choice of δ . Figure 3.11(a) shows that the fraction of the system which ends up in the second largest partition is proportional to δ . This correlation may be a good measure of finding an optimum value of δ for the simulation of kinetic arrest or gelation.

For completeness, we recall the case shown in Fig. 3.16(c) in which the finite size of the system leads to a simulation where only amorphous surfaces of two partitions exist, and where the absence of a crystalline surface leads to a long-lasting single-particle exchange between the surfaces.

Very low density and Brownian dynamics

Perez et al. [2011] performed Brownian dynamics simulations at very low packing fractions $\phi = 0.001$, using the same potentials as in this thesis [Eq. (3.1)] with no long-range repulsion (Sec. 3.4). Their system was quenched to temperatures $T = 0.10\sigma$ and $T = 0.20\sigma$, producing isolated spherical aggregates with well developed crystallinity. The aggregates did not merge into a single partition on the timescale of their long Brownian dynamics simulations. To compare with these results, and to investigate the role of collective Monte Carlo motion for a larger δ , we performed VMMC simulations on the same system. We used $\delta = 0.20\sigma$ and the Brownian scaling of proposed clusters according to Eq. (2.46). The result of the simulation was a system which merged very rapidly into one large fractal-like aggregate, without having enough time to develop spherical partitions with a high internal crystallinity.

Such a huge difference between Brownian dynamics and VMMC for large δ is consistent with the central result of this section, suggesting that for large δ , partitions grow much faster than they should with respect to the local order. We can thus expect that there is a suitable δ lower than 0.20σ , such that it guarantees a trajectory similar to that observed in the Brownian dynamics simulation. Note that $\delta \approx 0.20\sigma$ is commonly used in simulations of self-assembly with the symmetrized VMMC algorithm, and that this choice may thus not lead to realistic trajectories. However, the Brownian dynamics simulation itself can suffer from a deficiency in that it does not sufficiently capture the initial collective motion [appendix of Haxton and Whitlam [2012]]: the VMMC simulation for a large δ may be more physically realistic. We may verify this in the initial stages, but verification of the similarity between Brownian dynamics and VMMC at long times might be difficult, because we expect that the structure of the aggregate at late times will be similar for different

δ and different simulation methods [Foffi et al., 2005a]. A comparison with an experiment may be the best option. Stable aggregates in attractive colloids without long-range repulsion have been observed experimentally by Lu et al. [2006] at higher packing fractions $\phi = 0.04$ than those discussed above. However, a conclusion explaining their stability has not been reached, and polymers accumulated on the surface of the aggregates acting as a long-range repulsion seem to be a plausible reason for this result [Charbonneau and Reichman, 2007a; Whitelam and Geissler, 2008].

3.7 Conclusion

The efficiency of the VMMC can strongly depend on the choice of δ . For a fluid we have shown how different choices of δ affect the size distribution of proposed clusters and their acceptance probability. If δ is large, the simulation generates a lot of waste in the form of a large number of rejected clusters. If δ is small, negligible multiparticle cluster motion happens. The optimum δ is such that the largest clusters have a reasonable chance of being accepted. The optimum choice of δ depends on other parameters such as $k_B T/\epsilon$ or ϕ . We find that the acceptance probability scales exponentially with the size of the chosen cluster. This scaling seems to be independent of the choice of δ at densities close to the critical density and at temperatures not far above the inflexion point of the fluid-solid binodal.

For the low-density quench of attractive particles without a repulsive shoulder, the system condenses into isolated liquid partitions, shortly after the start of the simulation both in SPMC and VMMC. If δ is small, aggregates do not move, and the VMMC simulation is similar to the SPMC simulation with local crystallinity of isolated droplets being developed first. If δ is large, the trajectories produced by VMMC are entirely different, and liquid partitions quickly merge into a single, large, and disordered aggregate, with local crystallinity being developed afterwards.

Interestingly, the crystallization trajectories generated with different δ scale with time (at least approximately), and there are long-time structural features of the aggregates which are independent of the choice of δ . The displacement δ should be ideally chosen such that the growth of local order and growth of the largest aggregates are in harmony. Movement of intermediate sized clusters $1 \ll n_c \ll N/2$ happens only during a short time window, and becomes negligible at later times when the system merges into a single aggregate. This suggests that slow variation of δ during the course of the simulation may enhance the efficiency of the VMMC. Our extensive analysis show that the probability density, of generating or accepting the

cluster, changes rapidly during the simulation, indicating that an approach based on the integration of particles in accepted clusters might be needed to reliably associate a VMMC cycle with a physical time.

In addition to the tests varying δ and optimizing the VMMC simulations, we characterize products of quenches in short-range attractive systems, and analyse the corresponding crystallization pathways. If conditions of the quench are such that the dominant mechanism for structure evolution is the single-particle motion, and contribution from the collective modes of motion can be neglected or integrated out, the system is expected to evolve in three linear growth regimes correlated with local order. The first regime is the formation of aggregates with and without a crystal inside; the second regime is characterized by a dissolution of amorphous structures with a low internal crystal order; and the third regime is a slow growth of crystalline or polycrystalline units via adsorption of single particles from dissolving usually smaller crystals. The optimum conditions for this situation seem to be intermediately strong quenches of systems at low densities to temperatures which correspond to the intersections of the spinodal and the glass line, and which lie within the optimum temperature range for self-assembly gleaned from the dissipation ratios [Klotsa and Jack, 2013]. Linear growth regimes, with nearly identical quantitative properties, were observed in our short-range attractive long-range repulsive system under collective motion, and in the short-range attractive systems under single-particle motion. The results suggest that the expensive simulation of the former system may be approximated by a substantially cheaper simulation of the latter system.

Chapter 4

Collective moves for general potentials

In this chapter, we present a way of selecting the moving clusters through a wider class of virtual states, which is applicable even to purely repulsive interactions. In the context of this thesis, this generalization seems to have its main use in the efficient generation of rotational moves controlling nonclassical crystallization, and aligning the building blocks of a mesocrystal. The inspiration for the generalization comes from the suggestion to extend VMMC to hard spheres in the original paper of Whitlam and Geissler [2007], from the description of the problem by enumerating the number of possible avalanches [Krauth, 2006], and from our formal description of the symmetrized VMMC presented in Chap. 2.

4.1 Introduction

In what follows, it is convenient to use a special notation for the relative positions of particle pairs in their virtual states. This notation will now also include the fact that the VMMC can be applied to anisotropic particles. Let us denote the position and orientation vector of particle i as $x = (r_i, h_i)$, and the image of x under the map M as $Mx = (Mr_i, Mh_i)$. Similarly, for particle $j \neq i$, we will use $y = (r_j, h_j)$, $My = (Mr_j, Mh_j)$. We will consider four virtual states with relative positions of particle pairs denoted as follows. If (x, y) are the positions of particles in (i, j) in state μ , then (Mx, y) are the positions of (i, j) after applying M to i , with the corresponding pairwise energy $\epsilon_{ij}^{(\mu)}$. Similarly, if (x, y) are the positions of (i, j) in μ and M is applied to j , then the positions of (i, j) are (x, My) with energy denoted as $\epsilon_{ij'}^{(\mu)}$. If (x, y) are the positions of (i, j) in ν and M^{-1} is applied to i , then the

positions of (i, j) are $(M^{-1}x, y)$ with energy $\epsilon_{ij}^{(\nu)}$. And if (x, y) are the positions of (i, j) in ν , and M^{-1} is applied to j , then the positions of (i, j) are $(x, M^{-1}y)$, and the energy is $\epsilon_{ij'}^{(\nu)}$.

The particles of the cluster are selected recursively as described in Sec. 2.9. The essence of the generalized scheme is to propose links to all pairs interacting in one of the states

$$(x, y), (Mx, y), (x, My). \quad (4.1)$$

The linking probabilities defined in Sec. 2.3 for the original scheme can then be generalized as follows. Given a move map M , the probability that a link between (i, j) forms is given by

$$p_{ij}^{(\mu)} = 1 - \Phi \left(\Delta\epsilon_{i'j}^{(\mu)}, \Delta\epsilon_{ij'}^{(\mu)} \right), \quad (4.2)$$

where $\Delta\epsilon_{i'j}^{(\mu)}$ and $\Delta\epsilon_{ij'}^{(\mu)}$ are the energy gradients associated with the virtual moves of i and j in state μ , as defined in Eq. (2.8). The function Φ is a suitably chosen real function satisfying

$$0 < \Phi(x, y) \leq 1, \quad (4.3a)$$

$$\Phi(x, y) = \Phi(y, x), \quad (4.3b)$$

$$\Phi(0, 0) = 1, \quad (4.3c)$$

for all real numbers x and y . The probability that a link fails to form is defined in the same way as in Eq. (2.17) as $q_{ij}^{(\mu)} = 1 - p_{ij}^{(\mu)}$. Similarly for the reverse move, the probability of linking a pair (i, j) in state ν under M^{-1} is defined as

$$p_{ij}^{(\nu)} = 1 - \Phi \left(\Delta\epsilon_{i'j}^{(\nu)}, \Delta\epsilon_{ij'}^{(\nu)} \right), \quad (4.4)$$

with the corresponding link failure probability $q_{ij}^{(\nu)} = 1 - p_{ij}^{(\nu)}$. Linking probabilities defined in this way satisfy Eq. (2.3). This guarantees that if $R_{\mathcal{C}}^{(\mu)}$ is a realization of \mathcal{C} in μ under M in a free cluster selection, then \mathcal{C} can be selected via the same realization $R_{\mathcal{C}}^{(\nu)}$ in state ν under M^{-1} .

We will require that the probability of selecting and accepting a given realization in a free cluster selection is the same in states μ and ν . This is given by the

superdetailed balance condition of the form

$$\begin{aligned} & \exp(-\beta E_\mu) p^{(\mu)}(\{b\}) \prod_{(i,j) \in \mathcal{L}^{(\mu)}} p_{ij}^{(\mu)} \prod_{(i,j) \in \mathcal{F}^{(\mu)}} q_{ij}^{(\mu)} \prod_{(i,j) \in \mathcal{B}^{(\mu)}} q_{ij}^{(\mu)} W_{\text{acc}}^{(\mu \rightarrow \nu | R)} \\ &= \exp(-\beta E_\nu) p^{(\nu)}(\{b\}) \prod_{(i,j) \in \mathcal{L}^{(\nu)}} p_{ij}^{(\nu)} \prod_{(i,j) \in \mathcal{F}^{(\nu)}} q_{ij}^{(\nu)} \prod_{(i,j) \in \mathcal{B}^{(\nu)}} q_{ij}^{(\nu)} W_{\text{acc}}^{(\nu \rightarrow \mu | R)}, \quad (4.5) \end{aligned}$$

with terms having a similar meaning as in Eq. (2.22). The main difference between SDB conditions in Eqs. (2.22) and (4.5) is that the latter does distinguish between the realizations with and without a frustrated boundary pair, and takes into account all realizations. This is reflected by the fact that the product running over the boundary pairs considers the full link failure probabilities $q_{ij}^{(\mu)}$ rather than the pre-link failure probabilities $q_{i'j}^{(\mu)}$.

Similarly to Sec. 2.4, the SDB condition (4.5) can be simplified and the acceptance probability of moving particles \mathcal{C} by M can be expressed as

$$W_{\text{acc}}^{(\mu \rightarrow \nu | R)} = \min \left\{ 1, e^{-\beta(E_\nu - E_\mu)} \frac{\prod_{(i,j) \in \mathcal{B}^{(\nu)}} q_{ij}^{(\nu)}}{\prod_{(i,j) \in \mathcal{B}^{(\mu)}} q_{ij}^{(\mu)}} \right\}. \quad (4.6)$$

We assumed that the cluster was selected via the free cluster selection, i.e. via proposing links to all pairs satisfying conditions in Eq. (2.1). Similarly to what has been shown in Secs. 2.4 and 2.5, we can restrict the class of pairs to which a link is proposed, and show that the cluster acceptance probability (4.6) can still be used. The restriction of this selection is that links are proposed only to pairs interacting in one of the states expressed in Eq. (4.1). Table 4.1 then helps to identify those pairs to which a link is proposed in both states μ and ν , and to those pairs to which a link is proposed in only one of the states μ or ν . It shows that if (i, j) is a particle pair from the cluster in state μ , then the link can also be proposed to that pair in state ν . However, if (i, j) is a pair from the boundary of \mathcal{C} in one of the states μ or ν it is generally not guaranteed that (i, j) is a boundary in the other state. In analogy to the previous sections, we say that the realization has a symmetric core and an asymmetric boundary. Let us discuss this asymmetry closer, and let us distinguish between the following four different cases.

First, if (i, j) is a boundary pair in μ interacting in (x, y) or in (Mx, y) , then (i, j) is a boundary pair in ν interacting in (Mx, y) or in $(M^{-1}Mx, y)$, and a link is proposed to that pair in state ν .

Original state:	
(x, y)	
(Mx, y)	
(x, My)	
<hr/>	
Both i and j moves under M :	Only i moves under M :
(Mx, My)	(Mx, y)
$(M^{-1}Mx, My)$	$(M^{-1}Mx, y)$
$(Mx, M^{-1}My)$	$(Mx, M^{-1}y)$

Table 4.1: Given a pair of particles (i, j) in state μ to which a link is to be proposed, the table lists all their relative positions which can be considered during the selection of \mathcal{C} both in state μ under M and state ν under M^{-1} . The top three positions correspond to the original and two virtual states of (i, j) in state μ . The bottom three positions on the left are the original and two virtual states of (i, j) in state ν under the inverse map M^{-1} and under the assumption that both particles i and j are part of the cluster in state μ under M . The bottom three positions on the right are the original and two virtual states of (i, j) in state ν under M^{-1} , and under the assumption that (i, j) is a boundary pair of \mathcal{C} in state μ . The table aims to show that each relative position of (i, j) is equivalent to the relative position of (i, j) in one of the top three states, with the exception of the position $(Mx, M^{-1}y)$. These equivalences can be seen if we realize that the translational and rotational Monte Carlo move map M is an isometric operation, and that $M^{-1}M = I$ is an identical map.

Second, let (i, j) be a boundary pair in μ not interacting both in (x, y) and in (Mx, y) , but interacting in (x, My) . Then (i, j) is a boundary pair in ν not interacting both in (Mx, y) and in $(M^{-1}Mx, y)$, but is not necessarily interacting in $(Mx, M^{-1}y)$. The link is thus not necessarily proposed to that pair in state ν . We can discuss these two cases. If such pair is interacting in $(Mx, M^{-1}y)$, then a link is clearly proposed to that pair in ν . If such pair is not interacting in $(Mx, M^{-1}y)$, then a link, proposed to such a pair, is guaranteed to fail by the definition of the linking function in Eq. (4.4), and we can assume that a link is proposed to that pair although it is not.

Third, if (i, j) is a boundary pair in μ not interacting in any of the states (x, y) , (Mx, y) , (x, My) but interacting in $(Mx, M^{-1}y)$, then the link is not proposed to such a pair in state μ , but is proposed to that pair in state ν . The condition in Eq. (4.3c) again implies that such a pair fails to form in state μ , and we can again assume that a link is proposed to that pair is state μ , although it is not. This class of pairs needs to be detected and included in the bias of the acceptance probability.

Fourth, if (i, j) is a boundary pair in μ not interacting in none of the states (x, y) , (Mx, y) , (x, My) , and $(Mx, M^{-1}y)$, then the link has no chance to form both in states μ and ν , and we can again assume that the link is proposed to that pair although it is not.

The above considerations show that, if the pair selection is restricted only to pairs interacting in one of the states in Eq. (4.1), the set of pairs in the boundary of the cluster to which a link is proposed during the cluster selection is not the same in states μ and ν . However, it is also shown that the asymmetry does not violate the explicit formulation of the SDB condition, taking pair selection probabilities into account (see Chap. 2), and that the free cluster selection can be assumed, although the restricted selection is used in the algorithm.

We point out that all boundary pairs not interacting in any of the states (x, y) , (Mx, y) , (x, My) but interacting in $(Mx, M^{-1}y)$ must be detected after the cluster selection, and included in the boundary $\mathcal{B}^{(\nu)}$ to evaluate the denominator in the bias of the acceptance probability in Eq. (4.6).

4.1.1 Linking probabilities

We now consider two specific choices of linking function Φ . Firstly, we consider the case

$$\Phi \left(\Delta\epsilon_{i'j}^{(\mu)}, \Delta\epsilon_{ij'}^{(\mu)} \right) = \min \left\{ 1, \exp \left(-\beta \min \left\{ \Delta\epsilon_{i'j}^{(\mu)}, \Delta\epsilon_{ij'}^{(\mu)} \right\} \right) \right\}, \quad (4.7)$$

which as shown later is equivalent to the linking in the original symmetrized scheme [Whitelam et al., 2009a; Ružička and Allen, 2014]. The form (4.7) is such that if (i, j) is non-interacting in (x, y) and (Mx, y) , then $q_{ij}^{(\mu)} = q_{ij}^{(\nu)} = 1$. Thus when selecting the particles to \mathcal{C} , it is enough to consider only pairs (i, j) interacting in (x, y) or in (Mx, y) , and omit pairs interacting in (x, My) or in $(Mx, M^{-1}y)$. This speeds up the simulation, because fewer nearest neighbours need to be taken into account when using cell lists for the nearest neighbour search. A drawback of the choice (4.7) is that no clusters are formed in systems with purely repulsive potentials.

The second possible choice of Φ discussed here is

$$\Phi \left(\Delta\epsilon_{i'j}^{(\mu)}, \Delta\epsilon_{ij'}^{(\mu)} \right) = \min \left\{ 1, \exp \left(-\beta \max \left\{ \Delta\epsilon_{i'j}^{(\mu)}, \Delta\epsilon_{ij'}^{(\mu)} \right\} \right) \right\}, \quad (4.8)$$

which as opposed to (4.7) forms clusters not only in systems with attractive but also in systems with purely repulsive potentials. In the particular case of hard spheres or disks, any pair (i, j) , which interacts in one of the states (Mx, y) or (x, My) is linked.

