

CRYSTALLISATION ON A SPHERE

**Computational studies of two-dimensional
Lennard-Jones systems**

Jeroen M. Voogd

CRYSTALLISATION ON A SPHERE

Computational studies of two-dimensional Lennard-Jones systems

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam,
op gezag van de Rector Magnificus
prof. dr. J.J.M. Franse

ten overstaan van een door het College voor Promoties ingestelde
commissie in het openbaar te verdedigen in de Aula der Universiteit
op dinsdag 16 juni 1998 te 13.00 uur

door

Jeroen Marco Voogd

geboren te Heemskerk


promotor: Prof. dr. P.M.A. Sloot

co-promotor: Dr. R. van Dantzig

Commissie: Prof. dr. D. Frenkel
Prof. dr. L.O. Hertzberger
Prof. dr. P.V. Giaquinta
Prof. dr. M. Livny
Dr. J.P.J. Michels

Faculteit: Wiskunde, Informatica, Natuurkunde en Sterrenkunde
Kruislaan 403
1098 SJ Amsterdam
Nederland

This work is part of the research programme of the 'Stichting voor Fundamenteel Onderzoek der Materie (FOM)', which is financially supported by the 'Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO)'.



Advanced School for Computing and Imaging

This work was carried out in graduate school ASCI.
ASCI dissertation series number 32.

Part of this research is funded by the 'Stichting Nationale Computer Faciliteiten (NCF)' under number CRG-91-08.

ISBN 90-5776-004-5

Printed at PrintPartners Ipskamp, Enschede, The Netherlands.

Cover page: The fifth frequency breakdown (2562 particles) of a spherical icosahedron projected onto a hyperbolic surface. The shown Voronoi polygons of the particles are coloured according to the Lennard-Jones potential energy: bright is high energy, dark is low energy.

Contents

1	Introduction	7
2	Theoretical background and related work	13
2.1	Crystallisation	13
2.2	Potentials	16
2.3	Boundary conditions	18
2.4	Simulation methods	21
2.5	Ensembles	22
2.6	Analysis metrics and methods	22
2.6.1	Energy and specific heat	22
2.6.2	Pressure	23
2.6.3	Defect analysis	23
2.6.4	Symmetry analysis	25
2.6.5	Order analysis	27
2.7	Temperature $T > 0$	31
2.7.1	Thermodynamics	31
2.7.2	Melting	32
2.7.3	Freezing	33
2.8	Temperature $T = 0$	34
2.8.1	Many local minima	34
2.8.2	Built-in symmetry	35
2.9	Related work	36
2.9.1	Flat 2D space	37
2.9.2	Spherical 2D space	39
3	Computational methods	43
3.1	Optimisation methods for crystallisation	43
3.2	The choice of simulation methods	45
3.3	The simulated annealing method	51
3.4	The Metropolis Monte Carlo method	52
3.5	Implementation aspects	54
3.6	Tuning and comparison	58

4	Parallel algorithms	69
4.1	Possibilities for parallelism	69
4.2	Parallelisation of the simulation code	72
4.3	Time complexities	78
4.3.1	Sequential SA	79
4.3.2	Systolic SA	80
4.3.3	Data parallel Metropolis	82
4.3.4	Hybrid SA	83
4.3.5	Vectorised SA	85
4.4	Comparison of implementations	86
4.5	Future expectations	89
4.6	Conclusions	90
5	Crystallisation on a sphere ($T>0$)	93
5.1	Introduction	93
5.2	Identification of the phase transition	94
5.3	Positional order	98
5.4	Bond orientational order	102
5.5	Disclination ordering	108
5.6	Disclination clustering	112
6	Crystalline arrangements ($T=0$)	117
6.1	Introduction	117
6.2	From flat to curved 2D space	118
6.3	Global potential energy minima	135
6.3.1	Energy as a function of N	136
6.3.2	The radius as a function of N	140
6.3.3	Number and distribution of d-charges	141
6.3.4	The symmetry of the configurations	146
6.3.5	Enantiometric configurations	150
6.3.6	Dipole moment	152
6.4	Global and local energy minima	154
6.5	Arrangements under compression	162
6.6	Icosahedral arrangements	166
7	Discussion and conclusions	169
	Summary / Samenvatting	175
	Nawoord	183
A	Energy minima for $2 \leq N \leq 200$	185
B	Minima with built-in icosahedral symmetry	191

Chapter 1

Introduction

More than two decades ago W.S. Bont *et al.* [1] from the Netherlands Cancer Institute (NKI) performed measurements to determine the size distribution of vesicles, spherical bilayers consisting of lipid molecules and proteins, spontaneously formed *in vitro* from fragmented biomembrane material. Remarkably, in these experiments discrete vesicle sizes were observed. Even more remarkable was that regularity in the size distribution could be observed, indicating two geometrical series of vesicle sizes.

Discrete size distributions are also observed for other spherical systems in nature. Some types of viruses have a spherical capsid protecting their interior. This coat is a closely packed highly symmetric arrangement of protein structures, forming a spherical shell. Only specific sizes for the spherical viruses are observed. Another example is the new fullerene family of carbon compounds. The best known of these is the buckyball, a highly symmetric spherical cage structure built from sixty carbon atoms. Only specific numbers of carbon atoms form stable cage like structures.

These systems are formed in a self-assembly process. Self-assembly of closely packed spherical shells from elementary subunits is accompanied by the formation of complex geometrical structures with regular polyhedral symmetry. Given the molecular subunits, energy minimisation can force all emerging spherical shells to prefer identical size and shape. During the self-assembly process, the structure is in thermodynamic equilibrium with its environment from which it can take up subunits. Subunits that find a stable position in the structure form the basis for further assembly. The stable positions correspond locally to lowest energy states. During and after formation of the spherical shell, equilibration and energy minimisation can lead to rearrangements of the subunits, and even can change the number of subunits by expelling or taking up material. During this self-assembly process all subunits individually tend to search—until being stuck—lowest energy positions.

As a model for the formation of spherical self-assembly arrangements we can use a crystallisation process in two dimensions (2D) where on a sphere lowest

energy structures of N identical particles are formed. During crystallisation free energy minimisation determines which arrangements are formed. Here also the individual particles search for lowest energy positions consistent with their local constraints. The overall minimisation of free energy tends to lead to global order in the distribution of particles, resulting in a stable arrangement at zero temperature. Ideally, the particles obtain as much as possible identical lowest energy environments, leading to a high degree of symmetry. In nature low energy structures are often favoured. This might explain why in some cases symmetric assemblies occur more frequently than others [2]. Pushing to the extreme the relation between minimum energy arrangements and the tendency towards high symmetry, one may say that symmetry drives the pattern formation.