4.1.2 Early rejection scheme

Let us now show that the probability defined in Eq. (2.16) is equal to that defined in Eq. (4.7) for the early rejection scheme. Similarly to Eq. (2.21), we express the linking probability as

$$p_{ij}^{(\mu)} = \min \left\{ p_{i'j}^{(\mu)}, p_{ij'}^{(\mu)} \right\}.$$

By using definitions (2.15a), (2.15b) and relation

$$\begin{aligned} \min \left\{ \max \{0, 1 - e^x\}, \max \{0, 1 - e^y\} \right\} &= \min \left\{ 1 - e^{\min\{0, x\}}, 1 - e^{\min\{0, y\}} \right\} \\ &= 1 - \max \left\{ e^{\min\{0, x\}}, e^{\min\{0, y\}} \right\} = 1 - e^{\max\{\min\{0, x\}, \min\{0, y\}\}} \\ &= 1 - e^{\min\{0, \max\{x, y\}\}} = 1 - \min \left\{ 1, e^{\max\{x, y\}} \right\} = 1 - \min \left\{ 1, e^{-\min\{-x, -y\}} \right\}, \end{aligned}$$

valid for all real numbers x, y , this leads to

$$p_{ij}^{(\mu)} = 1 - \min \left\{ 1, \exp \left(-\beta \min \left\{ \Delta\epsilon_{i'j}^{(\mu)}, \Delta\epsilon_{ij'}^{(\mu)} \right\} \right) \right\},$$

which is identical to $p_{ij}^{(\mu)}$ defined by (4.2) and (4.7).

The early rejection scheme can also be formulated for hard particles. In such an algorithm, the links are proposed to all pairs interacting in one of the states (x, y) , (Mx, y) , (x, My) , or $(Mx, M^{-1}y)$. With the aid of Eq. (4.8), the rule for

linking can be formulated as follows. Form a link if (i, j) overlaps (is interacting) in one of the states (Mx, y) or (x, My) . If (i, j) does not overlap in (Mx, y) , (x, My) and in $(Mx, M^{-1}y)$, label it as outright failed. If (i, j) does not overlap in (Mx, y) , (x, My) but overlaps in $(Mx, M^{-1}y)$, label it as frustrated. Accept the cluster, if boundary \mathcal{B} is only composed of outright failed links. Reject the cluster, if there is a frustrated link in \mathcal{B} .

4.2 Results

4.2.1 Equilibrium sampling

Similarly to Sec. 2.11, the validity of the algorithm was tested by comparing the radial distribution functions obtained from VMMC and SPMC simulations.

Firstly, we test the algorithm on a 3D hard sphere fluid in the NpT -ensemble with pressure $p = 10$ and size $N = 1152$. Both SPMC and VMMC simulations had maximum displacement $\delta = 0.03$, length 10^7 sweeps, and started from an equilibrated state. The results are shown in Fig. 4.2(a). Similarly to what has been discussed in Sec. 2.12, the dominant motion in this VMMC simulation is still the single particle motion, and the match between the functions is a necessary requirement, but does not provide a strong evidence that the algorithm is correct.

An additional test has thus been performed on a system of particles with SALR interactions given by Eq. (3.1), with special parameters $A = 0.60$, $\xi = 2.0$, $\alpha = 18$, $\sigma = 1.0$, $\epsilon = 2.0$, with the potential being truncated and shifted at the cut-off distance $r_c = 3.0$. The system was simulated in the NVT -ensemble with size $N = 2000$, packing fraction $\phi = 0.01$, and temperature $T = 0.25$. The equilibrium system at those conditions is characterized by metastable partitions of around 13 particles, and the dominant motion in the VMMC simulation is the motion of those partitions. The metastable partitions can exchange particles and constantly rearrange, dissolve, and nucleate at the timescale of our simulations. This state thus provides a perfect state to compare VMMC with SPMC. We perform SPMC and VMMC simulations with length 10^7 sweeps, maximum displacement $\delta = 0.20$, and with various initial equilibrated states. The results are shown in Fig. 4.2(b), with an exact match observed again. The VMMC simulation shown in Fig. 4.2(b) uses the linking function of Eq. (4.8), a similar result (not shown) has also been obtained for linking function of Eq. (4.7). This provides a stronger evidence of the validity of the VMMC algorithm than the one provided in Sec. 2.11 for an equilibrium short-range attractive fluid.

4.2.2 Enhancement of acceptance probability

We now show that the simulation of repulsive interactions plays a key role in the generation of rotational VMMC moves. We show this in a simulation of the kinetically slowed down phase separation in the system with short-range attraction and long-range repulsion defined in Eq. (3.1) after a quench to the temperature $T = 0.25$. Fig. 4.3 shows the time evolution of generating and acceptance densities in the VMMC simulations taking into account both linking options of Sec. 4.1.1. We took $\delta = 0.20\sigma$ in both cases, and used the Brownian scaling of clusters defined by Eq. (2.46). It can be seen that linking the particles through the function in Eq. (4.8) increases the proportion of accepted moves both for rotations and translations, and that the fraction of accepted rotational moves is about two orders of magnitude higher. Such an enhancement of acceptance density is an example, where the generalized VMMC performs much better than the original VMMC algorithm.

4.2.3 Scaling of the dynamics

Figure 4.4 shows the time evolution of structure in the short-range attractive system following a quench. Different sizes of maximum displacement δ are used with the linking function defined in Eq. (4.8). Quantitatively, structural properties observed at a certain time in the simulation with the linking function of Eq. (4.7) and fixed maximum displacement δ are nearly identical to the structural properties observed at the same time with the linking function of Eq. (4.8) and maximum displacement $\delta/2$. A plausible reason for this is the fact that links are proposed not only to pairs interacting in (Mx, y) but also to pairs interacting in (x, My) , thus forming larger clusters for a given δ . The exact reason for this scaling was, however, not systematically investigated.

Figures 3.11(d) and 4.4(d) show that by using the linking function of Eq. (4.7) it is possible to achieve a larger local crystallinity than with function of Eq. (4.8) in a given CPU time. The CPU time is not shown, but the time at which the simulation stopped is similar for different δ , and is indicated by the end of the available data in the respective plots. Even though the simulations with potentials in Eq. (4.8) are generally twice as long as the one using the function in Eq. (4.7), the cluster acceptance probability is significantly larger for the potential in Eq. (4.8), especially for rotational moves.

Figure 4.4 also shows that for $\delta = 0.1\sigma$ and $\delta = 0.2\sigma$ there is almost no difference in the evolution of the local structure, indicating that the use of $\delta \leq 0.1\sigma$ is suitable for the approximation of the dynamical evolution.

Similarly to Fig. 3.11(f), it is confirmed again in Fig. 4.4(f) that the microscopic structure of the aggregates at long times is independent of the microscopic dynamics. Structure evolution at long times in simulations with $\delta \leq 0.025\sigma$ needs to be investigated. Nevertheless, Fig. 4.4(b) indicates that the isolated partitions move and coalesce much faster than those in Fig. 3.11(b) linked via Eq. (4.7) and $\delta = 0.05\sigma$. This points to the fact that $\delta = 0.025\sigma$ may be an effective value for simulations where crystals arrest in a gel of clusters or form a Wigner glass or a mesocrystal.

4.3 Discussion

The original aim of the VMMC simulation was to approximate the dynamics in short-range attractive systems with overdamped dynamics. Particles were selected to the moving cluster according to the local energy gradients; however, the links created with Eq. (4.7) have a chance to form only for the specific energetic situation, such that $\Delta\epsilon_{ij}^{(\mu)} > 0$ and $\Delta\epsilon_{i'j'}^{(\mu)} > 0$. In other situations, such as $\Delta\epsilon_{i'j'}^{(\mu)} > 0$ and $\Delta\epsilon_{ij}^{(\mu)} = 0$, the links do not form, although they should form. The new way of linking the particles defined in this chapter by Eq. (4.8) can link particles in such situations, but also links particles which would have not been linked in the classical mechanical picture. An example of such unrealistic linking is a particle pair with $\Delta\epsilon_{ij}^{(\mu)} = 0$ and $\Delta\epsilon_{i'j'}^{(\mu)} > 0$. The linking defined in Eq. (4.8) generates a different collection of moving clusters, which are kinetically realistic, but may be generated according to unphysical local energy gradients, or gradients which are unlikely. Clusters selected by linking in Eq. (4.7) are closer to realistic gradients, but capture only a certain class of possible physical collective motion. The linking realized via Eq. (4.8) generates another class of moving clusters, which preserves kinetics, but may be selected according to an unlikely local energetics. It can be expected that other functions Φ satisfying conditions (4.3), lead to different classes of kinetically realistic moving clusters, which may or may not be selected according to a realistic local energetics. It is possible to have a set of functions satisfying conditions (4.3), and use them randomly according to a certain distribution in order to select and accept moving clusters with a rich range of topologies. A simple example of such a multifunctional VMMC, is to take the function in Eq. (4.7) with probability 1/2, and the function in Eq. (4.8) with probability 1/2. Such a simulation clearly still samples from the Boltzmann distribution, and moving clusters have more diverse properties.

Considering different functions Φ satisfying conditions (4.3) may not be the only way of achieving a realistic dynamics. In fact, condition (4.3b) significantly

reduces the generation of realistic dynamical clusters, and even makes it impossible, because a link formed under the gradient $\Delta\epsilon_{ij}^{(\mu)}$ must simply form with a different probability than a link under the gradient $\Delta\epsilon_{ij'}^{(\mu)} > 0$. Although the condition (4.3c) greatly simplifies the acceptance probability, because it guarantees the equality between the link formation probabilities in Eq. (2.14), it seems better to return to the original formulation of the VMMC [Whitelam and Geissler, 2007], for which $p_{ij}^{(\mu)}$ is not necessarily equal to $p_{ij}^{(\nu)}$ for pairs in the cluster. In such a scheme, the linking probabilities would have to be chosen carefully, and links would be proposed not only to pairs interacting in state μ , but to any pairs interacting in one of the states in (4.1).

One may argue that proposing links to pairs interacting in (x, My) may not be needed in systems such as purely repulsive particles. However, there are still physical reasons to consider these interactions. If the system is suspended in a liquid, this interaction may result from hydrodynamic forces. In an ideal hard sphere system, the force may be thought of as an entropic force. Clearly, a particle pair interacting in (x, My) should not be linked with probability one, but with a lower probability, which is where the omission of condition (4.3c) is necessary. Related to this, it is worth pointing out that proposing links to pairs interacting in (x, My) offers a workaround to the problem with enumerating the number of possible avalanches, which are simply possible ways of selecting the cluster via a movement of a single by map M [Krauth, 2006]. The number of these avalanches is generally not the same under the forward and under the reverse move, which is why proposing links to pairs interacting in (Mx, y) and ignoring those in (x, My) is not sufficient, and violates detailed balance as also discussed for hard disks by Krauth [2006].

Finally, we briefly comment on technical aspects of the VMMC. First, the superdetailed balance condition should allow one to distinguish between straight and reflected moves [Bernard et al., 2009] even when selecting clusters in the VMMC. If reflection is used during the linking, the displacement M does not remain constant during the recursive selection of the cluster, and changes according to the tangents of the overlaps. The VMMC move might then become even more versatile and realistic. Second, a major disadvantage of the single cluster MC algorithms is that parallelization is not straightforward, because the size of the moving cluster is not known at the outset, and may span the entire system. A possible way of parallelizing the VMMC may follow the approach by Kaupuzs et al. [2010], designed for the Wolff cluster algorithm for lattice spin systems. Their parallelization does not separate the lattice into domains, but different threads are assigned to different pairs in the growing queue defined in Sec. 2.9. Third, rotations and translations are isometric

operations, and the isometric property is essential for simplifying the acceptance probability. Since the reflection operation is also an isometry, there is a good reason to expect a similar algorithm which incorporates the reflection operation for anisotropic particles [Sinkovits et al., 2012].

4.4 Conclusion

This chapter presented and analysed a complete set of virtual states which need to be considered to select the cluster via a general function linking not only attractive but also repulsive potentials. It has been shown that the simulation of repulsion is particularly useful in the generation of rotational moves. The acceptance probability of rotations increases by up to two orders of magnitude under the newly proposed linking. The high acceptance density is not only more efficient, but is desirable in the approximation of the dynamics by the MC scheme. The evolution of microscopic and macroscopic structural properties of the system is nearly identical in the new scheme as in the old scheme, but is about twice as fast when measured by the number of MC sweeps.

Combination of a wider class of linking functions is expected to provide a richer spectrum of moving cluster topologies thus enhancing the relaxation. We confirmed by yet another dynamics that the long-time structural properties of the aggregates are independent of the microscopic details of the simulation, and are also independent of the choice of the maximum displacement. We can thus expect that VMMC with the new way of linking yields physically realistic kinetically slowed down aggregates.

We conclude by mentioning the work of Hedges and Whitlam [2011], which suggests that access to the rotational degrees of freedom seems to affect the formation of amorphous drops prior to crystallization and hence to influence whether the crystallization path is classical or non-classical. Precise control over the collective rotations is thus likely to play an important role in the simulation of hierarchical self-assembly, where a gas and a liquid of clusters can coexist and form high density modulated phases such as cluster crystals or potentially mesocrystals [Sweatman et al., 2014].

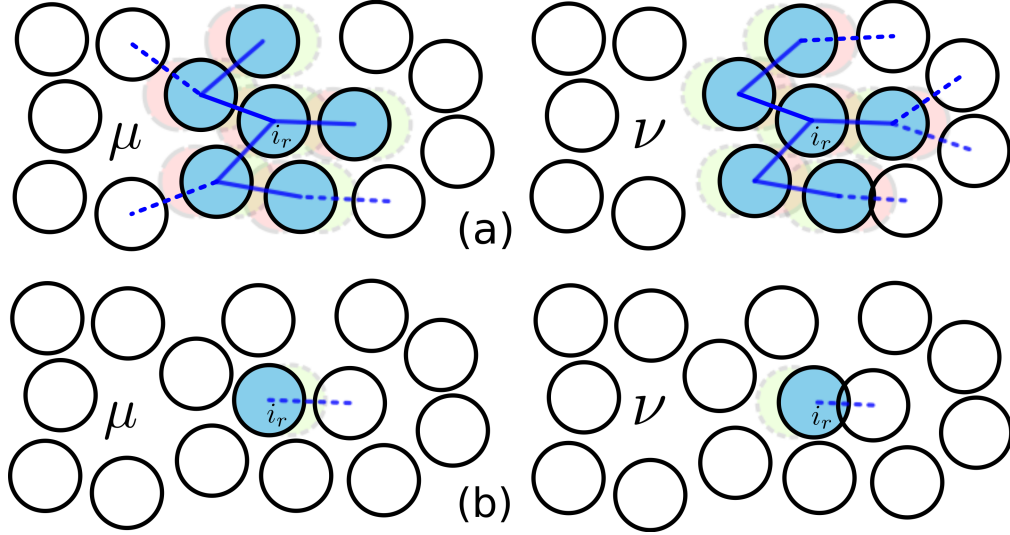


Figure 4.1: Illustration of the translational virtual move Monte Carlo on a system \mathcal{S} composed of soft repulsive disks where the range of interaction r_c is determined by the radius of the disk. State ν (right) is created from μ (bottom) by applying a translational move map M to the particles in \mathcal{C} (blue disks with full margin). Particles in virtual states corresponding to M are green (with short-dashed margin), and particles in virtual states corresponding to M^{-1} are red (with long-dashed margin). Full lines between particles represent formed links \mathcal{L} . Dashed lines are failed links defining the boundary of \mathcal{C} . There are no failed links internal to the cluster \mathcal{F} in this specific picture. As described in the main text, the cluster is selected by applying the map M to the root particle i_r , and by iterative linking of particles to the root. (a) The new way of selecting the moving cluster attempts to include all particles which interact in one of the states (x, y) , (Mx, y) , (x, My) . Linking function in Eq. (4.8) can be used in this case, and allows to create clusters even in a systems interacting only via a repulsive interaction. (b) Version of the algorithm described in Sec. 2.5 is used, where links are only proposed to pairs interacting in one of the states (x, y) or (x, My) , and function in Eq. (4.8) instead of function in Eq. (4.7) is used to link the particles to the cluster. In this particular case of repulsive interaction and no attraction, links cannot be formed and the maximum size of the moving cluster is a single particle.

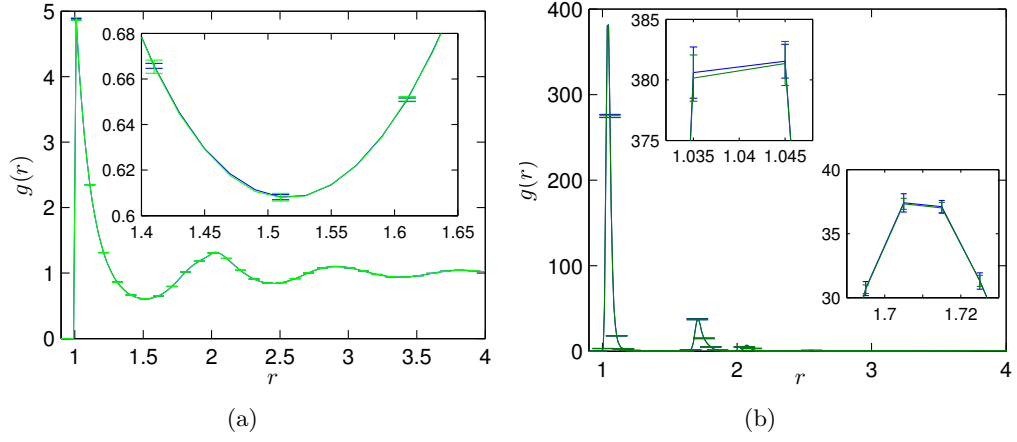


Figure 4.2: Radial distribution functions obtained from SPMC (blue) and VMMC (green) simulations. (a) Hard sphere fluid. The averages are over ten independent simulations. Inset: Magnification of the first minimum. (b) Equilibrium cluster fluid, with averages over six independent simulations. Insets: Magnifications of the first and second maximum. Averages are over ten independent simulations.