The geodesic connecting lines between neighbours on a spherical 2D system of N particles give a spherical polyhedral net with N vertices. On a spherical surface, perfectly regular polyhedral nets are only possible for specific numbers of particles; those corresponding to the five regular solids of Plato. Four particles can form the regular tetrahedron, six the regular octahedron with as dual eight particles forming the regular cube, twelve the regular icosahedron with as dual twenty particles forming the regular dodecahedron. For other N -values it is not possible to form perfectly symmetrical arrangements of particles distributed over the sphere. Beyond $N = 20$ the distribution can not be perfectly regular, but its degree of regularity is expected to show fluctuations based on the high symmetry of Plato's solids, which can be related to local minima in the mean (free) energy. If, for example, all 20 triangular faces of the icosahedron mapped on a sphere are broken down into 4 equal triangles, the "first frequency breakdown" [3] of an icosahedron, a structure with larger N can be formed with overall icosahedral symmetry (see cover for a fifth frequency breakdown configuration).

On a flat surface infinitely large arrangements can exist with perfect symmetry; all particles can have equal environments. On a sphere, as said, only Plato's solids with finite numbers of particles are perfectly regular. Unlike flat 2D geometry a sufficiently curved geometry imposes "defects" in a particle arrangement. The emergence of defects can be understood if one considers the geometrical mismatch that is introduced if a flat regular structure is made curved. Identical particles with isotropic interactions in flat 2D space strive towards a triangular lattice arrangement where each particle has six nearest neighbours. In a flat geometry arrangements are completely driven by local symmetry. The global symmetry is a direct consequence of straightforward replication of the local symmetry.

If an originally flat surface with an energetically optimal lattice is being curved, isotropy is lost. In order to regain isotropy, rearrangement of the particles must take place. The geometrical constraints lead to a frustrated system where the processes driving the global symmetry and the local symmetry are mutually interacting. The rearrangement causes some particles to have more or

less than six nearest neighbours. These particles are identified as topological defects or more precisely: disclinations. In particle distributions over the surface of a sphere a net amount of twelve disclinations with five instead of six nearest neighbours is induced. Local symmetry favours six-fold symmetry whereas the curvature induces disclinations with five-fold local symmetry. However, this does not imply at all that such a defected spherical 2D crystal is disordered. The deviation from a triangular lattice around a single disclination introduces strain. With several disclinations the system tends to minimise the overall strain by spreading the disclinations evenly. The disclinations try to settle as far apart as possible from each other, they effectively repel each other with a long range force generated by the lattice. On the surface of a sphere one may expect that the disclinations themselves form ordered structures at sufficiently low temperatures, leading to disclination arrangements such as in Plato's solids. In this way hierarchical ordering can arise.

The minimum energy arrangements are sensitive to the symmetry and shape of the interaction. This is found in calculations with various interactions. For example, J.J. Thompson [4] was the first to study global minimum energy (GEM) arrangements of N equally charged particles confined to the surface of a sphere in order to explain charge distributions in atoms. Another example is the problem of Tammes [5] where the minimum distance between any pair of the N particles is maximised. Tammes needed the solutions of this problem for his study on the distribution of holes in spherical pollen grains.

R. Buckminster Fuller, aiming at efficient design of geodesic domes [6], advocates that structural integrity is obtained if the structure consists entirely of triangles, resulting in polygonal pattern stabilisation. He applies the principle of “tensegrity” [6], where tension and compression are balanced. Distortions are spread over the entire surface such that only excessive stress can cause structural damage.

In all cases the emerging regularities tend to be related to the symmetries of the convex regular polyhedra, i.e. Plato's solids. The principles of induction as well as the breaking of such symmetries is still an open problem.

Several authors [7, 8, 9, 10] tried to find analytical methods to prove for arbitrary N which arrangements have the lowest energy. For $N \leq 20$ this was reasonably successful. For larger N , however, only for a few isolated cases where highly symmetric arrangements are formed, analytical methods could be used. In the past decades the problem of distributing particles optimally—given their interaction—on the surface of a sphere, received new attention due to the availability of fast computers. With the aid of computer simulations, minimum energy arrangements can be generated on a trial and error basis, although it can usually not be proven that a truly optimal solution has been reached.

In the physics literature various types of interactions are used in energy minimisation calculations and simulation experiments. Examples are the Coulomb

interaction, the logarithmic interaction, and interactions mimicking that of the carbon atom. Other optimisation problems include finding Tammes distributions and the covering of a spherical surface with as few (overlapping) circles as possible.

In this thesis we want to construct a model to study the properties of spherical particle distributions related to naturally occurring systems. Studies reported in literature deal mostly with interaction potentials that are purely repulsive. Then the radius of the sphere has to be fixed, or external pressure must be included in the model.

Our model is constructed [11] with Bont's size measurements of vesicles in mind. The basic substrate of vesicles is assembled from the lipid component. Lipid molecules have a hydrophilic headgroup and a hydrophobic tail section. Due to the properties of lipid molecules they spontaneously form bilayers that hide the hydrophobic tails from the aqueous environment except at the edge. Energetically it is then favourable for those bilayers to close upon themselves and form spherical bilayer shells; the vesicles. Including all the interactions present in a lipid molecule and determining the minimum energy configuration for N up to at least $\mathcal{O}(10^3)$ would result in a computationally too demanding problem. Instead we simplify the model while attempting to maintain the relevant features.

In the inner layer the lipid headgroups have the highest packing constraints and the most stable arrangement of those headgroups might determine the overall size (number of lipid headgroups, N) of the vesicle. Therefore we only model the inner headgroup layer of the vesicle and model the headgroups of the molecules as centrally interacting particles. In water lipid molecules separated by a small distance have—largely due to the hydrophobic interaction—an effectively attractive potential. If the lipid molecules are close together, however, steric hindrance results in a repulsive interaction. In our model the interaction is taken as the Lennard-Jones (LJ) interaction which approximates the required general shape. Although the particles are constrained to a spherical surface, the LJ interaction between particles is evaluated for interparticle distances (cords) through the sphere.

The model thus constructed has two parameters: the number of particles (N) and the curvature fixed by the radius of the sphere (R).

By applying computer simulations for our model, in general we want to learn how crystallisation emerges in two dimensions under spherical constraints and to study the stability of the found crystalline arrangements. Especially we want to determine whether some N values can form more stable structures than others, and which properties are connected to this stability. The disclinations in a particle distribution on a spherical surface have a different behaviour than on a flat surface. In this work we also want to explore how these disclinations are incorporated into the spherical arrangement and how they are distributed in energetically optimal arrangements.

In our simulations we use thermal equilibration to relax the spherical system. To find the energetically most optimal distribution we use simulated annealing, an analogy of physical annealing where first the material is heated and then allowed to cool slowly. Thus this method is a “natural solver” [12]. During the crystallisation phase we study the thermodynamic properties of the model system. Computational techniques are used to explore symmetry properties of crystalline structures in the energy minima. In particular, we study disclination patterns, bond-order correlation functions and several thermodynamical properties of interest.

The minimisation of the energy for a many-body system, in order to find the optimal crystalline arrangement, is computationally a hard problem. The energy landscape resulting from all possible configuration states, the state space, is far from smooth. There are many local minima irregularly distributed over the rugged energy landscape

In computer science the difficulty of problems can be classified according to the amount of time a solution requires as a function of problem size. Recently, for other particle systems, it has been shown [13] that the problem of determining if a particle will reach a certain region, tackled with a particle dynamics method at a single temperature, falls in the class of non-deterministic non-polynomial time (NP) problems. This means that the time to find the solution grows more rapidly than a polynomial function of the problem size. It is found that determining the ground state of a 3D cluster of identical atoms interacting through two-body central forces and determining the ground state of folded proteins, belongs to the class of NP-hard problems [14, 15].

Since finding the optimum arrangements for our model probably falls in the class of most difficult computational problems there is little hope to find an easy way out. To find global energy minimum (GEM) states, a method must be used that is optimally efficient in searching through all candidate states. The simulated annealing technique we use in our work conceptually closely matches the underlying physics of our problem. In addition it supports efficient computational schemes to reduce the execution time.

Since the amount of work needed for these search problems grows so rapidly with N we need an optimally efficient computational technique, implying some parallelism. Various forms—from very course grain to very fine grain—are applied in our simulations. We explore the possible implementations of the simulated annealing algorithm on High Performance Computing (HPC) systems such as vector supercomputers and (massive) parallel computers. For the bulk of our work we rely on High Throughput Computing (HTC) by using job parallelism on large numbers of distributed workstations in networks.

This thesis is organised as follows. In Chapt. 2 we summarise some theoretical background on crystallisation and related simulations found in literature. The choice of simulation methods is discussed in Chapt. 3. The parallel versions im-

plemented and tested for the chosen methods and their time complexity analysis are discussed in Chapt. 4. In Chapt. 5 the results from our simulations of the crystallisation process are presented. Then in Chapt. 6 the resulting zero temperature crystalline arrangements are analysed. Finally in Chapt. 7 a discussion on the results and their applications is given together with the conclusions.

Chapter 2

Theoretical background and related work

In this chapter we outline the theoretical background of our work. Relevant crystallisation and melting processes and their underlying theory is summarised. Specific properties of two-dimensional (2D) many-particle systems on flat and spherical surfaces, at finite and zero temperature are discussed. The simulation and analysis methods used are introduced. We end with a concise overview of related work found in the literature.

2.1 Crystallisation

For a many-particle system, freezing—the basic physical transformation of a liquid into either a crystalline or a glassy phase—can occur under suitable conditions of pressure (P) and temperature (T). The freezing transition already takes place for systems with purely repulsive two-body potentials [16] and for systems where an attractive term in the interaction potential is present.

Focussing here on crystallisation, an orderly structure is formed while minimising the free energy of the system at sufficiently low temperature. The free energy is minimised if the system reaches thermal equilibrium (TE). The definition of free energy depends on the thermodynamic constraints.

If the volume (V), number of particles (N) and temperature are constant, the Helmholtz free energy $A = U - TS$ drives the thermodynamics of the system [17], where, U is the internal energy and S the entropy. In a closed system at constant T and P , with only PV work allowed, the Gibbs free energy $G = U + PV - TS = A + PV$ is the appropriate function. In TE, A respectively G are in a minimum. If U , N and V are constant, the entropy is maximised. When the number of particles is allowed to fluctuate, while the chemical potential μ , the volume and the temperature are constant, then, in TE, the product PV is maximised.

During freezing the particles all independently search for the lowest energy po-

sition, the packing constraints being accounted for by the shape of the interaction potential and by the type of boundary conditions.

Around the freezing transition the effect of hysteresis can be observed in freezing-melting cycles. Freezing then sets in only after supercooling, and melting after overheating [18], analogously after overpressure and underpressure, respectively. If during crystallisation this temperature-pressure phase transition region is crossed too fast to allow continuous TE, defects tend to be frozen in. Defects are generally locally metastable arrangements of the crystal. The concentration of defects can be brought back to the equilibrium value by successive slow cycles of application and removal of heat. This restoration process is known as annealing. Removal of defects by annealing down to very low temperatures automatically causes the order in the system to increase.

The phase space of a system covers all attainable states (configurations of the particles) and velocities. The region of phase space to which the system is essentially confined during crystallisation, may be altered by applying a fixed or controlled compression. The transition temperature usually shifts upward. Minimising the free energy under higher density conditions favours more dense packing which in turn raises the mean free energy differences between different arrangements (with respect to kT). In general, as a function of density, the favoured arrangements may have different (degrees of) ordering and symmetry.

In general, the process of self-organisation of the particles during crystallisation towards a configuration with minimal energy is dependent on the thermodynamic conditions. If inequivalent particles are present due to a different interaction potential, or if the particles have inequivalent environments caused by the 3D topology of the 2D space, the equilibrium states can be quite complex in structure.

In this thesis we deal with 2D particle systems which are characterised by two different types of order parameters, and—correspondingly—by two types of defects. A positional order parameter is a measure for the distance over which particle positions are correlated. A bond orientational order parameter gives information about the distance over which bond angles are correlated.

In a flat 2D perfectly triangular crystal lattice, the positions of particles are given by lattice vectors, while bond angles are always oriented with differences $n\pi/3$, $n \in \mathbb{N}$. In this case a perfect correlation in positional and bond orientational order exists over the entire system.

Academically, phase transitions in 2D systems are of special interest due to the absence of conventional long-range positional order [19], even down to low temperatures. This can be understood within the framework of harmonic theory [20], where an expression can be derived for the vector difference between two lattice points minus the distance in the equilibrium lattice. It is found [21] that this vector difference diverges at infinite distance for all non-zero temperatures. We should not exaggerate this effect. Abraham [22] shows that the mismatch in position is only a few angstrom for a typical 2D lattice of the size of the uni-

verse. Thus in practice expected physical properties of an ordered phase do still show-up [23]. This positional ordering in 2D crystals is known as quasi-long-range order. The bond orientational order is not affected by the diverging vector difference.

Defects

Defects can be described geometrically and topologically. Geometrical displacements in a flat perfect lattice do not necessarily lead to defects. Glaser and Clark [24] define geometrical defects as interstitial holes with more than three sides. Their work is further discussed in Sect. 2.7.2. Usually the notion of topological defects is used in literature on 2D systems. In those defects particles are involved with a number of nearest neighbours different from that in the equilibrium situation. For topological defects—by definition—geometric displacements are not significant as long as the number of nearest neighbours of the defect remains the same. Typically such defects can be simply displaced, but not eliminated from the lattice without a global rearrangement. In general they are of two types: disclinations and dislocations (see Fig. 2.1).