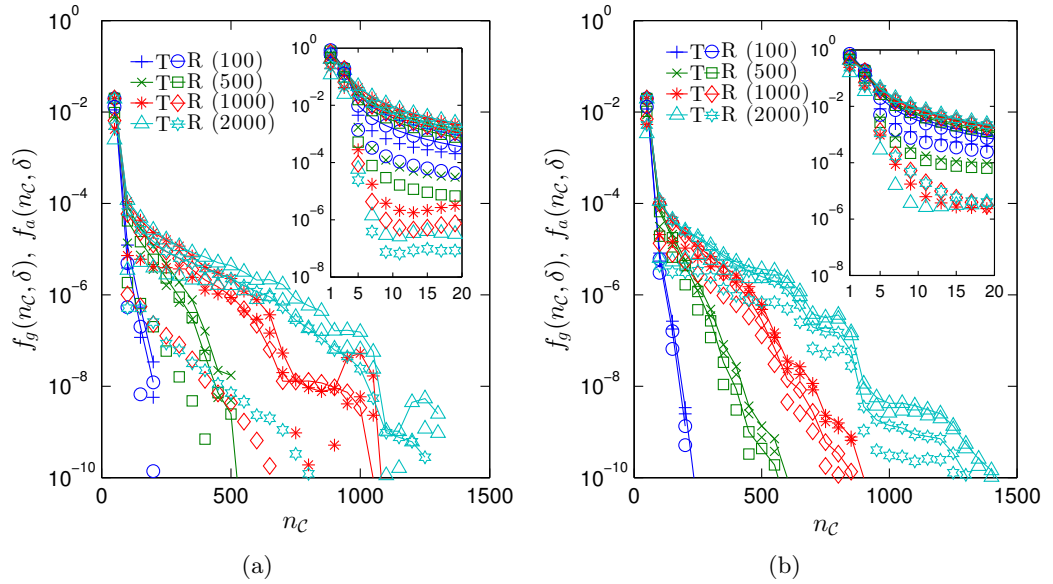


Figure 4.3: Acceptance density of rotations is up to two orders of magnitude higher when the alternative linking scheme is used to select the cluster. Meaning of the symbols is analogous to the previous figures.

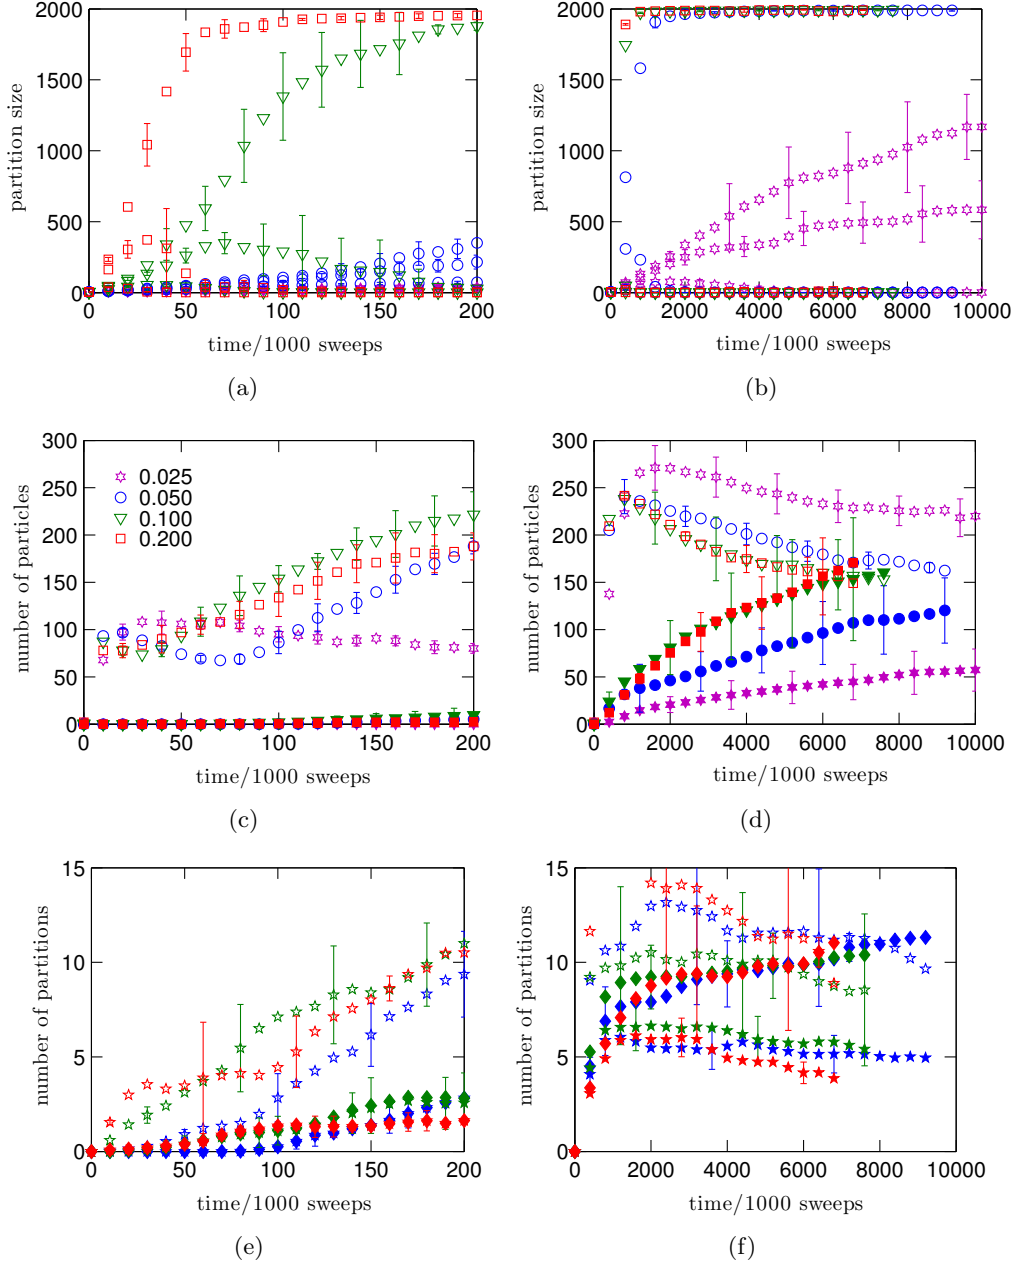


Figure 4.4: Dependence of structural evolution on the choice of maximum displacement in the VMMC simulation. Figures (a), (b), (c), and (d) have meaning analogous to that of the SPMC in Fig. 3.5. Figures (e) and (f) show the number of LSCs with no crystal inside (empty stars), the number of LSCs with a crystal inside (full stars), and the number of crystals (diamonds). Different colours correspond to different δ given in the legend of (c). All error bars are averages from ten independent simulations. Simulations with $\delta > 0.05\sigma$ were too slow, and did not reach the end of the time axis.

Chapter 5

Renormalized long-range repulsion

In this chapter, we study a short-range attractive and long-range repulsive system, where the repulsion depends on the size of the partitions, and on the position of particles with respect to the partitions accommodating the particle pair. In this type of interaction, particles within the same partition are related via a potential, which is different from the potential between particles from different partitions. We can no longer talk about pairwise potentials, but the term embedded potentials seems to be more appropriate. The renormalization is based on the one used by Xia et al. [2011] to study the self-assembly of polydisperse spherical nanoparticles of cadmium selenide to less polydisperse spherical clusters. The reduction of polydispersity was in fact the main result of their paper, providing a method for creating uniformly sized building blocks for hierarchical self-assembly. Another type of renormalization for very short-range attractive and long-range repulsive colloids was also proposed by Sciortino et al. [2004], with the difference that the clusters were treated as monodisperse spherical particles with a stronger repulsion.

5.1 Simulation model

In what follows, we define a model which aims to bridge the gap between the advanced model of Xia et al. [2011] and the simple model of Sciortino et al. [2004] with well established properties. An analogy between these two types of model has already been pointed out by Clancy [2011]. We use an attraction, which is shorter ranged than the one in Xia et al. [2011], but we keep the embedded and volume dependent renormalization. As opposed to Xia et al. [2011] we use a monodisperse

system only, and we compare the polydispersity of the aggregate with the polydispersity of the results using the same but non-normalized potentials of Sec. 3.2.1. We start again from the generalized Lennard-Jones potential equipped with Yukawa repulsion

$$V(r) = A \frac{e^{-r/\xi}}{-r/\xi} + 4\epsilon \left[\left(\frac{\sigma}{r} \right)^{2\alpha} - \left(\frac{\sigma}{r} \right)^\alpha \right], \quad (5.1)$$

which is cut off and shifted at a distance $r_c = 3.0$. To renormalize A we define the rate of change of repulsion with volume as

$$s \approx \frac{A_f - A_c}{R_f^3 - R_c^3}, \quad (5.2)$$

where quantities A_f , and A_c are defined and determined in *a priori* SPMC simulations with the potential in Eq. (5.1). These simulations start from a long-living state of several clusters as shown in Fig. 3.15 at late times. We perform several such simulations with different but fixed A . The smallest A such that the system dissolves on a reasonably short time-scale (10^5 VMMC sweeps) is defined as the repulsion amplitude of the fluid and denoted as A_f . The largest A such that the system does not dissolve on longer timescales (10^7 MC sweeps) is defined as A_c (repulsion amplitude of clusters). The parameter R_c is the average radius of gyration of the partition, and is fixed arbitrarily. We fix it as $R_c = \sigma$. The parameter R_f represents the radius of gyration of the non-normalized unit. In our case it is a single particle. We choose $R_f = \sigma/\sqrt{3}$. The parameters A_f and A_c are thus determined in separate simulations using the potential in Eq. (5.1) for different values of A (non-normalized), and observing for which A partitions grow, and for which A they dissolve.

To renormalize the repulsion A in Eq. (5.1) the system needs to be split into partitions. Similarly to Sec. 3.2.2, by a partition we understand a set of neighbouring particles \mathcal{P}_l , with neighbours being defined as two particles with distance less than 1.4σ . The system is then divided into n_p partitions, symbolically

$$\mathcal{S} = \bigcup_{l=1}^{n_p} \mathcal{P}_l = \bigcup_{i=1}^N \mathcal{P}_{k_i}, \quad (5.3)$$

where k_i is the number of the partition with particle i , and \mathcal{P}_{k_i} is the partition with particle i .

The idea behind the following phenomenological renormalization of the repulsion is that A in the pairwise potential (5.1) is chosen depending on the relative positions of particles i and j , with respect to their partitions \mathcal{P}_{k_i} and \mathcal{P}_{k_j} . Let R_{k_i}

and R_{k_j} denote the radii of gyration of \mathcal{P}_{k_i} and \mathcal{P}_{k_j} . Having obtained A_f and A_c previously, we can calculate the rate of amplitude change in Eq. (5.2), and define the renormalized amplitude as

$$A = \begin{cases} A_c & \text{if } k_i = k_j, \\ A_c + s(R_{k_i}^3 + R_{k_j}^3) & \text{if } k_i \neq k_j. \end{cases} \quad (5.4)$$

The parameters of the model were chosen as close as possible to those of Sec. 3.2.1 in order to compare the effects of renormalization with the previous results. We fixed the screening length as $\xi = 2.0\sigma$ similarly as in Toledano et al. [2009]; Sciortino et al. [2004]. The range of attraction 0.10σ is given by $\alpha = 18$ in analogy with Vliegthart et al. [1999]. The depth of the well is taken as $\epsilon = 1.0$, and the potential is cut-off at a distance equal to $r_c = 3.0\sigma$. Our *a priori* simulations for the given choice of parameters showed that $A_f \approx 0.25$ and $A_c \approx 0.08$. We took $R_c = \sigma$ for simplicity, and because we want the clusters to be sufficiently small. The potential was cut off and a shift eliminating the discontinuity at r_c was applied.

To simulate the evolution of the system after the quench we used the generalized VMMC algorithm with a maximum displacement $\delta = 0.20\sigma$. Without testing, we assume that the dynamics scales with respect to δ , and that the resulting structure is independent of the choice of δ . The VMMC algorithm is only defined for pairwise potentials, and the concept of virtual moves cannot be directly applied to the system interacting via potentials in Eq. (5.4), because there is no *a priori* information about the final state during the cluster selection, and hence no information about the corresponding pairwise energy change.

We thus use the non-normalized potential in Eq. (3.1) with $A = A_c$ for the linking potential, and the renormalized potential in Eq. (5.4) for accepting the clusters. This may not be fully microscopically realistic, but the sampling still follows the Boltzmann distribution, because the non-normalized potential only appears in the bias of the cluster acceptance probability, and is the same in the old and in the new state. Using two different potentials in the VMMC simulations is similar to the use of fictitious and real potentials, as was discussed by Whitlam [2011]. Another limitation of the renormalized potentials defined in Eq. (5.4), is that the definition makes sense only if the maximum size of the partition is smaller than one half of the minimum box dimension. This may not be a problem at low packing fractions, but if the compression and the dynamical arrest of the partitions are of interest, the model may become undefined and meaningless.

Let us mention that even the SPMC moves require that the system is parti-

tioned in the old and in the new state, to determine the energy given by potentials in Eq. (5.4). Partitioning of the whole system is not efficient, and the system is instead partitioned locally. This can be done efficiently by using dynamical biconnectivity techniques [Henzinger, 1995]. If particles are moved collectively, we use the local partitioning for every single particle in the cluster.

5.2 Results

Simulations with the potential in Eq. (5.1), and amplitude A defined in Eq. (5.4), proceeded in the same way as in the previous chapters. They started from a high temperature fluid, without high energetic overlaps, which were instantaneously quenched to temperature $T = 0.25$. The size of the system was $N = 2000$, at packing fraction $\phi = 0.01$. We monitored again the sizes of the largest partitions as well as the evolution of the local order.

Fig. 5.2 compares the structure evolution in the SALR system with and without the volume dependent renormalization. Figures 5.2(a) and 5.2(b) show the size of the first, second and the tenth largest partitions. It can be seen that the volume dependent renormalization leads to the stabilization of smaller aggregates, and that partitions have roughly the same size compared to highly polydisperse aggregates in the non-normalized SALR potentials.

Figures 5.2(c)-(f) show that there is an initial time window where local order develops in the same way, independently of whether the long-range repulsion is renormalized by the volume. It can be seen from Figs 5.2(c) and 5.2(d) that the two models start to deviate only at the end of the initial time window, where the low ordered structure starts to relax to a higher ordered structure in the non-normalized model, and this relaxation is inhibited in the renormalized model. Figures 5.2(e) and 5.2(f) show again that the structure evolution is independent of the renormalization in the same initial time window, implying that the number of partitions with crystals inside is similar in the renormalized and non-normalized model at the end of the window. Figure 5.2(f) also shows that a LSC with a crystal inside usually only contains one such crystal, and that the number of those partitions remain constant at later times. This confirms that the crystallinity does not evolve in time in the later stages, but it also suggests that the partitions with crystal inside do not merge, and remain isolated thus preserving the monodispersity. The number of LSCs with no crystal inside keeps growing, reflecting the slow growth in local structure shown in Fig. 5.2(d).

Fig. 5.3 shows that the generation and acceptance densities evolve in time,

implying that phase separation is not complete, and equilibrium is not reached. The acceptance density is close to the generation density both for translations and rotations and even for large clusters. This is a consequence of the generalized linking of Eq. (4.8), the spherical shape of the partitions in the system, and the fact that the moving clusters span the entire partitions. This is confirmed by the fact that the peaks in the generation and acceptance densities at late times correspond to the size of the largest partitions shown in Fig. 5.2(b).




Figure 5.1: Time evolution of the system composed of particles with short-range attraction and renormalized long-range repulsion. The simulation is started from a quenched fluid state. The particle coloring is the same as in Fig. 3.3.