By local changes in the geometry at the expense of free energy (local excitations) no single disclination can arise. Also no single dislocation can emerge. Dislocation pairs, however, can be generated by local changes in the geometry.




Figure 2.1: Disclination (left) and dislocation (right). A disclination gives a mismatch in orientational order, a dislocation gives a mismatch between endpoints when circumnavigated (after Allen and Tildesley [25]).

Still considering a flat 2D lattice, a disclination looks like a particle having a “wrong” number of nearest neighbours [19]. It can be shown that on a flat surface for any irregular array of points, the mean number of nearest neighbours as determined by the Voronoi construction (see Sect. 2.6.3), is six. This can be proven theoretically [26] and verified computationally using Monte Carlo (MC) simulations of randomly generated Voronoi polygons [27]. The number of nearest neighbours of a particle is called its Coordination Number (CN). A disclination

can be characterised by a mismatch in orientation as it is circumnavigated [19]. This mismatch in orientation disrupts the bond orientational order (see Fig. 2.1).

A dislocation may be viewed as an extra row of particles stuck partway into the crystal, thus locally disrupting positional order. Alternatively, it can also be viewed as a pair of tightly bound disclinations. Dislocations are characterised by a Burger’s vector [19], which represents the amount by which a path of an equal number of bonds along the regular lattice axis, around the dislocation fails to close. Arguments from elasticity theory [19, 20] indicate that the energy of a single dislocation diverges logarithmically with system size. However, important to be noted, a bound pair of dislocations has a finite energy, and does not disrupt the long distance order. Such pairs tend to emerge even at low temperatures.

For reasons of understanding and discussion, a topological “charge” can be assigned to a disclination of which the magnitude depends on its CN. Throughout this thesis the concept disclination charge is denoted as d-charge. A particle with CN nearest neighbours is a disclination with d-charge = CN-6. Thus a particle with CN 7 is said to have d-charge +1, a particle with CN 5 has d-charge -1. As with electrically charged particles, disclinations with like-sign d-charges tend to repel each other while particles with unlike-sign d-charges tend to attract each other. A bound disclination pair, necessarily of unlike-sign d-charges ± 1 , thus constitutes a single dislocation.

We can relate the physical properties of a flat 2D crystalline solid to the thermodynamic arrangement of the particles, as well as to the thermodynamic arrangement of the defects. A description in terms of defects—less numerous than the particles—requires less degrees of freedom, and thus allows in principle a more efficient description. The underlying regular (d-charge 0) crystal lattice then acts as a medium in which the defects interact with each other. In particular, dislocations can interact strongly by a long range attractive potential increasing as the square of the separation [19]. For free dislocations the mutual interaction is screened and grows only logarithmically with their separation.

2.2 Potentials

Details of the interaction potential between the particles can be of considerable influence on the free energy, and thus on the thermodynamics of the N -particle system. For example, at sufficiently low temperatures, the rotational symmetry of the potential can be important. Potentials with covalent binding and a discrete rotational symmetry (like the covalent binding in carbon atoms) can give rise to intricate structures like buckyballs [28], nanotubes, fullerene-tori and nanoconic tips [29, 30], while van der Waals type spherically symmetric potentials (like for noble gases) form Bravais lattices [20].

Also for the fluid phase the shape of the interaction potential, in particular its core and tail, are of importance in understanding the behaviour of a system.

For example the structure of dense fluids [18] is mainly determined by the short range repulsive part of the potential.

For the studies presented in this thesis we mainly use two types of potentials: the Lennard-Jones (LJ) potential and the (generalised) Coulomb potential. The first has a repulsive core and an attractive tail. The Coulomb potential is purely repulsive. Both are spherically symmetric two-body potentials only dependent on the distance between particles.

The LJ potential (see Fig. 2.2) between two particles i and j is

$$V_{ij} = 4\epsilon\left\{\left(\frac{\sigma}{r_{ij}}\right)^{12} - \left(\frac{\sigma}{r_{ij}}\right)^6\right\}. \quad (2.1)$$

The physical constants ϵ and σ fix the depth and the size of the energy well respectively. To parameterise a LJ system most conveniently in computer simulations, the energy and length units are scaled. Energies are measured in units ϵ and distances in $2^{\frac{1}{6}}\sigma$. Then we have

$$V_{ij} = \frac{1}{r_{ij}^{12}} - \frac{2}{r_{ij}^6}. \quad (2.2)$$

The temperature is also scaled with the parameter ϵ such that the physical temperature T becomes kT/ϵ . For simplicity the scaled temperature is denoted as T . The mass of the particles is taken unity. At short range the LJ repulsive core drops rather fast with r^{-12} , at longer distance the attraction goes to zero due to the r^{-6} term. This long distance behaviour makes it possible to speed up calculations by approximating the energy using a cut-off distance, ignoring interaction (and thus particles) beyond that distance. Another computationally attractive property of the LJ potential is that it can be evaluated cheaply in terms of computer cycles since evaluation of the 6th and 12th powers are easily calculated from the distance squared.

The Coulomb potential between two equally charged particles is given by

$$V_{ij} = \frac{\epsilon}{r_{ij}}. \quad (2.3)$$

The constant ϵ contains the charges and an energy constant. Usually the interaction energy is scaled to unity. This potential falls off slowly, we can not use a cut-off distance in this case.

An often used generalisation of the Coulomb potential, allowing effectively to vary the steepness of the repulsion and the range of the interaction, is

$$V_{ij} = \frac{1}{r_{ij}^n}, \quad (2.4)$$

where $n = 1, 2, \dots$. In the case of $n = 1$ this is the normal (scaled) Coulomb potential. With increasing n the potential becomes more short ranged, the core

becomes “harder”. In the case of $n = \infty$, particle pairs reach an infinite energy if the separation distance is smaller than unity and zero otherwise; in 2D the particles are converted into “hard discs”.

The logarithmic potential, a special case (denoted as $n=0$) of the generalised Coulomb potential [31], is

$$V_{ij} = \ln(r_{ij}). \quad (2.5)$$

A remarkable property of this potential is that for the lowest energy arrangement the dipole moment must vanish (see Bergersen et al. [32]).

In Fig. 2.2 we show the LJ potential and the generalised Coulomb potential for $n=0, 1, 12$ and ∞ .

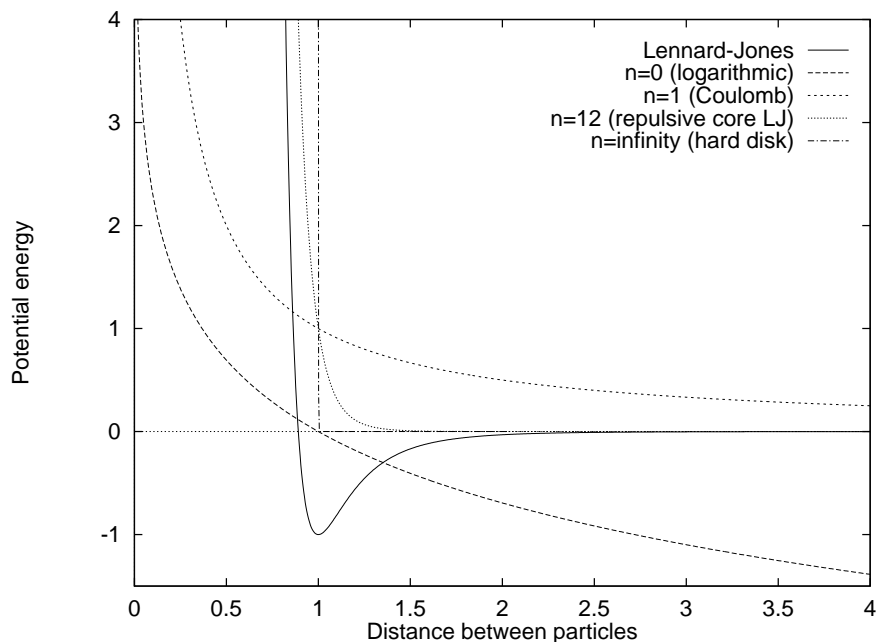


Figure 2.2: The Lennard-Jones potential and generalised Coulomb potentials of order 0 (logarithmic), 1 (Coulomb), 12 (repulsive core LJ) and ∞ (hard disc).

2.3 Boundary conditions

In 2D many-particle systems edge effects can be relatively important. The fraction of particles at the edge is proportional to $N^{-1/2}$ [18]. Depending on research purposes, several types of boundary conditions can be used. If edge effects are to be explicitly included in the study, the use of free boundaries is appropriate. If the effect of enclosing walls is important, a finite region of space with appropriately imposed boundaries, interacting with the particles, is used. In these two cases the boundaries have physical meaning. If the system size to be studied is

small—for instance for computational reasons—then bulk effects can be studied through the use of (artificial) periodic boundary conditions. To eliminate the explicit use of boundary conditions and still be able to study bulk properties in 2D, spherical boundary conditions can also be used. But then the underlying spatial topology has changed. This last type of boundary condition is at the heart of our studies; it is discussed below at some length. Here it is not used as an artificial boundary condition, but as part of the model we use to mimic natural spherical many-particle systems.

The shape of the used interaction potential is of influence on how any boundary conditions can be applied. For purely repulsive potentials some form of restrictive boundary conditions must be present to keep the particles from moving outward to infinity. Either wall boundaries must be fixed (the density should be kept constant), or an external pressure has to be applied. The crystallisation process and the lattice formed during crystallisation usually depend on the type and shape [25] of boundary conditions applied.

Free boundaries

For sufficiently large system size N , with free boundaries a major part of the system tends to reach the bulk equilibrium structure. With free boundaries no artificial restrictions are introduced in the simulation. This type of boundaries can only be used if the potential is capable of keeping the system together (avoiding evaporation) in the temperature range of interest; it must have a sufficiently attractive term (in terms of kT). For the LJ potential this is possible [33]. For a long range purely repulsive potential the particles move away towards infinity even at zero temperature.

Fixed wall boundaries

For studies where the influence of walls is important, fixed wall boundary conditions can be used to confine the system to a finite region. The packing constraints can severely influence the formed lattice structure. Usually the walls also have an interaction potential with limited range (such as used in the 3D study of surface wetting [34]). With fixed wall boundary conditions for small regions the optimal packing has a strong dependence on the system size [35], in general the optimal lattice structure is not known. The size of the area enclosed by wall boundaries determines the average density in the system and can provide external pressure.

Periodic boundary conditions

In order to describe an infinite system in terms of a finite number of degrees of freedom, Periodic Boundary Conditions (PBC) can be applied. This type of boundary condition is artificial and care must be taken that it does not influence the system too severely. In flat 2D space most common is the square (or rectangular) PBC [18], topologically the surface of a 3D toroid. In this case N particles are initially placed in a square of size L . This is the unit cell of a simple lattice of replica systems. Every time a particle moves out of the original cell (say at $x = +L/2$) its periodic image will move into the cell (at $x = -L/2$). Although PBC effectively eliminate edge effects, they do introduce some effects of their own. In general correlations are introduced among particles separated by distances of order L since fluctuations with wavelength larger than L are suppressed. A periodic 2D fluid can no longer be isotropic but has a built-in symmetry derived from the boundary conditions. Moreover, crystal lattices that are commensurate with the boundary conditions are clearly favoured over others.

The size of the unit cell is not necessarily fixed and can be adapted to the ensemble used. For long range purely repulsive potentials the boundaries must be fixed at positions that give the desired average density, or they may fluctuate while maintaining a given external pressure.

Spherical boundary conditions

To study bulk properties for 2D systems with a relatively small value of N , Spherical Boundary Conditions (SBC) can be used in stead of PBC. This type of boundary condition is again purely artificial. Hansen *et al.* [36] use this strategy to calculate the self-diffusion in a classical electron gas. For a system of particles with attraction, the size of the sphere can either be fixed or adjusted to keep a specific thermodynamic quantity constant. For purely repulsive potentials the radius must be fixed or an external pressure is needed if the radius R of the sphere is left free. For the (generalised) Coulomb potential it can be shown that—in the $kT \rightarrow 0$ limit— R is not of influence on the stability of crystalline structures. All distances, and thus potential energies, scale with powers of R . Therefore the energy sequence and the associated arrangements remain fixed for any fixed R .

At sufficiently high temperatures, the change of the underlying topology from PBC to SBC has only a limited effect on liquid and gaseous bulk properties. It does, however, affect properties of the crystalline arrangement in an essential way. This is due to the necessity of incorporating inherent, unbound disclinations. It can be shown by using a theorem due to Euler [37, 38] (stating that the number of faces minus the number of edges plus the number of vertices equals 2), that in a distribution of points on a spherical surface—in terms of topological defects (see Sect. 2.1)—a net d-charge of -12 must be present [39]. In addition to the inherent disclinations, lattice defects may be present, induced by local excitation. There

is a tendency for local excitations to cluster near the inherent d-charges due to net attraction and local “d-charge polarisation” of the lattice surrounding the inherent d-charges.