5.3 Discussion

The original renormalization of the potentials [Xia et al., 2011] is more complex. It takes into account the fact that the dielectric constant in the medium inside and outside the partition is generally not the same, and it distinguishes between the repulsive strength A and the screening length ξ inside and outside the partition. To preserve the simplicity of our model, we do not make this distinction. Another

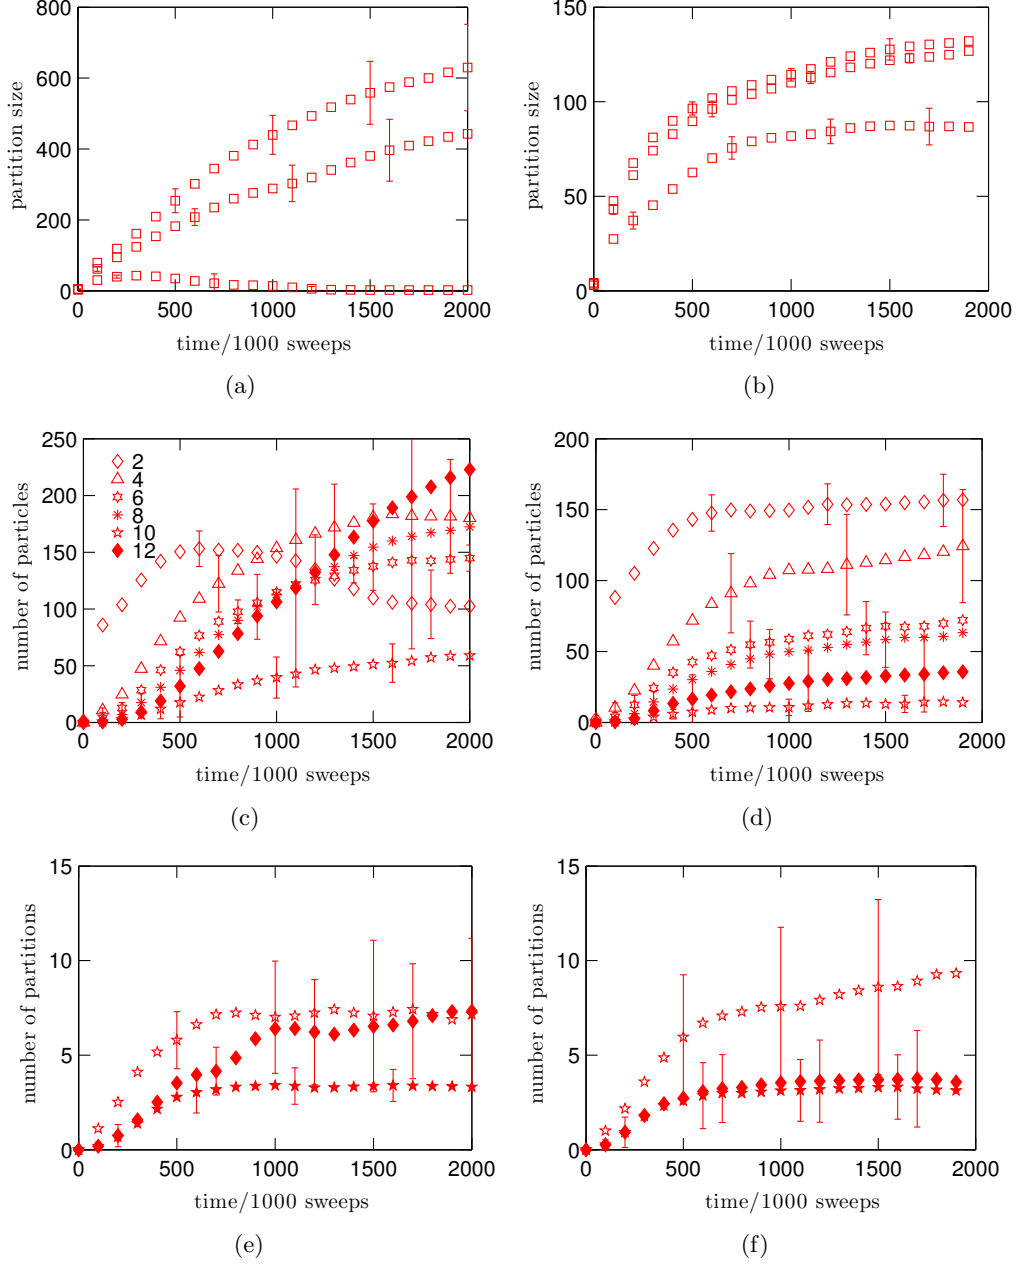


Figure 5.2: Size of the largest partitions and parameters of the local structure for the system interacting via non-normalized SALR potentials (left), and SALR with renormalized amplitude (right). Maximum Monte Carlo displacement is $\delta = 0.20\sigma$. 5.2(a) and 5.2(b) show the size of the first, second and tenth largest partitions in the system. Meaning of the other plots is the same as in Fig. 3.4.




Figure 5.3: Generation and acceptance densities at different times of coarsening in the SALR system with the renormalized amplitude, and VMMC with $\delta = 0.20\sigma$.

reason for this simplification is that the difference between the dielectric constants can be neglected [Royall et al., 2003] for short screening lengths. Whether or not this assumption is the cause for the slow relaxation of local structure shown in Fig. 5.2(d) is not clear and requires further investigation.

The model of Xia et al. [2011] also distinguished between the partitions of size smaller and larger than the critical size. It was assumed that the partitions with size smaller than $0.8R_c$ do not dissolve, and the screening length and the repulsive amplitude were chosen differently than for larger partitions. Their motivation for this assumption was the classical nucleation picture with R_c representing the critical size of the nucleus. We again did distinguish between those sizes. The model of Xia et al. [2011] also distinguished between the pairwise interactions of particles which are in the disordered surfaces of the forming clusters, and particles which are inside the clusters. The corresponding interactions were determined experimentally or computationally such that they match the experiment. To keep our model reasonably simple, these complex interactions were again not considered.

As for the simulation technique, the system of Xia et al. [2011] was repartitioned at regular time intervals. This means the dynamics between the partitioning steps was only approximate and was following potentials not fully corresponding to the potentials defined in Eqs (5.1) and (5.2).

Let us now address the question of whether the partition-based renormalization is realistic. Renormalizing the repulsive part of the pairwise potential according

to the partition size is a possible way of expressing that repulsions between large partitions are stronger than repulsions between small partitions. This is meant to represent the fact that large partitions repel each other stronger than small partitions. There are different models to represent different physics behind the renormalized long-range forces. First, the simplest model assumes that each particle within the partition carries a point Coulomb charge or another long-range pairwise force. The model then simply adds up the long-range effect of this potential close the surface of the partition where this effect is large Sciortino et al. [2004]. Another, more sophisticated type of renormalization [Denton, 2008], is applied to models where multicomponent ion-mixture is responsible for charge stabilization in colloidal suspensions. Although there can be an overlap between the previous model [Sciortino et al., 2004], the model of Denton [2008] averages over the degrees of freedom of salt ions and counterions, while taking into account the size of macroion, and thickness of the spherical macroion shell representing strongly associated counterions. This model seems to be a promising alternative to the model of Sciortino et al. [2004] who integrates over the long-range effect of the standard Yukawa interactions of particles in the spherical partition, which represent screened electrostatic pairwise repulsions. Finally, another common physical source of renormalization is the accumulation of a larger number of polymers or other particles on the surface of large partitions. This repulsion may have form of steric repulsion [Likos et al., 1998] in the case of micelles or polymers where polymeric chains decorate the particle, or generally accumulation repulsion [Louis et al., 2002], where polymer or other particles of the system interact with the self-assembling colloids. The renormalized long-range repulsion can thus have different physical origins, which may compete with each other. Although the linear renormalization considered here or in Xia et al. [2011] is only a first possible approximation, it may thus be a sufficient model for engineering applications.

Nevertheless, we also mention that the model of Xia et al. [2011] or the one of this chapter can still be modified to better represent the physical reality, while preserving the computational cost. In such model, the renormalization of the amplitude according Denton [2008] would be done with respect to the centre of mass of the partition as was suggested by Sciortino et al. [2004], and repartitioning the system at each step as is done here on in Xia et al. [2011]. A disadvantage of this model is again that it is only restricted to the partitions which are mostly spherical during the course of the simulation. The advantage of such model is that the cut-off may still remain relatively short compared to the non-normalized SALR model, because the long-range repulsive effect of particles which are outside the cut-off range is included in the renormalizing integral. Such a model has not been

implemented by us, but it would be interesting to test whether such renormalization leads to a decrease in polydispersity of the aggregates. It is worth noting that as opposed to the renormalization in Eq. (5.4), the renormalization integrating over the amplitudes does not require the provision of the optimum size of the partition R_c as an input parameter for the model.

In both approaches, the renormalization of the repulsion requires a calculation of the dynamical biconnectivity [Henzinger, 1995], which may be computationally expensive and which needs to be done in order to determine the partitions and their centres of mass. The use of the original, non-normalized SALR model, with long-range cut-off is thus still viable and may be the best option to investigate the polydispersity of partitions in quenched SALR systems. Controlled quenches [Royall and Malins, 2012], adjusting the size of the Monte Carlo displacement (Sec. 3.6), may be alternative approaches to reduce the polydispersity of the aggregates, although we know that these systems have equilibrium states with highly polydisperse aggregates [Stradner et al., 2004; Toledano et al., 2009]. This problem is addressed in Secs. 3.6 and 6.1.1. The advantage of the original non-normalized model is that it is considerably easier to implement, especially because biconnectivity routines are not required. Other advantages are that the potentials are defined at high packing fractions, the long-range repulsion remains microscopically realistic, and that partitions do not need to be assumed to be spherical. In agreement with this conclusion, we note that Sweatman et al. [2014] used simple non-normalized SALR potentials simulating large stable partitions with low polydispersity, and showed that their size mainly depends on the strength of attraction and repulsion. Their work emphasizes that the simulation of SALR systems with kinetic (translational and rotational) single particle and collective Monte Carlo moves may lead to long living states of polydisperse partitions, and that smaller, but still very large, partitions do not dissolve on the timescale of the simulation. Their MC simulation thus systematically disperses the smaller partitions in the attempt to accelerate the kinetic evolution, and to find the optimum partition size given by the energy minimum of the system. We note that Charbonneau and Reichman [2007b] used non-kinetic moves with the aim to reach the equilibrium of the cluster phase, while still observing a relatively high polydispersity. However, the two models can only hardly be compared, since Charbonneau and Reichman [2007b] considers a different system with the range of attraction, and the partition size being about two orders of magnitude smaller than the one of Sweatman et al. [2014]. The recent work of Mani et al. [2014] reports formation equilibrium partitions in SALR using simple kinetically realistic simulation. In contrast to Sweatman et al. [2014], their partition sizes are again much smaller,

and the range of attraction is shorter. Moreover, the equilibrium partitions of Mani et al. [2014] can only be achieved at certain temperatures.

5.4 Conclusion

It has been confirmed that the volume dependent renormalization greatly reduces the polydispersity, which is an essential requirement for hierarchical self-assembly. However, it is also pointed out that this renormalization is a crude and computationally expensive approximation which requires *a priori* simulations, a parameter defining the size of the desired partition, and partitioning of the system at each Monte Carlo step. Moreover, the model cannot be used to study gelation or arrest at high densities. The model is related to the model of Sciortino et al. [2004], who suggests another type of renormalization, integrating over long-range repulsions. It is concluded that advanced simulation protocols of the non-normalized potentials may be a simpler model leading to the states of monodisperse clusters, and that a way of defining a computationally tractable and accurate physical model describing the stability or long-time coexistence of monodisperse clusters of cadmium selenide [Xia et al., 2011] remains an open question.

Chapter 6

Discussion

The final discussion focuses on how the Monte Carlo algorithms developed in this thesis may be applied to simulate mesocrystallization on the colloidal scale [Sciortino et al., 2004; Xia et al., 2011] or how to use them to simulate oriented attachment [Leite and Ribeiro, 2012].

6.1 Alignment in the Wigner glass

In the first part of the discussion, we describe a hypothetical crystallization pathway, which is similar to mesocrystallization. This pathway proceeds via alignment of nanocrystals in a low density arrested state of matter called the Wigner glass and subsequent oriented attachment realized by compressing the Wigner glass. We use the results of this thesis to suggest how to make such a pathway realistic in a simulation and possibly also in an experiment. Briefly, the pathway is initiated by a moderately deep quench of a low density fluid which leads to a phase separation of the system into monodisperse spherical crystalline aggregates. The crystals do not merge and are assumed to have enough time to crystallographically align. The resulting low density system of aligned clusters is slowly compressed to form a single crystalline block.

6.1.1 Formation of monodisperse aggregates

Phase separation of the system into isolated crystals can be achieved at low packing fractions and by moderately deep quenches into a suitable temperature, obtained either from the analysis of fluctuation-dissipation ratios [Jack et al., 2007; Klotsa and Jack, 2011] or from the estimate of the intersection of the glass line and the spinodal [Zaccarelli, 2007]. The temperature must be low enough such that the spinodal

decomposition drives the phase separation, but high enough such that bond breaking and bond making can regularly happen [Grant et al., 2011; Grant and Jack, 2012], and crystallinity within the dense regions can evolve. A simple quench only leads to low quality crystals or polycrystals. A controlled quench [Royall and Malins, 2012] or variation of interactions according to some kind of real time feedback [Klotsa and Jack, 2013] can improve the crystal yield. The size of the system must be reasonably small such that the timescales on which the non-crystalline partitions (liquid drops) dissolve can be achieved by the computer simulation. On the other hand, the system size must be large enough such that a sufficient number of crystals is formed, and further self-assembly into larger structures can be studied. The number of crystals (or equivalently their size) can be controlled by other parameters of the system such as the packing fraction or the shape of the potential.

A common and necessary prerequisite for hierarchical self-assembly or for mesocrystallization is to require that the self-assembled aggregates have low polydispersity, to provide new building blocks for further self-assembly of larger aggregates. We observed that monodispersity may be achieved in different ways.

First, polydispersity may simply be a consequence of the simulation parameters such as the size of the maximum displacement δ . If δ is small enough, and the quench starts from an equilibrium fluid, the system is expected to condense into a large number of liquid drops. These drops exist for a long time even in the absence of long-range repulsions. If δ is large, the polydispersity of the aggregates is seen to be increased. We thus expect that it is preferable to perform simulations with a low δ , at least at those stages of phase separation where most particles form parts of larger partitions.

Second, polydispersity can also be reduced by a suitable renormalization of the long-range repulsion. In Chap. 5, the amplitude of the pairwise potential was renormalized by the size of the parent partitions, which leads to a great reduction of polydispersity of the aggregates developed after the quench. The renormalization, however, is so approximate that it may be considered non-physical or *ad hoc* with the aim to achieve monodispersity. Other, more physical partition-based renormalization models may be derived from integration of long-range repulsion [Sciortino et al., 2004].

Third, if the aim is only to obtain partitions with a low polydispersity, the original non-normalized model defined in Eq. (3.1) may be sufficient, provided the simulation starts from a uniformly distributed set of particles in the fluid phase, and provided a small maximum displacement δ is used at the beginning of the simulation. Controlling the nucleation at early times is, indeed, a way of experimentally

achieving low polydispersity (Sec. 1.6). The two previously described renormalization models are based on a dynamical biconnectivity routine [Henzinger, 1995], which may be computationally expensive or hard to implement. The advantage of the non-normalized model in Eq. (3.1) is that biconnectivity is not needed, and so the model can be easily implemented. In fact, the computational cost of evaluating the long-range repulsion is comparable to the cost of biconnectivity. Nevertheless, as discussed later, if the monodisperse partitions are to be compressed into a cluster crystal or a mesocrystal, a renormalization of the repulsion representing the polymers accumulated on the partitions is needed, to avoid reorganization of the partitions into larger domains during the compression.

Fourth, as recently emphasized by Sweatman et al. [2014] or by Zhang et al. [2012], the polydispersity in SALR system, where the range of attraction is large, may only be a kinetic phenomenon, and a systematic dispersion of smaller partitions may accelerate the simulation with the aim to reach the equilibrium state characterized by partitions with low polydispersity. It is emphasized here that this may only be valid for longer-range attractions, as the simulation of the short-range attractive model of Charbonneau and Reichman [2007b], which uses non-kinetic moves equilibrating the system, still shows clusters with a high polydispersity.

Let us stress that to achieve timescales on which most drops crystallize and liquid drops dissolve may require a long simulation. Our analysis also suggests that an acceleration might also be achieved, if the expensive simulation of SALR systems by VMMC is approximated by a SPMC simulation of the same system but without the long-range repulsion. A similar approximation may be possible for the renormalized potentials. The collective moves would then be used only sparingly, in order to equilibrate the partition-partition distances. This approximation is made possible by facts that the structure evolution happens mainly via single-particle movements and exchange, and that the contribution of collective modes of motion to long-time structural changes can be neglected (integrated out).

We also point out that if the long timescale cannot be achieved, we can construct the state of isolated partitions artificially. Basin hopping techniques [Wales, 2003] allow us to predict the parameters of SALR potentials and a suitable temperature at which the spherical partitions are stable [Mossa et al., 2004]. These crystalline partitions can then be periodically replicated in a simulation box such that nearest-neighbour inter-partition distances are roughly equal, and the partitions are rotated such that their relative lattices have random orientations. Although the long-range nature of the potential may make this state of isolated partitions unstable [Toledano et al., 2009], and particles within the partitions may reorganize to

form a more stable state [Malins et al., 2011], we assume that this effect is negligible at low densities. If the effect is large, we may attempt to equilibrate the system, or chose other properties and parameters of the system.

Generation of monodisperse partitions, both artificially and via quenches, is a preparation protocol, which results in a fluid of clusters or in a non-equilibrium (stationary) system, which has the properties of a cluster fluid for a long time. The protocol can be such that the Wigner glass instead of a fluid of clusters is produced. This can be achieved if the quench is performed slowly, and at suitable packing fractions. The assumption about the possibility of reaching a Wigner glass composed of monodisperse partitions will be discussed in what follows.

6.1.2 Alignment of partitions

To illustrate a possible way of aligning the crystal lattices of the isolated partitions, we will assume that the system is in a kinetically slowed down (arrested) state with properties similar or identical to the properties of a Wigner glass. In particular, we will make use of the numerical evidence that the isolated partitions within a Wigner glass do not move via translational motion, and that the main degrees of freedom are collective rotations [Toledano et al., 2009]. This is indirectly confirmed by our VMMC simulations of fluids of partitions showing that as long as the partitions are spherical, large rotations are accepted more often than large translations. We may thus expect that the clusters within the Wigner glass have enough time to rotate until they align.

In such a crystal alignment process, the rotational entropy contributing to the free energy would compete with the potential energy of the system, and the transition between the aligned and rotationally disordered Wigner glass might be a first order transition with a distinct free energy barrier between the two states. The waiting time needed for the system to cross such barrier can be related to the timescale on which the properties of the Wigner glass persist. In the following, we will compare those timescales to the timescales of an Ostwald ripening process. This allows us to characterize the Wigner glass, and to tell whether it is stable enough for the alignment to occur.

In particular, the timescale of an Ostwald ripening process is given by the time needed for the detachment of a particle from its partition, and the average time needed to diffuse to another partition. We characterize the timescales of the stability of the Wigner glass as the timescale needed for two partitions to merge. These two timescales can be seen as the stability timescales, and we can compare them with the average time needed for all partitions to align. If any of the two timescales needed

for stability is shorter than the timescale needed for alignment, the partitions grow or merge before they align. (Unless the collisions are effective as will be discussed later). Hence, a necessary condition for a system of isolated monodisperse crystalline partitions to align is that the timescale for stability against aggregation or against growth via Ostwald ripening must be longer than the timescale needed for alignment.

We note that in a system where the phase of uniformly sized clusters is thermodynamically stable, a steady exchange of single particles between the partitions as well as continuous splitting and merging of the partitions might occur. The reversible nature of these changes, and the fact that the partitions are stable, suggests that the earlier defined timescale for stability of a cluster phase can be arbitrary with respect to the timescales needed for alignment.

6.1.3 Compression into a mesocrystal

Wigner crystals are solids whose density is far too low to form a mesocrystal. A higher density may be achieved if the Wigner crystal or the cluster fluid are compressed or crunched [Mossa et al., 2002]. Crunching can be an alternative to quenching: both can reach a certain state point of the phase diagram [Royall and Malins, 2012]. Here, it is assumed that, depending on whether the crystals within the initial Wigner crystal are aligned or not, the compression leads either to a low density cluster crystal [Sciortino et al., 2004; Toledano et al., 2009] or to a mesocrystal. A similar compression of a very low density cluster fluid was used by Xia et al. [2011] to obtain a high density arrested state of uniformly sized CdSe nanoparticles, which were reminiscent of the cluster crystal or a glass of clusters.

Now the compression of the system may not be straightforward. It is known that the size of the wave vector corresponding to the first peak in the structure factor has been reported to be invariant with respect to the packing fraction [Stradner et al., 2004; Cardinaux et al., 2007]. This means that the position of the peak should remain constant during a sufficiently slow compression, and that slow crunches of the SALR system into lower densities may have a fundamental significance. However, it is known that partitions reorganize [Malins et al., 2011] and their size grows with increasing packing fraction [Sciortino et al., 2004]. Also, higher densities may be accompanied with transitions into modulated phases [Archer and Wilding, 2007]. Moreover, it is expected that the growth is not only given by the partition-partition interactions or translational entropy, but also by the temperature dependence of the interfacial tension [Stradner et al., 2004]. These results suggest that a fast compression or a non-trivial renormalization of the potential similar to that in Sec. 5 or in Xia et al. [2011] is necessary to produce a low density cluster crystal or a

mesocrystal composed of monodisperse building blocks.

We also point out that the re-entrant percolation in the non-normalized SALR system [Sciortino et al., 2005] may indicate another way of reaching higher packing fractions by compressions in systems where the potential has another shape.

6.2 Alignment via oriented attachment

In the previous section we considered a possible mesocrystallization pathway where crystal alignment was based on the fact that crystal partitions can easily rotate within the Wigner glass, and can align even before they are in close contact. A subsequent compression of the aligned system then ensures low density and the integrity of the desired mesocrystalline structure.

While parts of this mechanism are realistic both experimentally and in a simulation, the possibility of alignment within the Wigner glass is newly discussed here and is only hypothetical. Let us now address the role of collective rotational motion in another mechanism of mesocrystal formation which, as has been described previously by Leite and Ribeiro [2012], is based on the orientation dependent particle collisions.

In this model, partitions are represented by a system of equally sized and shaped nanocrystals with a crystal lattice defining their orientation. The system behaves like a gas of nanocrystals undergoing phase separation or gelation depending on the strength of interaction. The interaction between the partitions is special in that its strength is highly dependent on whether the colliding partitions are crystallographically aligned or misoriented. If the collision between the nanocrystals is such that their lattice orientations match at contact (epitaxial alignment), the overall attractive interaction between the partitions is expected to be strong, irreversible, and the collision is referred to as *effective*. If the collision is such that the lattices are not aligned (nonepitaxial attachment), the energy associated with the creation of the interface is not supposed to be high, the collisions may be reversible and termed *non-effective*.

We will distinguish between reversible and irreversible collisions, and will consider two possible relative rotations. First, partitions collide and split in a non-effective fashion until an effective collision occurs. In this case rotations happen while partitions are isolated, i.e. do not form a single partition through any short-range contact between the two partitions. Second, partitions remain attached after a non-effective collision, and relative rotations of the two partitions (more precisely two subpartitions within a single partition) happen until they align.

In systems where long-range repulsion is strong enough (dispersed colloidal state), the crystalline partitions are expected to collide until the collision is such that the respective crystalline surfaces are crystallographically aligned. We know that the binding and unbinding probability is exponentially dependent on the energy of attachment [Grant et al., 2011], and that even a small energy difference between the non-aligned and aligned attachment may lead to the former being able to break much easier than the latter.

Irreversibility of the effective collisions and reversibility of the non-effective collisions are thus plausible assumptions. Nevertheless, the probability of an ideal collision where the crystal surfaces and their edges are aligned is low. The nanocrystals seem to be more likely to form an interface when they are part of a single partition. The driving force for alignment can again result from favourable energetics, interfacial tension [Leite and Ribeiro, 2012], or from the internal strains [Chen et al., 2014].

Our simulations do not lead to a perfect collection of uniformly sized and shaped nanocrystals, but we are able to obtain a system composed of isolated aggregates with a high internal crystallinity and with a well developed crystal surface. We propose improvements to the model, and to the simulation protocol, such that higher crystallinity and lower polydispersity may be achieved. This allows us to assume an equilibrium fluid composed of similarly shaped crystals with random lattice orientations, and point out the overlaps in mechanisms observed both in our model and in the collision model. We discuss in Sec. 3.6 that the renormalization of the partitions into larger spheres leads not only to shorter attraction range relative to the spheres, but also to the increase of its strength due to the larger number of microscopic contacts between particles constituting the two colliding partitions. The number of those contacts is generally highest if the attachment is epitaxial. The epitaxial alignment is thus energetically and not entropically favoured.

Our VMMC simulations showed that the intrapartition collective rearrangements are unlikely, especially if the partitions have high crystallinity. This suggests that the simulations do not capture the large collective rearrangements known as grain-rotation-induced grain growth mechanisms. This may be due to the simulation parameters, but also due to other properties of the system. The centres of rotations may be chosen in a more realistic or in a more efficient way. The geometry of the pairwise potential also plays a role. For example, Whitelam et al. [2009a] reported that better assembly by collective motion can be achieved in viral capsids than in spherical systems. It is also clear that flat surfaces composed of cubical nanoparticles [Smallenburg et al., 2012] can glide in contact more easily than rough

surfaces composed of spheres. The temperature choice is also worth specifying. It should not be higher than the melting temperature, but must be high enough such that self-recrystallization of the grain boundaries can happen after the oriented attachment [Leite and Ribeiro, 2012]. In the recent simulation study of Bjerre et al. [2013], the optimum temperature for grain-rotation-induced grain mechanism has been identified in the vicinity of the melting temperature. It would be interesting to compare this temperature with the temperature region discussed in this thesis.

Finally, precise control over the rotations in the generalized VMMC may also affect the non-classical crystallization process. If the crystalline partitions are represented by anisotropic particles, with specific directional interactions defined by their lattice orientation, the results of Hedges and Whitelam [2011] suggest that a cluster liquid is more likely to precede the cluster solid, if clusters rotate sluggishly. Hedges and Whitelam [2011] also claim that partitions must be aligned in order to crystallize, and that the specific inter-partition interactions must have an optimum value, in analogy with what has been discussed above.

6.3 Micelle as a building block

We will now briefly use the framework of this thesis to suggest another possible class of crystallization pathways leading to synthetic crystals with micellar inclusions [Kim et al., 2011] or to mesocrystals such as the sea urchin spine [Seto et al., 2012], where inclusions of biomolecules and amorphous calcium carbonate are present in the crystalline blocks of the mesocrystal. Star polymers, micelles and other sterically stabilized colloids are soft units with well established potentials, and with known basic properties of the phase diagram [Zaccarelli, 2007; Sweatman et al., 2014].

We will assume a mixture of biomolecules and calcium carbonate, and that the amorphous calcium carbonate is accumulated within the biomolecule or copolymer thus forming the micelles. We further assume that the copolymer provides a long-range repulsion to the micelle, sterically stabilizing the liquid drop of the amorphous calcium carbonate, and that the micelle has a short-range attraction having one of the origins discussed in Sec. 1.3.

Similarly to what has been shown for SALR potentials, the system can phase separate into clusters of micelles, which may survive for a long time and undergo a further self-organization into higher ordered structures, such as the Wigner glass, low density cluster crystal, or a cluster fluid. Such higher ordered assembly may include additional directional forces of the copolymer micelles, which may be represented by attractive or repulsive patches decorating the soft particles and mediating the

self-assembly. A characteristic feature of the assembled state is that the atomistic structure of the calcium carbonate drops present in the micelles is still amorphous.

Depending on the result of such assembly, the aggregates either form large isolated structures, which recrystallize on the atomistic scale into calcite crystals with micellar inclusions [Kim et al., 2011] or form an organized network (cluster crystal) of such amorphous structures and the crystallinity propagates within that network to form the mesocrystal [Seto et al., 2012]. This interpretation suggests that the mechanism behind the incorporation of micelles into calcite crystals is similar in the synthetically produced crystal and in the biological mesocrystal, and also that the macroscopic self-assembly of amorphous drops precedes the microscopic crystallization. Within this picture, it is also plausible that biomolecules or copolymers stabilizing and self-assembling the amorphous drops of calcium carbonate play an additional role in the recrystallization of the calcite crystal, because the bimolecular inclusions are oval shaped with a preferential direction. We stress that this mesocrystalline route is different from the originally proposed route of oriented attachment between the crystalline particles.

Chapter 7

Conclusion

This thesis is an attempt to use the recent results of colloidal physics in order to propose a minimal model and a computer simulation protocol which would exhibit the kind of self-assembly processes observed during mesocrystal formation.

We start from the prototypical model for protein or colloidal crystallization, short-range attractive spheres, and we focus on low concentrations, because those are characteristic features of solutions in which mesocrystals nucleate, grow, and self-assemble. We then use a well known observation in colloidal physics that if systems under those conditions are quenched to a suitable temperature region, a strong clustering into liquid drops within the system emerges, and the speed of crystallization within these drops can be easily controlled by the temperature. We then model the self-assembly across different scales, and we study the competition between the evolution of local order and the order characterized by the size of the aggregates.

The simulations in this thesis are able to reach the timescale at which the local crystallinity starts developing across the drops, but not the timescale at which the crystal order stops evolving. This points to the fact that equilibrium is not reached due to the presence of kinetically slowed phase separation caused by the proximity of the glass transition in the locally dense regions. We show that long-range repulsion is needed to stabilize the droplets against their aggregation into a single partition. The system size, which was tractable in our simulation, led only to a few metastable clusters, and a system with a size of at least an order of magnitude larger would be needed to model the formation of a cluster crystal or a mesocrystal.

As part of the minimalist approach to our model, the potentials under consideration were spherical. The monomers building up real mesocrystals are believed to be particles interacting via anisotropic interactions, which are not straightforward

to model by molecular dynamics based algorithms. This thesis was thus mainly focused on the development of the Virtual Move Monte Carlo algorithm, which is an algorithm allowing an easy control of single particle and collective movements of particles interacting via complex anisotropic potentials, and could potentially be used in the simulation of mesocrystallization.

The advantage of the Virtual Move Monte Carlo cluster algorithm over similar cluster algorithms is that the cluster is selected dynamically according to the pairwise energy gradients. This allows us to guarantee a physically more realistic topology of the moving clusters, and to approximate the kinetic evolution of the system by Monte Carlo simulation in situations where collective modes of motion cannot be neglected. The derivation and implementation of the algorithm involves novel and non-trivial concepts in the design of Monte Carlo schemes, which we believe were not sufficiently described in the previous papers. The first part of this thesis thus provides a concise description of the simulation method, along with the formal proof of its validity. It is then shown that the maximum displacement size δ significantly affects the cluster size distribution and the efficiency of the algorithm. The moving clusters tend to span the entire system for large δ , and the cluster motion is negligible for small δ . The choice of the optimum δ may depend on the properties of the system and the aims of the simulation. In an equilibrium fluid, the moving cluster size distribution may follow a universal exponential behaviour. In the simulation of phase separation, the moving cluster size distribution develops over time towards larger clusters, and intermediate sized rearrangements become rare. It is proposed that a slow lowering of δ at those stages is likely to enhance the efficiency of the simulation. In the simulation of stabilized isolated aggregates, the cluster generation and acceptance probabilities also evolve in time for different sizes of the moving clusters. It is suggested that a slow increase of δ during the course of the simulation may enhance the sampling in the multiscale modelling of hierarchical self-assembly. Interestingly, simulations of kinetically slowed down phase separation led to the same product, independently of the choice of δ , and in spite of entirely different trajectories leading to the final kinetically arrested states. This suggests that the products of VMMC simulation may have universal significance, and may be physical.

We further discussed the importance of collective rotation operations for oriented attachment mechanisms. We showed that in most situations the collective rotation operations are ineffective in that their acceptance probability density is significantly lower than the generation density. Collective rotational degrees of freedom in the standard symmetrized VMMC (Chap. 2) are thus not sufficiently represented.

It turned out that this is a consequence of the inability of the standard VMMC to link particles to the cluster through purely repulsive collisions. Chapter 4 solves this problem by listing the complete set of virtual states and by generalizing the link formation functions, which capture repulsive interactions, and which increase the acceptance rate of rotations by up to two orders of magnitude.

Formation of isolated amorphous or crystalline aggregates is a common stage of mesocrystallization pathways. Chapter 3 analyses possibly the simplest way of preparing such amorphous or crystalline aggregates by quenching attractive spheres at low densities from a high to a suitably chosen low temperature. It is confirmed that in the absence of collective motion, the system phase separates into isolated partitions. Three linear growth regimes are then observed. The first regime is characterized by a rapid formation of drops with and without a crystal inside. The second regime is given by the dissolution of drops without a crystal inside. The last regime is slowest, and large crystals grow at the expense of smaller ones by a single particle exchange similar to the Ostwald ripening. It is shown that the choice of δ systematically affects the polydispersity of the metastable drops. It is then also demonstrated that, if the collective motion is allowed, the system rapidly aggregates into a single amorphous partition, and that crystallinity within that partition evolves at about the same pace as in the absence of collective motion.

It is shown that the addition of only a weak long-range repulsion to the short-range attraction increases significantly the lifetime of the drops. Whether or not these drops are thermodynamically stable is not confirmed by our simulations, and the ongoing slow growth of the drops suggests that phase separation into a single crystal is underway. The important result is that the lifetime of the drops is long enough for local crystallinity to develop, and that these states are similar to the metastable states of isolated drops forming a mesocrystal. One of the main obstacles to hierarchical self-assembly at this scale is the persisting polydispersity of the drops. In analogy to the previous work of Xia et al. [2011], it is shown in Chap. 5 that the polydispersity can be significantly decreased by an *ad hoc* volume dependent renormalization. At the same time, it is stressed that this type of renormalization is very approximate, simulationally expensive, and not easy to implement.

This thesis presents a general way of applying Monte Carlo algorithms to coarsening and self-assembly at colloidal, and possibly on atomistic, scales with the focus on how to interpret these results in modelling different stages of various types of mesocrystallization pathways. We have shown that the VMMC, and its generalized version in particular, may be suitable for the simulation of oriented attachment. The algorithm is expected not only to simulate the aligning and rotating nanocrystals,

but may also capture the propensity to dissolve via single particle or small cluster motion or to capture the role of single particle motion in the recrystallization of the mesocrystal. This may be used to model oriented attachment both via effective and non-effective collisions. To model oriented attachment via grain-rotation-induced grain coalesce mechanisms may also be possible by VMMC, but the parameters of the model and the simulation may need to be carefully adjusted.

The VMMC algorithm, or more precisely the collective motion, proved to be essential in the formation of gels or disordered aggregates, and our simulations confirm previous observations that the crystallinity within these kinetically slowed down structures may still evolve. Mesocrystals formed in weakly flocculated states may follow a similar route. A gel or a prestructured amorphous network develops first, and local crystallinity propagates afterwards within that network. In such situations the VMMC may be used both to simulate the network formation, and the subsequent evolution of local crystal order at the atomistic scale.

Appendix A

The inconsequential nature of the asymmetric core in Sec. 2.5

In this appendix, we show why the asymmetric core of the realization of the cluster, in the algorithm of Sec. 2.5, does not matter. Let us now assume that $R_{\mathcal{C}}^{(\mu)}$, with $\mathcal{L}^{(\mu)}$, $\mathcal{F}^{(\mu)}$, $\mathcal{B}^{(\mu)}$, is the realization of \mathcal{C} , generated by the restricted recursive selection of Sec. 2.5.

Let (i, j) be a pair from $\mathcal{L}^{(\mu)}$. Any such pair is interacting in μ or in μ_i , and it must hold that $i \in \mathcal{C}$, $j \in \mathcal{C}$, $p_{ij}^{(\mu)} > 0$. We start with the trivial case: if (i, j) is interacting in μ , then (i, j) is interacting in ν , and a link is, indeed, proposed to that pair in ν . If (i, j) is not interacting in μ , but interacting in μ_i , then, with the aid of Eq. (2.21), we can write

$$p_{ij}^{(\mu)} = \min \left\{ p_{i'j}^{(\mu)}, p_{ij'}^{(\mu)} \right\} = p_{ij}^{(\nu)} > 0. \quad (\text{A.1})$$

Equations (2.6), (2.15a), and (2.15b) then imply that $\epsilon_{i'j}^{(\nu)} > 0$ [(i, j) is interacting in ν_i], and hence that a link can also be proposed to (i, j) in state ν . Hence, given $\mathcal{L}^{(\mu)}$ in $R_{\mathcal{C}}^{(\mu)}$, we can assume the existence of $\mathcal{L}^{(\nu)}$ such that $\mathcal{L}^{(\nu)} = \mathcal{L}^{(\mu)}$.