We apply the SBC not as a method to avoid using PBC but as part of the model to mimic natural spherical systems (see Chapt. 1). In the presented studies we distinguish—as much as possible—between isolated inherent disclinations (possibly with defect cluster), and isolated defects due to thermal excitation, randomly distributed and without long range interaction.

Strain effects, distorting the lattice, cause an effective repulsive potential extending over the lattice medium between isolated inherent (like-sign d-charge) disclinations. At low temperatures these disclinations therefore tend to settle as far as possible from each other, thereby creating a tendency for global spherical symmetry of the defect distribution. It is, for example, possible to construct a distribution of particles over a spherical surface with six particles of d-charge -2 having a tendency for regular octahedral symmetry. However to spread the strain put on the lattice as evenly as possible, we expect a tendency towards 12 defects with d-charge -1, settling in the energetically most optimal distribution, spread over the sphere according to regular icosahedral symmetry.

Lattice defects due to thermal excitation, have net d-charge zero since these consist of pairs of unlike-sign d-charges. It can be expected that in principle the distribution of exited disclination clusters on a sphere resembles the one on a flat surface. Due to the attraction between unlike-sign d-charges, clustering of disclinations tends to occur. At sufficiently large N and sufficiently high temperatures exotic and complex clusters emerge. Because of the tendency towards arrangements with icosahedral symmetry for the inherent defects, at TE all the induced clusters of disclinations tend to conform to this symmetry.

2.4 Simulation methods

The studies in this thesis comprise 2D simulations of an N particle system on a spherical surface. To study the behaviour of the system at finite temperature (see Chapt. 5) we use different types of simulation methods: stochastic and deterministic. Most of the shown results are obtained with the (stochastic) MC simulation method. Some special simulations are performed using the (deterministic) Molecular Dynamics (MD) method. For finding zero temperature crystalline arrangements (see Chapt. 6) we use a combination of two techniques: the stochastic Simulated Annealing (SA) and the deterministic Steepest Descent (SD) methods, related to the above mentioned methods. Simulated annealing is a series of MC simulations with decreasing temperature. The Steepest Descent method can be viewed as a limiting case of Molecular Dynamics (infinite quenching and variable time steps). These methods are further discussed in Sect. 3.2.

2.5 Ensembles

In the MC simulations we use two ensembles: the canonical (constant NVT) and the isothermal-isobaric (constant NPT) ensemble. In the MD simulations we use the microcanonical (constant NVE) ensemble, the 2D volume V being the available spherical surface area.

For the canonical ensemble N and V —and consequently the density ρ —are constant. Simulations in this ensemble are relatively simple compared to e.g. constant pressure P simulations; only the positions of the particles change during the simulation. In this ensemble density fluctuations tend to be suppressed. Density fluctuations are important when studying the excitation of defects [19]. Suppression of excitation defects may have a severe effect on system properties, especially around a phase transition.

In the isothermal-isobaric ensemble P is kept constant while V is allowed to fluctuate. The radius of the sphere must be allowed to change, and the particles must be allowed to move along the spherical surface.

In the microcanonical ensemble used in our MD implementation, T can fluctuate. When this method is used for annealing, the velocities are rescaled downward, and thus the average T of the system decreases. The system can still escape from local energy minima by trading kinetic energy for potential energy. At high T sufficient kinetic energy is available to cross potential energy barriers. As the kinetic energy is reduced, more and more barriers turn out to be too high to be crossed. Then parts of the phase space get out of reach for the system. For sufficiently low temperatures the kinetic energy is too low for crossing barriers: the system is trapped. The region of state space which—after applying SD—leads to the same energy minimum is called the capture basin of that minimum.

2.6 Analysis metrics and methods

In this section we discuss the methods we use to analyse the data. On a sphere not all methods constructed for flat 2D space can be applied. The distance metric depends on the property being evaluated. Below we discuss how we determine the following properties: energy, specific heat, pressure, Voronoi polygons, clusters of d-charges, symmetry, enantiometry, centre of mass, radial distribution functions, the structure factor and orientational correlation functions.

2.6.1 Energy and specific heat

A rough way to locate the transition region from liquid to solid is to calculate the average potential energy E as a function of temperature T . The energy is averaged for a sufficiently large sample of arrangements each weighted by its probability. The constant volume specific heat $C_v = (\partial E / \partial T)_v$, actually represents

the slope in $E(T)$. The values of C_v are determined from $E(T)$ by the following method. A line is fitted to κ successive points. The slope of the line is taken as the derivative at the average temperature of the κ points. In this way the data are smoothed while determining the derivative.

2.6.2 Pressure

For the calculation of the pressure several methods can be used. We use a method described by Eppenga and Frenkel [40] for hard platelets and adapted for potentials with an attractive part [41, 42]. This method does not explicitly use coordinate vectors or forces. The pressure is evaluated in the canonical ensemble using the derivative of the Helmholtz free energy with respect to volume. It involves trial volume changes rather like the test particle determination of the chemical potential [43]. In this case, the volume change is a trial only: it does not actually take place. The size of the volume change is $\Delta V = V' - V$, where V' is the trial value for the volume. If ΔV is constant in magnitude and sign, and the ratio V'/V is close to unity, the pressure can be evaluated as

$$\beta P = \frac{N}{V} + \frac{1}{\Delta V} \ln \langle \exp(-\beta \Delta E) \rangle, \quad (2.6)$$

where $\beta = 1/(kT)$ with k as Boltzmann's constant and where ΔE is the energy difference resulting from the trial volume change. It is known [40] that for hard core potentials only volume reductions are appropriate, while for soft core potentials both reductions and expansions are necessary. If we take ΔV sufficiently small the difference between pressures calculated for different ΔV (also of different sign) is much less than the corresponding deviation in pressure.

2.6.3 Defect analysis

Voronoi analysis

In order to describe the structure of an arrangement in terms of bond-angle correlations and defects (see Sect. 2.1) we first need to identify the nearest neighbours of each particle. The Voronoi construction [44] is a most natural way to uniquely define a set of neighbours for each point in a (possibly disordered) n -dimensional array.

In our case this construction provides a tessellation of the sphere with Voronoi polygons, each following the curvature of the surface. A Voronoi polygon, also called Wigner-Seitz cell [26, 20], surrounding a given particle i , is the polygonal 2D region of points nearer to i than to any other particle. The number of sides of a Voronoi polygon gives the number of neighbours of the corresponding particle, the coordination number CN. The dual of the Voronoi tessellation, the Delaunay triangulation [44], consists of the triangular network formed by shortest

connecting lines between neighbouring particles. Three adjacent neighbour cells meet in each vertex of the graph formed by the polygon sides. This is not always true [24]. In a flat perfect square lattice, for example, the vertices in the Voronoi construction are points where four cells intersect. Normally, in thermally disordered arrangements such four-(and higher) coordinated vertices occur, however, with essentially zero probability [26].