Let (i, j) be a pair from $\mathcal{F}^{(\mu)}$. We distinguish between two cases, and decompose $\mathcal{F}^{(\mu)}$ as

$$\mathcal{F}^{(\mu)} = \mathcal{F}^{*(\mu)} \cup \mathcal{F}^{\dagger(\mu)}, \quad (\text{A.2})$$

where $\mathcal{F}^{\dagger(\mu)}$ denotes pairs (i, j) not interacting in μ but interacting in μ_i such that $p_{ij}^{(\mu)} = 0$. The set $\mathcal{F}^{*(\mu)}$ then denotes all other pairs in $\mathcal{F}^{(\mu)}$, and is thus composed either from pairs which interact in μ , or from pairs which do not interact in μ , but interact in μ_i such that $p_{ij}^{(\mu)} > 0$. Considerations, similar to those done for $\mathcal{L}^{(\mu)}$, imply that there exists $\mathcal{F}^{*(\nu)}$ such that $\mathcal{F}^{*(\nu)} = \mathcal{F}^{*(\mu)}$. However, it is generally not

guaranteed that if (i, j) is from $\mathcal{F}^{\dagger(\mu)}$, then (i, j) ($i, j \in \mathcal{C}$) is also interacting in ν_i .

The fact that for each $\mathcal{B}^{(\mu)}$ there exists $\mathcal{B}^{(\nu)}$ such that $\mathcal{B}^{(\nu)} = \mathcal{B}^{(\mu)}$ follows from the properties of the cluster selection conditions (2.1), from the existence of $\mathcal{L}^{(\nu)} = \mathcal{L}^{(\mu)}$, and from the fact that if a boundary pair is interacting in μ or in μ_i , it is also interacting in ν or in ν_i .

We have thus shown that it is only due to $\mathcal{F}^{\dagger(\mu)}$ that pairs to which a link is proposed in states μ cannot be guaranteed to be the same in ν , and we cannot ensure the existence of $R_{\mathcal{C}}^{(\nu)}$ such that $R_{\mathcal{C}}^{(\nu)} = R_{\mathcal{C}}^{(\mu)}$. Nevertheless, identity (2.14) implies that any pair proposed to $\mathcal{F}^{\dagger(\mu)}$ has no chance to form in state ν , and we can assume that a link was proposed to that pair although it was actually not. Hence, the set of pairs to which a link is proposed in the recursive selection of the cluster can in practice be assumed to be the same. The restricted cluster selection can be seen to be equivalent to the free cluster selection, in other words the cluster can be accepted with probability given by Eq. (2.27), even if the links are only proposed to pairs interacting in μ or in μ_i , and not to all pairs satisfying (2.1a) and (2.1b).

Appendix B

Derivation of the acceptance probability of the algorithm of Sec. 2.8

To derive the bias in the cluster acceptance probability needed to correct the asymmetry of the realization of a cluster selected by the procedure described in Sec. 2.8, we define the linking probabilities as

$$P_{i'j}^{(\mu)} = I_{ij}^{(\mu)} p_{i'j}^{(\mu)}, \quad Q_{i'j}^{(\mu)} = 1 - P_{i'j}^{(\mu)}, \quad (\text{B.1})$$

where $I_{ij}^{(\mu)} = 1$, if (i, j) is interacting in μ , and $I_{ij}^{(\mu)} = 0$, if (i, j) is not interacting in μ .

For simplicity, we assume again that $R_{\mathcal{C}}^{(\mu)}$ is a realization resulting from the free recursive selection of \mathcal{C} , but using the linking probabilities (B.1). Since the probabilities (B.1) satisfy relations (2.3a) and (2.3b), one can assume the existence of $R_{\mathcal{C}}^{(\nu)}$ such that $R_{\mathcal{C}}^{(\nu)} = R_{\mathcal{C}}^{(\mu)}$, allowing us to require the SDB between $R_{\mathcal{C}}^{(\mu)}$ and $R_{\mathcal{C}}^{(\nu)}$, and then simplify the acceptance probability, in a similar way as in Sec. 2.4, into the form

$$W_{\text{acc}}^{(\mu \rightarrow \nu | R)} = \min \left\{ 1, \exp(-\beta(E_{\nu} - E_{\mu})) \frac{\prod_{(i,j) \in \mathcal{B}^{(\nu)}} Q_{i'j}^{(\nu)}}{\prod_{(i,j) \in \mathcal{B}^{(\mu)}} Q_{i'j}^{(\mu)}} \right\}. \quad (\text{B.2})$$

The Rosenbluth factor in Eq. (B.2) can be simplified as follows. Let us decompose

the boundary of \mathcal{C} into

$$\mathcal{B}^{(\mu)} = \mathcal{B}^{(\mu\nu)} \cup \mathcal{B}^{(\bar{\mu}\nu)} \cup \mathcal{B}^{(\mu\bar{\nu})} \cup \mathcal{B}^{(\bar{\mu}\bar{\nu})}, \quad (\text{B.3})$$

where $\mathcal{B}^{(\mu\nu)}$ are pairs from \mathcal{B} which interact both in states μ and ν , $\mathcal{B}^{(\bar{\mu}\nu)}$ are boundary pairs which do not interact in μ but interact in ν , $\mathcal{B}^{(\mu\bar{\nu})}$ are boundary pairs which interact in μ but not in ν , and $\mathcal{B}^{(\bar{\mu}\bar{\nu})}$ are boundary pairs which do not interact both in μ and in ν . Boundary $\mathcal{B}^{(\nu)}$ in state ν can be decomposed similarly. Equation (2.7) and definition (2.8) then imply that

$$\Delta\epsilon_{i'j}^{(\mu)} = -\Delta\epsilon_{i'j}^{(\nu)} \quad \forall (i, j) \in \mathcal{B}^{(\mu\nu)}, \quad (\text{B.4a})$$

$$\Delta\epsilon_{i'j}^{(\nu)} = -\epsilon_{ij}^{(\nu)} \quad \forall (i, j) \in \mathcal{B}^{(\bar{\mu}\nu)}, \quad (\text{B.4b})$$

$$\Delta\epsilon_{i'j}^{(\mu)} = -\epsilon_{ij}^{(\mu)} \quad \forall (i, j) \in \mathcal{B}^{(\mu\bar{\nu})}. \quad (\text{B.4c})$$

Using Eq. (2.2), the total energy of the system in state μ can then be expressed as

$$E_\mu = \sum_{(i,j) \in \mathcal{L}^{(\mu)}} \epsilon_{ij}^{(\mu)} + \sum_{(i,j) \in \mathcal{F}^{(\mu)}} \epsilon_{ij}^{(\mu)} + \sum_{(i,j) \in \mathcal{B}^{(\mu\nu)}} \epsilon_{ij}^{(\mu)} + \sum_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}} \epsilon_{ij}^{(\mu)} + \sum_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}} \epsilon_{ij}^{(\mu)} + \sum_{(i,j) \in \mathcal{B}^{(\bar{\mu}\bar{\nu})}} \epsilon_{ij}^{(\mu)} + \sum_{(i,j) \in \mathcal{X}^{(\mu)}} \epsilon_{ij}^{(\mu)}. \quad (\text{B.5})$$

Similarly, the total energy of \mathcal{S} in state ν is

$$E_\nu = \sum_{(i,j) \in \mathcal{L}^{(\nu)}} \epsilon_{ij}^{(\nu)} + \sum_{(i,j) \in \mathcal{F}^{(\nu)}} \epsilon_{ij}^{(\nu)} + \sum_{(i,j) \in \mathcal{B}^{(\mu\nu)}} \epsilon_{ij}^{(\nu)} + \sum_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}} \epsilon_{ij}^{(\nu)} + \sum_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}} \epsilon_{ij}^{(\nu)} + \sum_{(i,j) \in \mathcal{B}^{(\bar{\mu}\bar{\nu})}} \epsilon_{ij}^{(\nu)} + \sum_{(i,j) \in \mathcal{X}^{(\nu)}} \epsilon_{ij}^{(\nu)}. \quad (\text{B.6})$$

Since, $\mathcal{L}^{(\mu)} = \mathcal{L}^{(\nu)}$, $\mathcal{F}^{(\mu)} = \mathcal{F}^{(\nu)}$, $\mathcal{X}^{(\mu)} = \mathcal{X}^{(\nu)}$, and $\epsilon_{ij}^{(\mu)} = \epsilon_{ij}^{(\nu)}$, for pairs (i, j) from these sets, we can express the energy difference as

$$E_\nu - E_\mu = \sum_{(i,j) \in \mathcal{B}^{(\mu\nu)}} \epsilon_{ij}^{(\nu)} - \sum_{(i,j) \in \mathcal{B}^{(\mu\nu)}} \epsilon_{ij}^{(\mu)} + \sum_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}} \epsilon_{ij}^{(\nu)} - \sum_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}} \epsilon_{ij}^{(\mu)}, \quad (\text{B.7})$$

where we have also used the fact that

$$\sum_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}} \epsilon_{ij}^{(\mu)} = \sum_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}} \epsilon_{ij}^{(\nu)} = \sum_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}} \epsilon_{ij}^{(\nu)} = \sum_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}} \epsilon_{ij}^{(\mu)} = 0.$$

Let us separate the terms in Eq. (B.7) for which $\epsilon_{ij} > 0$ from those for which $\epsilon_{ij} \leq 0$,

and identify the corresponding partial sums with $+$ and $-$ respectively, so that

$$E_\nu - E_\mu = \sum_{(i,j) \in \mathcal{B}^{(\mu\nu)}} \epsilon_{ij}^{(\nu)} - \sum_{(i,j) \in \mathcal{B}^{(\mu\nu)}} \epsilon_{ij}^{(\mu)} + \sum_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}}^+ \epsilon_{ij}^{(\nu)} + \sum_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}}^- \epsilon_{ij}^{(\nu)} - \sum_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}}^+ \epsilon_{ij}^{(\mu)} - \sum_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}}^- \epsilon_{ij}^{(\mu)}. \quad (\text{B.8})$$

Hence,

$$\exp(-\beta(E_\nu - E_\mu)) = \frac{\prod_{(i,j) \in \mathcal{B}^{(\mu\nu)}} \exp(-\beta\epsilon_{ij}^{(\nu)}) \prod_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}}^+ \exp(-\beta\epsilon_{ij}^{(\nu)}) \prod_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}}^- \exp(-\beta\epsilon_{ij}^{(\nu)})}{\prod_{(i,j) \in \mathcal{B}^{(\mu\nu)}} \exp(-\beta\epsilon_{ij}^{(\mu)}) \prod_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}}^+ \exp(-\beta\epsilon_{ij}^{(\mu)}) \prod_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}}^- \exp(-\beta\epsilon_{ij}^{(\mu)})}, \quad (\text{B.9})$$

where the $+$ and $-$ symbols have a meaning analogous to that in Eq. (B.8). The next step is to decompose the products of link failure probabilities in (B.2), in a similar way. Using Eq. (2.19), the ratio of products can be expressed as

$$\frac{\prod_{(i,j) \in \mathcal{B}^{(\nu)}} Q_{i'j}^{(\nu)}}{\prod_{(i,j) \in \mathcal{B}^{(\mu)}} Q_{i'j}^{(\mu)}} = \frac{\prod_{(i,j) \in \mathcal{B}^{(\nu)}}^+ \exp(-\beta\Delta\epsilon_{i'j}^{(\nu)})}{\prod_{(i,j) \in \mathcal{B}^{(\mu)}}^+ \exp(-\beta\Delta\epsilon_{i'j}^{(\mu)})} = \frac{\prod_{(i,j) \in \mathcal{B}^{(\mu\nu)}}^+ \exp(-\beta\Delta\epsilon_{i'j}^{(\nu)}) \prod_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}}^+ \exp(-\beta\Delta\epsilon_{i'j}^{(\nu)})}{\prod_{(i,j) \in \mathcal{B}^{(\mu\nu)}}^+ \exp(-\beta\Delta\epsilon_{i'j}^{(\mu)}) \prod_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}}^+ \exp(-\beta\Delta\epsilon_{i'j}^{(\mu)})}. \quad (\text{B.10})$$

Similar to the above, $+$ denotes products over those boundary pairs (i, j) for which the corresponding energy differences are positive, i.e. for which $\Delta\epsilon_{i'j}^{(\nu)} > 0$ or $\Delta\epsilon_{i'j}^{(\mu)} > 0$. An analogous notation is also used in what follows. By using Eq. (B.4a), the first

fraction in (B.10) can be expressed

$$\prod_{(i,j) \in \mathcal{B}^{(\mu\nu)}}^+ \exp\left(-\beta\Delta\epsilon_{i'j}^{(\nu)}\right) \prod_{(i,j) \in \mathcal{B}^{(\mu\nu)}}^- \exp\left(-\beta\Delta\epsilon_{i'j}^{(\nu)}\right) = \prod_{(i,j) \in \mathcal{B}^{(\mu\nu)}} \exp\left(-\beta\Delta\epsilon_{i'j}^{(\nu)}\right) = \frac{\prod_{(i,j) \in \mathcal{B}^{(\mu\nu)}} \exp\left(-\beta\epsilon_{ij}^{(\mu)}\right)}{\prod_{(i,j) \in \mathcal{B}^{(\mu\nu)}} \exp\left(-\beta\epsilon_{ij}^{(\nu)}\right)}, \quad (\text{B.11})$$

and the second fraction in (B.10) can be expressed as

$$\frac{\prod_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}}^+ \exp\left(-\beta\Delta\epsilon_{i'j}^{(\nu)}\right)}{\prod_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}}^+ \exp\left(-\beta\Delta\epsilon_{i'j}^{(\mu)}\right)} = \frac{\prod_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}}^- \exp\left(\beta\epsilon_{ij}^{(\nu)}\right)}{\prod_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}}^- \exp\left(\beta\epsilon_{ij}^{(\mu)}\right)}. \quad (\text{B.12})$$

Hence

$$\frac{\prod_{(i,j) \in \mathcal{B}^{(\nu)}} Q_{i'j}^{(\nu)}}{\prod_{(i,j) \in \mathcal{B}^{(\mu)}} Q_{i'j}^{(\mu)}} = \frac{\prod_{(i,j) \in \mathcal{B}^{(\mu\nu)}} \exp\left(-\beta\epsilon_{ij}^{(\mu)}\right)}{\prod_{(i,j) \in \mathcal{B}^{(\mu\nu)}} \exp\left(-\beta\epsilon_{ij}^{(\nu)}\right)} \frac{\prod_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}}^- \exp\left(\beta\epsilon_{ij}^{(\nu)}\right)}{\prod_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}}^- \exp\left(\beta\epsilon_{ij}^{(\mu)}\right)}. \quad (\text{B.13})$$

By combining expressions (B.9) and (B.13), the acceptance probability (B.2) reduces to

$$W_{\text{acc}}^{(\mu \rightarrow \nu | R)} = \min \left\{ 1, \frac{\prod_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}}^+ \exp\left(-\beta\epsilon_{ij}^{(\nu)}\right)}{\prod_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}}^+ \exp\left(-\beta\epsilon_{ij}^{(\mu)}\right)} \right\}, \quad (\text{B.14a})$$

provided there is no frustrated link in the boundary $\mathcal{B}^{(\mu)}$. If $\mathcal{B}^{(\mu)}$ contains a frustrated link, the cluster is rejected, i.e.

$$W_{\text{acc}}^{(\mu \rightarrow \nu | R)} = 0. \quad (\text{B.14b})$$

Expressions (B.14) are identical to the original expression for the cluster move acceptance probability [Whitelam et al., 2009a], although the notation here is slightly different. The notation here is more explicit, in that the product in the numerator runs only over those pairs in the boundary of \mathcal{C} that are outside the interaction region in the original state, but end up in the interaction region in the final state. Again,

this set of pairs is denoted by $\mathcal{B}^{(\bar{\mu}\nu)}$. Moreover, the notation also explicitly restricts the product to run only over those pairs in $\mathcal{B}^{(\bar{\mu}\nu)}$ such that $\epsilon_{ij}^{(\nu)} > 0$. Similarly, the product in the denominator runs over boundary pairs (i, j) which are interacting in the original state but not interacting in the final state, and have $\epsilon_{ij}^{(\mu)} > 0$.

Recall that Eq. (B.14) was derived under the assumption of free cluster selection. Let us now restrict the class of pairs to which a link is proposed in the recursive selection of \mathcal{C} . If links are proposed only to pairs interacting in state μ , and $R_{\mathcal{C}}^{(\mu)}$ is the corresponding realization of \mathcal{C} , then it is clear that there exists a realization $R_{\mathcal{C}}^{(\nu)}$ such that $\mathcal{L}^{(\nu)} = \mathcal{L}^{(\mu)}$ and $\mathcal{F}^{(\nu)} = \mathcal{F}^{(\mu)}$; however, one cannot expect $\mathcal{B}^{(\nu)}$ equal to $\mathcal{B}^{(\mu)}$. Since links proposed to pairs $\mathcal{B}^{(\nu)} \setminus \mathcal{B}^{(\mu)}$ or more generally to the corresponding subset of pairs $\mathcal{A} \setminus \mathcal{B}^{(\mu)}$ are doomed to fail by definition (B.1), it is practical not to propose links to those pairs, i.e. to the pairs that do not interact in state μ . We can thus assume that links were proposed to all pairs (i, j) , $i \in \mathcal{C}$, $j \notin \mathcal{C}$, although they were not, and we can still use the acceptance probability (B.14). Note that to evaluate the denominator in Eq. (B.14a), the set of pairs $\mathcal{B}^{(\nu)} \setminus \mathcal{B}^{(\mu)}$ still needs to be detected. Since $\mathcal{B}^{(\nu)} \setminus \mathcal{B}^{(\mu)}$ consists of pairs (i, j) , $i \in \mathcal{C}$, $j \notin \mathcal{C}$ such that (i, j) is not interacting in μ but interacting in μ_i , the computational cost of this algorithm is comparable to the cost of the algorithm generating realizations of clusters with symmetric boundary and asymmetric core (Sec. 2.5) in the limit of high cluster acceptance probability.