In order to set the boundaries for the polygonal region around particle i for a given arrangement of the system [45], we begin by selecting particles Q_i in the neighbourhood of particle i , usually the twelve closest. On the midpoints of the lines connecting i to all $j \in Q_i$, perpendicular lines are placed. All pairs of these lines are intersected with each other to find the set of points S . The point S_{jk} , with $j, k \in Q_i$, is the intersection of the lines perpendicular to ij and ik (see Fig. 2.3).




Figure 2.3: The Voronoi construction: in S_{jk} three Voronoi polygons meet.

Obviously, on a sphere the lines that we refer to are great circles that intersect at two diametrically opposite points. The intersection point closest to i is chosen. The points S_{jk} are put through a selection process: if a point S_{jk} happens to be closer to a point $l \in Q_i$ (different from j and k) than to i , it is removed from S . The points left after selection are the vertices of the Voronoi polygon for particle i . This recipe is applied for all N particles of the system, yielding a unique Voronoi tessellation.

Cluster analysis

Disclinations—identified by their d-charge—have a tendency to cluster. In particular a dislocation is in fact a dimer cluster of unlike-sign d-charges. But also larger clusters consisting of several dislocations are observed [24]. Giarritta *et al.* [39] find for a hard disc system on a spherical surface at intermediate densities that most disclinations are grouped together; the number of isolated d-charges ± 1 in their study never exceeds 2% of the total number of disclinations.

To monitor the clustering of defects we list and mark for each disclination its nearest neighbours that also have d-charge $\neq 0$. This step is repeated recursively for all particles in the list. Then one cluster is identified. For disclinations that are not yet marked this procedure is repeated using a different mark. The thus identified clusters can be characterised by the net d-charge, the d-charge distribution and the total number of d-charges occurring in that cluster.

2.6.4 Symmetry analysis

Applicable symmetry operators

For the system at hand, the allowed symmetry operators are rotation, reflection (with a mirror plane) and inversion (point mirroring through the centre of the sphere).

The following symmetry groups are possible [46]: those having 1) no rotational symmetry axis, 2) exactly one rotational symmetry axis, 3) one main rotational symmetry axis with an additional system of two fold rotational symmetry axes perpendicular to the main axis, and finally 4) the symmetry groups of the regular polyhedra. Moreover, all these groups can have reflection symmetries and inversion symmetry in addition to the identity operator. We use the Schoenflies [47] notation for the names of the symmetry groups.

The groups with no rotational symmetry are: C_1 with only the identity operator, and corresponding to complete absence of symmetry; C_i with inversion; C_s with reflection.

The groups with one rotational symmetry axis are: C_n with n rotations; C_{nh} with n rotations and a mirror through the rotational axis, (if n is even the inversion operator is also present); C_{nv} with n rotation operations and a mirror operator perpendicular to the rotational axis; and finally for completeness we only mention the group S_n containing rotations and rotation-reflections.

It is possible to construct spherical distributions such that $n \in \mathbb{N}$, in practice n is limited to the range 1, 2, ..., 6.

The groups with one n -fold axis and n 2-fold ($n \geq 2$) axes at right angles with the n -fold axis are: D_n with the rotations of the main axis and those of the 2-fold axes; D_{nh} with the rotations of group D_n and mirrors which contain the main axis, (if n is even the inversion operator is also present); D_{nd} with the rotations of group D_n and mirrors that bisect the angle between two neighbouring 2-fold axes, (if n is odd the inversion operator is also present).

The groups of the regular polyhedra are: T with the rotations for the tetrahedron; T_d and T_h which contain the rotations of the tetrahedron and additional reflections; T_h also contains the inversion operator; O with the rotations of the octahedron; O_h with the rotations for the octahedron and additional reflections and including inversion; I with the rotations for the icosahedron; I_h with the rotations of the icosahedron and additional reflections and inversion.

The group with the lowest symmetry is clearly C_1 , the group with the highest symmetry, if we restrict n to the range 1, 2, ..., 6 is I_h since this group contains the highest number of symmetry operators.

Finding the symmetry group

In symmetry analyses the aim is to search for the highest symmetry group that can be applied to a given N particle arrangement mapping that arrangement onto itself. In our case the arrangement is generated by a computer optimisation algorithm. This means that the positions have a certain inaccuracy. Therefore the mapping criterion needs a given tolerance allowing for minor deviations. The symmetries are determined in several steps. First the rotational symmetry is considered, then the reflection and inversion symmetries.

There are three possible locations for rotational symmetry axes through the centre of the sphere. An axis can be through a particle, halfway an edge between two neighbouring particles and in the centre of any face spanned by neighbouring particles.

To find the rotational symmetry axis, a series of rotation operators is applied to the arrangement for all three location types described above. After applying a given operator the new arrangement is matched with the original. If the complete arrangements can be mapped (within the tolerance) onto each other, that operator is part of the symmetry group.

After finding all the rotational symmetries present in the system the name of the symmetry group can be given. We introduced the symmetry groups such that they can be identified by their typical combination of symmetry axes [46]. If there is no or only one rotational axis, the group name is C . If there are more rotational axes, and one main axis, the rest of the axes being two-fold and at right angles with the main axis, the group name is D . The notation of these groups also has a subscript number that gives the symmetry of the highest symmetry axis. Finally, there are the regular polyhedral groups, T for tetrahedral, O for octahedral and I for icosahedral.

Concerning reflection symmetry, there are four different types of mirror planes. They can be perpendicular to a rotational axis, denoted with the subscript h , or parallel to the axis, denoted with v . It is also possible to bisect the angle between two rotational axes in which case the mirror plane is denoted with the subscript d . If no rotational axis is present there can still be a mirror plane which is then denoted by s .

For groups with one or more rotational axes the possible mirror planes can be checked by testing for the first three types of mirror planes. If no rotational axis is present an exhaustive search is necessary to detect a mirror plane.

Inversion symmetry is easily found by inverting all the particles and checking whether the original and new arrangement match within the allowed tolerance.

Enantiometric arrangements

Arrangements can be enantiometric. In such a case two distinct stable solutions exist which have exactly the same energy and other properties but they cannot be mapped onto each other by rotation operators. These solutions have opposite chirality or handedness. Two enantiometric arrangements can be transformed into each other by an improper isometry consisting of a rotation combined with a reflection in a plane perpendicular to the axis of rotation.

To check whether an arrangement has this enantiometric property the following algorithm [32] can be applied. First make the inverse of the original arrangement by inverting all particles. If this inverted arrangement can be made to match with the original arrangement by using rotation operations it is not enantiometric, otherwise it is.