Appendix C

Validation of the early rejection schemes

In what follows, we validate the early rejection schemes used in Chap. 2 in a way which is similar to that derived in Frenkel and Smit [2002] for a single-particle move.

C.1 Algorithm of Sec. 2.4

Here, we attempt to clarify why rejection of clusters selected via the free cluster selection, and having frustrated links in the boundary, does not violate sampling from the Boltzmann distribution. Let us firstly explain why clusters with frustrated boundary pairs are rejected. If a boundary pair is frustrated in μ under M , then $p_{i'j}^{(\mu)} > 0$, and Eq. (2.9) implies that $p_{i'j}^{(\nu)} = 0$, which means that (i, j) is necessarily outright failed in state ν under M^{-1} . It is thus impossible to construct the realization of \mathcal{C} in state ν with the same boundary pair, which would be frustrated, and the cluster must be rejected. Now, pair (i, j) is frustrated in μ with probability lower than 1, and it might not be completely obvious why rejection of clusters with frustrated links in the boundary preserves the correct sampling. To demonstrate this more clearly, let us compare the cluster move acceptance probabilities for the early rejection scheme in a way which is analogous to Frenkel and Smit [2002]. It can be easily seen that the cluster move acceptance probability (2.27) is equivalent to the probability

$$W_{\text{acc}}^{(\mu \rightarrow \nu | R)} = \prod_{(i,j) \in \mathcal{B}^{(\mu)}} q_{i'j}^{(\mu)}, \quad (\text{C.1})$$

which can be understood as the probability of not having any frustrated link in the boundary. Similarly, the acceptance probability for the reverse move can be written

as

$$W_{\text{acc}}^{(\mu \rightarrow \nu | R)} = \prod_{(i,j) \in \mathcal{B}^{(\nu)}} q_{i'j}^{(\nu)}. \quad (\text{C.2})$$

Since links between (i, j) $i, j \in \mathcal{C}$ can be seen to be selected to \mathcal{C} with the same probability in states μ and ν , relation (2.26) implies that

$$\frac{W_{\text{acc}}^{(\mu \rightarrow \nu | R)}}{W_{\text{acc}}^{(\nu \rightarrow \mu | R)}} = \exp(-\beta(E_\nu - E_\mu)), \quad (\text{C.3})$$

which is what is expected in a valid early rejection scheme.

C.2 Algorithm of Sec. 2.6

We will show that rejecting the proposed clusters with frustrated links in the boundary, selected by the procedure described in Sec. 2.6, still leads to the correct sampling. We define the probability of accepting a realization $R_{\mathcal{C}}^{(\mu)}$ in a different way than in Eqs. (2.22) and (2.34) as

$$W_{\text{acc}}^{(\mu \rightarrow \nu | R)} = \prod_{(i,j) \in \mathcal{B}^*}^+ \exp(-\beta \Delta \epsilon_{i'j}^{(\mu)}) \min \left\{ 1, \prod_{(i,j) \in \mathcal{B}^\dagger} \exp(-\beta \Delta \epsilon_{i'j}^{(\mu)}) \right\}, \quad (\text{C.4})$$

which is easily seen to be equivalent to Eq. (2.34a). Indeed, the first product determines the probability of generating a boundary formed of outright failed links, and the second term (conditional upon the first) is the acceptance probability, provided \mathcal{B}^* is without frustrated links. By using relation (2.8) we can express similarly the acceptance probability for the reverse move as

$$W_{\text{acc}}^{(\nu \rightarrow \mu | R)} = \prod_{(i,j) \in \mathcal{B}^*}^- \exp(\beta \Delta \epsilon_{i'j}^{(\mu)}) \min \left\{ 1, \prod_{(i,j) \in \mathcal{B}^\dagger} \exp(-\beta \Delta \epsilon_{i'j}^{(\nu)}) \right\}. \quad (\text{C.5})$$

Now, it follows from Eq. (2.9) that

$$\text{if } \prod_{(i,j) \in \mathcal{B}^\dagger} \exp(-\beta \Delta \epsilon_{i'j}^{(\mu)}) \leq 1, \text{ then } \prod_{(i,j) \in \mathcal{B}^\dagger} \exp(-\beta \Delta \epsilon_{i'j}^{(\nu)}) \geq 1, \quad (\text{C.6})$$

$$\text{and if } \prod_{(i,j) \in \mathcal{B}^\dagger} \exp(-\beta \Delta \epsilon_{i'j}^{(\mu)}) > 1, \text{ then } \prod_{(i,j) \in \mathcal{B}^\dagger} \exp(-\beta \Delta \epsilon_{i'j}^{(\nu)}) < 1. \quad (\text{C.7})$$

Let us consider the case (C.6); then Eqs. (C.4), (C.5), (2.29) and (2.31) imply that

$$\begin{aligned} \frac{W_{\text{acc}}^{(\mu \rightarrow \nu | R)}}{W_{\text{acc}}^{(\nu \rightarrow \mu | R)}} &= \frac{\prod_{(i,j) \in \mathcal{B}^*}^+ \exp(-\beta \Delta \epsilon_{i'j}^{(\mu)}) \prod_{(i,j) \in \mathcal{B}^\dagger} \exp(-\beta \Delta \epsilon_{i'j}^{(\mu)})}{\prod_{(i,j) \in \mathcal{B}^*}^- \exp(\beta \Delta \epsilon_{i'j}^{(\mu)})} \\ &= \prod_{(i,j) \in \mathcal{B}} \exp(-\beta \Delta \epsilon_{i'j}^{(\mu)}) = \exp(-\beta(E_\nu - E_\mu)), \end{aligned} \quad (\text{C.8})$$

which is what must be valid in a correct early rejection scheme. One can show the same for the case (C.7). This guarantees that the scheme samples from the Boltzmann distribution.

C.3 Algorithm of Sec. 2.8

We now verify the validity of the early rejection scheme of the algorithm in Sec. 2.8. It can be seen that the cluster move acceptance probability (B.14) is equivalent to the probability

$$W_{\text{acc}}^{(\mu \rightarrow \nu | R)} = \prod_{(i,j) \in \mathcal{B}^{(\mu)}} Q_{i'j}^{(\mu)} \min \left\{ 1, \frac{\prod_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}}^+ \exp(-\beta \epsilon_{ij}^{(\nu)})}{\prod_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}}^+ \exp(-\beta \epsilon_{ij}^{(\mu)})} \right\}, \quad (\text{C.9})$$

which can be understood as the probability of accepting a cluster move, conditional upon not having any frustrated link in the boundary. Similarly, the acceptance probability for the reverse move can be written as

$$W_{\text{acc}}^{(\nu \rightarrow \mu | R)} = \prod_{(i,j) \in \mathcal{B}^{(\nu)}} Q_{i'j}^{(\nu)} \min \left\{ 1, \frac{\prod_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}}^+ \exp(-\beta \epsilon_{ij}^{(\mu)})}{\prod_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}}^+ \exp(-\beta \epsilon_{ij}^{(\nu)})} \right\}. \quad (\text{C.10})$$

By using Eqs. (B.13) and (B.9), and distinguishing between the two cases

$$\begin{aligned} \prod_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}}^+ \exp\left(-\beta\epsilon_{ij}^{(\mu)}\right) &\leq \prod_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}}^+ \exp\left(-\beta\epsilon_{ij}^{(\nu)}\right) , \\ \prod_{(i,j) \in \mathcal{B}^{(\mu\bar{\nu})}}^+ \exp\left(-\beta\epsilon_{ij}^{(\mu)}\right) &> \prod_{(i,j) \in \mathcal{B}^{(\bar{\mu}\nu)}}^+ \exp\left(-\beta\epsilon_{ij}^{(\nu)}\right) , \end{aligned}$$

one can show that

$$\frac{W_{\text{acc}}^{(\mu \rightarrow \nu | R)}}{W_{\text{acc}}^{(\nu \rightarrow \mu | R)}} = \exp\left(-\beta(E_\nu - E_\mu)\right) , \quad (\text{C.11})$$

as desired.

Bibliography

- M. Alimohammadi and K. A. Fichthorn. Molecular dynamics simulation of the aggregation of titanium dioxide nanocrystals: preferential alignment. *Nano Lett.*, 9:4198, 2009.
- M. P. Allen and D. Quigley. Some comments on Monte Carlo and molecular dynamics methods. *Mol. Phys.*, 111:3442, 2013.
- V. J. Anderson and H. N. W. Lekkerkerker. Insights into phase transition kinetics from colloid science. *Nature*, 416:811, 2002.
- A. J. Archer and N. B. Wilding. Phase behavior of a fluid with competing attractive and repulsive interactions. *Phys. Rev. E*, 76:031501, 2007.
- A. J. Archer, D. Pini, R. Evans, and L. Reatto. Model colloidal fluid with competing interactions: Bulk and interfacial properties. *J. Chem. Phys.*, 126:014104, 2007.
- S. Babu, J.-C. Gimel, T. Nicolai, and C. De Michele. The influence of bond rigidity and cluster diffusion on the self-diffusion of hard spheres with square well interaction. *J. Chem. Phys.*, 128:204504, 2008.
- A. S. Barnard and P. Zapol. A model for the phase stability of arbitrary nanoparticles as a function of size and shape. *J. Chem. Phys.*, 121:4276, 2004.
- C. Bechinger, F. Sciortino, and P. Zihlerl. *Physics of Complex Colloids*, volume 184. IOS Press, 2013.
- E. P. Bernard and W. Krauth. Addendum to 'Event-chain Monte Carlo algorithms for hard-sphere systems'. *Phys. Rev. E*, 86:017701, 2012.
- E. P. Bernard, W. Krauth, and D. B. Wilson. Event-chain Monte Carlo algorithms for hard-sphere systems. *Phys. Rev. E*, 80:056704, 2009.
- A. Bhattacharyay and A. Troisi. Self-assembly of sparsely distributed molecules: An efficient cluster algorithm. *Chem. Phys. Lett.*, 458:210, 2008.

- N. S. Bieler, T. P. J. Knowles, D. Frenkel, and R. Vácha. Connecting Macroscopic Observables and Microscopic Assembly Events in Amyloid Formation Using Coarse Grained Simulations. *PLoS Comput. Biol.*, 8:e1002692, 2012.
- T. Biyikoglu and J. Leydold. Faberkrahn type inequalities for trees. *J. Comb. Theory, Ser. B*, 97:159, 2007.
- M. Bjerre, J. M. Tarp, L. Angheluta, and J. Mathiesen. Rotation-induced grain growth and stagnation in phase-field crystal models. *Phys. Rev. E*, 88:020401, 2013.
- J.-M. Bomont, J.-L. Bretonnet, and D. Costa. Temperature study of cluster formation in two-yukawa fluids. *J. Chem. Phys.*, 132:184508, 2010.
- J. M. Bomont, J. L. Bretonnet, D. Costa, and J. P. Hansen. Communication: Thermodynamic signatures of cluster formation in fluids with competing interactions. *J. Chem. Phys.*, 137:011101, 2012.
- M. Broccio, D. Costa, Y. Liu, and S.-H. Chen. The structural properties of a two-yukawa fluid: Simulation and analytical results. *J. Chem. Phys.*, 124:084501, 2006.
- M. Camargo, S. A. Egorov, and C. N. Likos. Cluster formation in star-linear polymer mixtures: equilibrium and dynamical properties. *Soft Matter*, 8:4177, 2012.
- F. Cardinaux, A. Stradner, P. Schurtenberger, F. Sciortino, and E. Zaccarelli. Modeling equilibrium clusters in lysozyme solutions. *Europhys. Lett.*, 77:48004, 2007.
- M. E. Cates, M. Fuchs, K. Kroy, W. C. K. Poon, and A. M. Puertas. Theory and simulation of gelation, arrest and yielding in attracting colloids. *J. Phys.: Condens. Matter*, 16:S4861, 2004.
- P. Charbonneau and D. R. Reichman. Systematic characterization of thermodynamic and dynamical phase behavior in systems with short-ranged attraction. *Phys. Rev. E*, 75:011507, 2007a.
- P. Charbonneau and D. R. Reichman. Phase behavior and far-from-equilibrium gelation in charged attractive colloids. *Phys. Rev. E*, 75:050401, 2007b.
- B. Chen and J. I. Siepmann. A novel Monte Carlo algorithm for simulating strongly associating fluids: Applications to water, hydrogen fluoride, and acetic acid. *J. Phys. Chem. B*, 104:8725, 2000.

- B. Chen, K. Lutker, J. Lei, J. Yan, S. Yang, and H. Mao. Detecting grain rotation at the nanoscale. *Proc. Natl. Acad. Sci. (USA)*, 111:3350, 2014.
- P. Clancy. Nanoparticles: Self-assembly finds its own limits. *Nat. Nanotechnol.*, 6: 540, 2011.
- H. Cölfen and M. Antonietti. *Mesocrystals and Nonclassical Crystallization*. Wiley, 2008.
- C. J. Dalmaschio, C. Ribeiro, and E. R. Leite. Impact of the colloidal state on the oriented attachment growth mechanism. *Nanoscale*, 2:2336, 2010.
- P. F. Damasceno, M. Engel, and S. C. Glotzer. Predictive self-assembly of polyhedra into complex structures. *Science*, 337:453, 2012.
- A. de Candia, E. Del Gado, A. Fierro, N. Sator, M. Tarzia, and A. Coniglio. Columnar and lamellar phases in attractive colloidal systems. *Phys. Rev. E*, 74:010403, 2006.
- C. De Michele. Simulating hard rigid bodies. *J. Comput. Phys.*, 229:3276, 2010.
- R. Demichelis, P. Raiteri, J. D. Gale, D. Quigley, and D. Gebauer. Stable prenucleation mineral clusters are liquid-like ionic polymers. *Nat. Commun.*, 2:590, 2011.
- A. R. Denton. Charge renormalization, effective interactions, and thermodynamics of deionized colloidal suspensions. *J. Phys.: Condens. Matter*, 20:494230, 2008.
- A. D. Dinsmore, P. L. Dubin, G. M. Grason, D. J. Kraft, J. Hilhorst, M. A. P. Heinen, M. J. Hoogenraad, B. Luigjes, W. K. Kegel, N. Meilhac, et al. Clustering in complex fluids. *J. Phys. Chem. B*, 115:7173, 2011.
- K. A. Fichthorn. Molecular phenomena in colloidal nanostructure synthesis. *Mol. Simulat.*, 40:134, 2014.
- G. Foffi, K. A. Dawson, S. V. Buldyrev, F. Sciortino, E. Zaccarelli, and P. Tartaglia. Evidence for an unusual dynamical-arrest scenario in short-ranged colloidal systems. *Phys. Rev. E*, 65:050802, 2002.
- G. Foffi, C. De Michele, F. Sciortino, and P. Tartaglia. Arrested phase separation in a short-ranged attractive colloidal system: A numerical study. *J. Chem. Phys.*, 122:224903, 2005a.

- G. Foffi, C.D. Michele, F. Sciortino, and P. Tartaglia. Scaling of dynamics with the range of interaction in short-range attractive colloids. *Phys. Rev. Lett.*, 94:78301, 2005b.
- D. Frenkel. Speed-up of Monte Carlo simulations by sampling of rejected states. *Proc. Natl. Acad. Sci. (USA)*, 101:17571, 2004.
- D. Frenkel and B. Smit. *Understanding Molecular Simulation: From Algorithms to Applications*. Academic Pr, 2nd ed., 2002.
- A. Furukawa and H. Tanaka. Key role of hydrodynamic interactions in colloidal gelation. *Phys. Rev. Lett.*, 104:245702, 2010.
- D. Fusco and P. Charbonneau. Crystallization of asymmetric patchy models for globular proteins in solution. *Phys. Rev. E*, 88:012721, 2013.
- M. Godinho, C. Ribeiro, E. Longo, and E. R. Leite. Influence of microwave heating on the growth of gadolinium-doped cerium oxide nanorods. *Cryst. Growth Des.*, 8:384, 2008.
- J. Grant and R. L. Jack. Quantifying reversibility in a phase-separating lattice gas: An analogy with self-assembly. *Phys. Rev. E*, 85, 2012.
- J. Grant, R. L. Jack, and S. Whitlam. Analyzing mechanisms and microscopic reversibility of self-assembly. *J. Chem. Phys.*, 135:214505, 2011.
- J. Groenewold and W. K. Kegel. Anomalous large equilibrium clusters of colloids. *J. Phys. Chem. B*, 105:11702, 2001.
- J. H. Harding, D. M. Duffy, M. L. Sushko, P. M. Rodger, D. Quigley, and J. A. Elliott. Computational techniques at the organic-inorganic interface in biomineralization. *Chem. Rev.*, 108:4823, 2008.
- T. K. Haxton and S. Whitlam. Design rules for the self-assembly of a protein crystal. *Soft Matter*, 8:3558, 2012.
- L. O. Hedges and S. Whitlam. Limit of validity of Ostwald’s rule of stages in a statistical mechanical model of crystallization. *J. Chem. Phys.*, 135:164602, 2011.
- M. R. Henzinger. Fully dynamic biconnectivity in graphs. *Algorithmica*, 13:503, 1995.
- R. L. Jack, M. F. Hagan, and D. Chandler. Fluctuation-dissipation ratios in the dynamics of self-assembly. *Phys. Rev. E*, 76:021119, 2007.

- P. K. Jha, V. Kuzovkov, B. A. Grzybowski, and M. O. de la Cruz. Dynamic self-assembly of photo-switchable nanoparticles. *Soft Matter*, 8:227, 2012.
- J. Kaupužs, J. Rimšāns, and R. V. N. Melnik. Parallelization of the Wolff single-cluster algorithm. *Phys. Rev. E*, 81:026701, 2010.
- S. J. Khan, C. M. Sorensen, and A. Chakrabarti. Kinetics and morphology of cluster growth in a model of short-range attractive colloids. *J. Chem. Phys.*, 131:194908, 2009.
- C. E. Killian, R. A. Metzler, Y. U. T. Gong, I. C. Olson, J. Aizenberg, Y. Politi, F. H. Wilt, A. Scholl, A. Young, A. Doran, et al. Mechanism of calcite co-orientation in the sea urchin tooth. *J. Am. Chem. Soc.*, 131:18404, 2009.
- Y. Y. Kim, K. Ganesan, P. Yang, A. N. Kulak, S. Borukhin, S. Pechook, L. Ribeiro, R. Kröger, S. J. Eichhorn, S. P. Armes, et al. An artificial biomineral formed by incorporation of copolymer micelles in calcite crystals. *Nat. Mater.*, 10:890, 2011.
- D. Klotsa and R. L. Jack. Predicting the self-assembly of a model colloidal crystal. *Soft Matter*, 7:6294, 2011.
- D. Klotsa and R. L. Jack. Controlling crystal self-assembly using a real-time feedback scheme. *J. Chem. Phys.*, 138:094502, 2013.
- D. J. Kraft, R. Ni, F. Smalenburg, M. Hermes, K. Yoon, D. A. Weitz, A. van Blaaderen, J. Groenewold, M. Dijkstra, and W. K. Kegel. Surface roughness directed self-assembly of patchy particles into colloidal micelles. *Proc. Natl. Acad. Sci. (USA)*, 109:10787, 2012.
- W. Krauth. *Algorithms and Computations*. Oxford University Press, 2006.
- K. Kroy, M.E. Cates, and W.C.K. Poon. Cluster mode-coupling approach to weak gelation in attractive colloids. *Phys. Rev. Lett.*, 92:148302, 2004.
- D. P. Landau and K. Binder. *Guide To Monte Carlo Simulations In Statistical Physics*. Cambridge Univ Pr, 2009.
- E. R. Leite and C. Ribeiro. *Crystallization and Growth of Colloidal Nanocrystals*. Springer, 2012.
- C. N. Likos. Colloidal interactions: From effective potentials to structure. *Physics of Complex Colloids*, 184:1, 2013.

- C. N. Likos, H. Löwen, M. Watzlawek, B. Abbas, O. Jucknischke, J. Allgaier, and D. Richter. Star polymers viewed as ultrasoft colloidal particles. *Phys. Rev. Lett.*, 80:4450, 1998.
- J. W. Liu and E. Luijten. Rejection-free geometric cluster algorithm for complex fluids. *Phys. Rev. Lett.*, 92:035504, 2004.
- J. W. Liu and E. Luijten. Generalized geometric cluster algorithm for fluid simulation. *Phys. Rev. E*, 71:066701, 2005.
- A.A. Louis, E. Allahyarov, H. Löwen, and R. Roth. Effective forces in colloidal mixtures: From depletion attraction to accumulation repulsion. *Phys. Rev. E*, 65:061407, 2002.
- P. J. Lu, J. C. Conrad, H. M. Wyss, A. B. Schofield, and D. A. Weitz. Fluids of clusters in attractive colloids. *Phys. Rev. Lett.*, 96:28306, 2006.
- P. J. Lu, E. Zaccarelli, F. Ciulla, A. B. Schofield, F. Sciortino, and D. A. Weitz. Gelation of particles with short-range attraction. *Nature*, 453:499, 2008.
- A. Malins, S. R. Williams, J. Eggers, H. Tanaka, and C. P. Royall. The effect of inter-cluster interactions on the structure of colloidal clusters. *J. Non-Cryst. Solids*, 357:760, 2011.
- E. Mani, W. Lechner, W. K. Kegel, and P. G. Bolhuis. Equilibrium and non-equilibrium cluster phases in colloids with competing interactions. *Soft Matter*, 10:4479, 2014.
- D. Moldovan, V. Yamakov, D. Wolf, and S. R. Phillpot. Scaling behavior of grain-rotation-induced grain growth. *Phys. Rev. Lett.*, 89:206101, 2002.
- L. J. Moore, R. D. Dear, M. D. Summers, R. P. A. Dullens, and G. A. D. Ritchie. Direct observation of grain rotation-induced grain coalescence in two-dimensional colloidal crystals. *Nano Lett.*, 10:4266, 2010.
- S. Mossa, G. Ruocco, F. Sciortino, and P. Tartaglia. Quenches and crunches: does the system explore in ageing the same part of the configuration space explored in equilibrium? *Philos. Mag. B*, 82:695, 2002.
- S. Mossa, F. Sciortino, P. Tartaglia, and E. Zaccarelli. Ground-state clusters for short-range attractive and long-range repulsive potentials. *Langmuir*, 20:10756, 2004.

- F. Niedermayer. General cluster updating method for Monte Carlo simulations. *Phys. Rev. Lett.*, 61:2026, 1988.
- M. G. Noro and D. Frenkel. Extended corresponding-states behavior for particles with variable range attractions. *J. Chem. Phys.*, 113:2941, 2000.
- G. Odriozola. Revisiting the phase diagram of hard ellipsoids. *J. Chem. Phys.*, 136:134505, 2012.
- W. Z. Ostwald. Studies on formation and transformation of solid materials. *Z. Phys. Chem.*, 22:289, 1897.
- T. E. Ouldridge, A. A. Louis, and J. P. K. Doye. DNA nanotweezers studied with a coarse-grained model of DNA. *Phys. Rev. Lett.*, 104:178101, 2010.
- E. Peters et al. Rejection-free Monte Carlo sampling for general potentials. *Phys. Rev. E*, 85:026703, 2012.
- T. Perez, Y. Liu, W. Li, J. D. Gunton, and A. Chakrabarti. Pathways of cluster growth and kinetic slowing down in a model of short-range attractive colloids. *Langmuir*, 27:11401, 2011.
- D. Pini, G. Jialin, A. Parola, and L. Reatto. Enhanced density fluctuations in fluid systems with competing interactions. *Chem. Phys. Lett.*, 327:209, 2000.
- F. Romano, C. De Michele, D. Marenduzzo, and E. Sanz. Monte Carlo and event-driven dynamics of Brownian particles with orientational degrees of freedom. *J. Chem. Phys.*, 135:124106, 2011.
- F. Romano, A. Hudson, J. P. K. Doye, T. E. Ouldridge, and A. A. Louis. The effect of topology on the structure and free energy landscape of DNA kissing complexes. *J. Chem. Phys.*, 136:215102, 2012.
- P. W. K. Rothmund. Folding dna to create nanoscale shapes and patterns. *Nature*, 440:297, 2006.
- C. P. Royall and A. Malins. The role of quench rate in colloidal gels. *Faraday Discuss.*, 158:301, 2012.
- C. P. Royall, M. E. Leunissen, and A. Van Blaaderen. A new colloidal model system to study long-range interactions quantitatively in real space. *J. Phys.: Condens. Matter*, 15:S3581, 2003.

- Š. Ružička and M. P. Allen. Collective Translational and Rotational Monte Carlo Move for Attractive Particles. *Phys. Rev. E*, 89:033307, 2014.
- J. Russo, P. Tartaglia, and F. Sciortino. Association of limited valence patchy particles in two dimensions. *Soft Matter*, 6:4229, 2010.
- E. Sanz and D. Marenduzzo. Dynamic Monte Carlo versus Brownian dynamics: A comparison for self-diffusion and crystallization in colloidal fluids. *J. Chem. Phys.*, 132:194102, 2010.
- S. Sastry. Liquid limits: Glass transition and liquid-gas spinodal boundaries of metastable liquids. *Phys. Rev. Lett.*, 85:590, 2000.
- T. Schilling, H. J. Schöpe, M. Oettel, G. Opletal, and I. Snook. Precursor-mediated crystallization process in suspensions of hard spheres. *Phys. Rev. Lett.*, 105:25701, 2010.
- F. Sciortino. Primitive models of patchy colloidal particles. A review. *Collect. Czech. Chem. Commun.*, 75:349, 2010.
- F. Sciortino, S. Mossa, E. Zaccarelli, and P. Tartaglia. Equilibrium cluster phases and low-density arrested disordered states: the role of short-range attraction and long-range repulsion. *Phys. Rev. Lett.*, 93:55701, 2004.
- F. Sciortino, P. Tartaglia, and E. Zaccarelli. One-dimensional cluster growth and branching gels in colloidal systems with short-range depletion attraction and screened electrostatic repulsion. *J. Phys. Chem. B*, 109:21942, 2005.
- R. P. Sear and W. M. Gelbart. Microphase separation versus the vapor-liquid transition in systems of spherical particles. *J. Chem. Phys.*, 110:4582, 1999.
- J. Seto, Y. Ma, S. A. Davis, F. Meldrum, A. Gourrier, Y. Y. Kim, U. Schilde, M. Sztucki, M. Burghammer, S. Maltsev, et al. Structure-property relationships of a biological mesocrystal in the adult sea urchin spine. *Proc. Natl. Acad. Sci. (USA)*, 109:3699, 2012.
- D.W. Sinkovits, S.A. Barr, and E. Luijten. Rejection-free Monte Carlo scheme for anisotropic particles. *J. Chem. Phys.*, 136:144111, 2012.
- F. Smalenburg, L. Filion, M. Marechal, and M. Dijkstra. Vacancy-stabilized crystalline order in hard cubes. *Proc. Natl. Acad. Sci. (USA)*, 109:17886, 2012.

- A. G. Stack, P. Raiteri, and J. D. Gale. Accurate rates of the complex mechanisms for growth and dissolution of minerals using a combination of rare-event theories. *J. Am. Chem. Soc.*, 134:11, 2011.
- P. J. Steinhardt, D. R. Nelson, and M. Ronchetti. Bond-orientational order in liquids and glasses. *Phys. Rev. B*, 28:784, 1983.
- E. Stiakakis, G. Petekidis, D. Vlassopoulos, C. N. Likos, H. Iatrou, N. Hadjichristidis, and J. Roovers. Depletion and cluster formation in soft colloid-polymer mixtures. *Europhys. Lett.*, 72:664, 2005.
- A. Stradner, H. Sedgwick, F. Cardinaux, W. C. K. Poon, S. U. Egelhaaf, and P. Schurtenberger. Equilibrium cluster formation in concentrated protein solutions and colloids. *Nature*, 432:492, 2004.
- D. Stranski and Z. Totomanov. Rate of formation of (crystal) nuclei and the Ostwald step rule. *Z. Phys. Chem. Abt. A*, 163:399, 1933.
- P. Šulc, F. Romano, T. E. Ouldridge, L. Rovigatti, J. P. K. Doye, and A. A. Louis. Sequence-dependent thermodynamics of a coarse-grained DNA model. *J. Chem. Phys.*, 137:135101, 2012.
- M. B. Sweatman, R. Fartaria, and L. Lue. Cluster formation in fluids with competing short-range and long-range interactions. *J. Chem. Phys.*, 140:124508, 2014.
- R. H. Swendsen. How the maximum step size in Monte Carlo simulations should be adjusted. *Phys. Procedia*, 15:81, 2011.
- R. H. Swendsen and J. S. Wang. Nonuniversal Critical-Dynamics in Monte-Carlo Simulations. *Phys. Rev. Lett.*, 58:86, 1987.
- Z. Tang, Z. Zhang, Y. Wang, S. C. Glotzer, and N. A. Kotov. Self-assembly of CdTe nanocrystals into free-floating sheets. *Science*, 314:274, 2006.
- P. R. ten Wolde and D. Frenkel. Enhancement of protein crystal nucleation by critical density fluctuations. *Science*, 277:1975, 1997.
- P. R. ten Wolde, M. J. Ruiz-Montero, and D. Frenkel. Numerical calculation of the rate of crystal nucleation in a lennard-jones system at moderate undercooling. *J. Chem. Phys.*, 104:9932, 1996.
- V. Testard, L. Berthier, and W. Kob. Influence of the glass transition on the liquid-gas spinodal decomposition. *Phys. Rev. Lett.*, 106:125702, 2011.

- V. Testard, L. Berthier, and W. Kob. Intermittent dynamics and logarithmic domain growth during the spinodal decomposition of a glass-forming liquid. *arXiv:1309.1587*, 2013.
- J. C. F. Toledano, F. Sciortino, and E. Zaccarelli. Colloidal systems with competing interactions: from an arrested repulsive cluster phase to a gel. *Soft Matter*, 5: 2390, 2009.
- A. Troisi, V. Wong, and M. A. Ratner. Self-assembly on multiple length scales: A Monte Carlo algorithm with data augmentation. *J. Chem. Phys.*, 122:024102, 2005.
- G. A. Vliegenthart, J. F. M. Lodge, and H. N. W. Lekkerkerker. Strong weak and metastable liquids structural and dynamical aspects of the liquid state. *Phys. A*, 263:378, 1999.
- D. Wales. *Energy landscapes: Applications to clusters, biomolecules and glasses*. Cambridge University Press, 2003.
- A. F. Wallace, L. O. Hedges, A. Fernandez-Martinez, P. Raiteri, J. D. Gale, G. A. Waychunas, S. Whitlam, Jillian F. Banfield, and J. J. De Yoreo. Microscopic evidence for liquid-liquid separation in supersaturated CaCO_3 solutions. *Science*, 341:885, 2013.
- R. H. Walmsley. The exact canonical equation of state for a finite hard rod gas with square-well attractive interactions. *J. Chem. Phys.*, 88:4473, 1988.
- S. Whitlam. *private communication*.
- S. Whitlam. Control of pathways and yields of protein crystallization through the interplay of nonspecific and specific attractions. *Phys. Rev. Lett.*, 105:088102, 2010.
- S. Whitlam. Approximating the dynamical evolution of systems of strongly interacting overdamped particles. *Mol. Simulat.*, 37:606, 2011.
- S. Whitlam and P. L. Geissler. Avoiding unphysical kinetic traps in Monte Carlo simulations of strongly attractive particles. *J. Chem. Phys.*, 127:154101, 2007.
- S. Whitlam and P. L. Geissler. Avoiding unphysical kinetic traps in Monte Carlo simulations of strongly attractive particles (vol 127, art no 154101, 2007). *J. Chem. Phys.*, 128:219901, 2008.

- S. Whitelam, E. H. Feng, M. F. Hagan, and P. L. Geissler. The role of collective motion in examples of coarsening and self-assembly. *Soft Matter*, 5:1251, 2009a.
- S. Whitelam, C. Rogers, A. Pasqua, C. Paavola, J. Trent, and P. L. Geissler. The impact of conformational fluctuations on self-assembly: Cooperative aggregation of archaeal chaperonin proteins. *Nano Lett.*, 9:292, 2009b.
- S. Whitelam, I. Tamblyn, P. H. Beton, and J. P. Garrahan. Random and ordered phases of off-lattice rhombus tiles. *Phys. Rev. Lett.*, 108:035702, 2012.
- U. Wolff. Collective Monte-Carlo updating for spin systems. *Phys. Rev. Lett.*, 62:361, 1989.
- L. B. Wright and T. R. Walsh. Efficient conformational sampling of peptides adsorbed onto inorganic surfaces: insights from a quartz binding peptide. *Phys. Chem. Chem. Phys.*, 15:4715, 2013.
- D. Wu, D. Chandler, and B. Smit. Electrostatic analogy for surfactant assemblies. *J. Phys. Chem.*, 96:4077, 1992.
- Y. Xia, T. D. Nguyen, M. Yang, B. Lee, A. Santos, P. Podsiadlo, Z. Tang, S. C. Glotzer, and N. A. Kotov. Self-assembly of self-limiting monodisperse supraparticles from polydisperse nanoparticles. *Nat. Nanotechnol.*, 6:580, 2011.
- V. M. Yuwono, N. D. Burrows, J. A. Soltis, and R. L. Penn. Oriented aggregation: formation and transformation of mesocrystal intermediates revealed. *J. Am. Chem. Soc.*, 132:2163, 2010.
- E. Zaccarelli. Colloidal gels: equilibrium and non-equilibrium routes. *J. Phys.: Condens. Matter*, 19:323101, 2007.
- E. Zaccarelli and W. C. K. Poon. Colloidal glasses and gels: The interplay of bonding and caging. *Proc. Natl. Acad. Sci. (USA)*, 106:15203, 2009.
- I. Zhang, C. P. Royall, M. A. Faers, and P. Bartlett. Phase separation dynamics in colloid-polymer mixtures: the effect of interaction range. *Soft Matter*, 9:2076, 2013.
- T. H. Zhang, J. Klok, R. H. Tromp, J. Groenewold, and W. K. Kegel. Non-equilibrium cluster states in colloids with competing interactions. *Soft Matter*, 8:667, 2012.

- Z. Zhang and S. C. Glotzer. Self-assembly of patchy particles. *Nano Lett.*, 4:1407, 2004.
- L. Zhou and P. O'Brien. Mesocrystals – properties and applications. *J. Phys. Chem. Lett.*, 3:620, 2012.