Centre of mass

The centre of mass of an equilibrium spherical arrangement one could expect to be at the centre of the sphere, here taken as the origin of the coordinate system. This is not necessarily the case, partly depending on the two-body potential involved. Energetically optimal arrangements with the logarithmic potential in fact always have their centre of mass in the origin [32]. For the spherical 2D LJ system the centre of mass usually does *not* coincide with the origin (see Sect. 6.3.6). In simulations with the Coulomb potential it is assumed that a uniform neutralising background charge is present. If the charged particles have positions such that the centre of mass does not coincide with the centre of the sphere a dipole moment is present. Erber and Hockney [48] find for spherical 2D Coulomb systems dipole moments of $O(1)$ for any N .

2.6.5 Order analysis

During crystallisation the freedom of movement of the particles in the system decreases. The particles individually search for the lowest energy positions consistent with their local constraints. The overall minimisation of free energy tends to lead to global order in the distribution of the particles. The order in the system can be measured by order parameters. Two-dimensional particle systems are characterised by two different types of order parameters: positional and bond orientational order. The concentration of different types of defects disrupting local and global order as a function of temperature can be used as an order parameter. Here we are especially interested in the ordering of topological defects, in particular d-charges. To study this aspect we use a positional order parameter that only keeps track of the d-charge distribution.

Radial distribution function

The radial distribution function (RDF), denoted as $g(r)$, is a pair correlation function. It gives the normalised probability of finding a particle at radial distance r of a reference particle measured along the spherical surface. On a sphere with radius R the RDF is evaluated as

$$\rho g(r) = \frac{\mathcal{N}(r)}{\mathcal{A}(r)}, \quad (2.7)$$

where $\rho = N/(4\pi R^2)$ is the density, $\mathcal{N}(r)$ the average number of particles located within a spherical ring (centred at a reference particle) with distance r , thickness $dr \ll r$, and area $\mathcal{A}(r) = 2\pi R \sin(r/R)dr$.

The RDF is an indicator for the positional organisation of the system. For a disordered system $g(r) \approx 1$ everywhere at large distance. A liquid has a $g(r)$ that mildly fluctuates around unity at short distances, where the particle positions are correlated, and is unity at large distance. A crystalline solid is characterised by sharp peaks in $g(r)$, indicating that only at specific discrete distances from a particle other particles are found with appreciable probability. On a sphere the maximum distance along the sphere possible is half the circumference, thus $r_{\max} = \pi R$. At distances approaching r_{\max} the statistical uncertainty increases: the surface of the rings at those distances contain relatively few particles, and large fluctuations are found in the RDF for a single configuration at non-zero temperature. Therefore, if the RDF has to be precise up to the largest possible distances many different configurations generated with the same thermodynamic conditions have to be averaged.

The structure factor

The structure factor $S(\mathbf{k})$ is related to the susceptibility which measures the response of the arrangement to a periodic external potential with wave-vector \mathbf{k} . Thus it describes the Fourier components of density fluctuations in space. $S(\mathbf{k})$ is a quantity (like the RDF) to probe the positional order in the system [49]. For the flat 2D case the structure factor is defined as

$$S(\mathbf{k}) = \frac{1}{N} \left\langle \sum_{i=1}^N \sum_{j=1}^N e^{i\mathbf{k}\cdot\mathbf{r}_i} e^{i\mathbf{k}\cdot\mathbf{r}_j} \right\rangle, \quad (2.8)$$

where \mathbf{k} is the wave vector and where \mathbf{r}_i and \mathbf{r}_j are the position vectors of two particles. The bracket pair stands for the ensemble average.

A more practical formula is derived from Eq. 2.8 [49] by averaging over all directions of the \mathbf{k} -vector. We find that

$$S(k) = 1 + \frac{2}{N} \sum_{r_{ij}} J_0(kr_{ij}), \quad (2.9)$$

where $r_{ij} = r_i - r_j$ and J_0 is the zeroth order Bessel function of the first kind. The disadvantage is that information on the angle dependence of \mathbf{k} is lost. The $S(k)$ is related to $g(r)$ by the Fourier transform

$$S(k) = 1 + 4\pi\rho \int_0^\infty r^2 \frac{\sin(kr)}{kr} (g(r) - 1) dr. \quad (2.10)$$

On a sphere the interpretation of the wave vector \mathbf{k} is not straightforward. It can be interpreted as a wave wrapped around the surface of the sphere. Therefore on a sphere we have to be careful with using the structure factor. One possible solution is to calculate the structure factor as if we are working on a flat surface. Then, however, we have problems with wave vectors corresponding to reciprocal lattice vectors where the curvature of the sphere is appreciable.

The height and position of the first peak of $S(k)$ plays an important role in perturbation theories of liquids around the freezing transition (see Sect. 2.7.3). In Euclidian space the first peak in $S(k)$ corresponds to the the first reciprocal lattice vector. On a sphere, no Bravais type of lattice is possible, and thus no reciprocal lattice vectors exist. However, for large N , where the inherent d-charges are a minority, most particles favour a regular triangular lattice structure. Thus in our case lattice vectors correspond to the (distorted) triangular arrangement of those (d-charge=0) particles.

Defect radial distribution function

For systems with spherical symmetry it is convenient to split up the RDF for particles with specific combinations of d-charges (see also Giarritta et al. [39]). Such new RDF's, that together constitute the total RDF, are denoted as $g_{\alpha\beta}(r)$ where α and β are labels for the various d-charges. In this thesis these radial distribution functions are called defect radial distribution functions (DRDF's). The DRDF for a particle distribution on a sphere with radius R is defined as

$$\rho_\alpha \rho_\beta g_{\alpha\beta}(r) = \frac{\mathcal{N}_{\alpha\beta}^{\text{pairs}}(r)}{\mathcal{A}(r)}, \quad (2.11)$$

where for example $\rho_\alpha = \mathcal{N}_\alpha / (4\pi R^2)$ is the average partial density of particles with d-charge α , where $\mathcal{N}_{\alpha\beta}^{\text{pairs}}(r)$ is the average number of distinct pairs of particles with d-charges α and β lying within a spherical ring (centred at a reference particle) with curved radius r , thickness dr and where the area $\mathcal{A}(r) = 2\pi R \sin(r/R) dr$. At non-zero temperature the averaging of \mathcal{N}_α and $\mathcal{N}_{\alpha\beta}^{\text{pairs}}(r)$ in Eq. 2.11 is taken over the arrangements sampled for a fixed density ρ . Note that $g_{\alpha\beta}(r)$ is equal to $g_{\beta\alpha}(r)$.

This split-up according to d-charge actually constitutes the spectral decomposition of the full pair distribution function $g(r)$ in the form

$$g(r) = \sum_{\alpha,\beta} x_\alpha x_\beta g_{\alpha\beta}(r), \quad (2.12)$$

where $x_\alpha = \rho_\alpha / \rho$.

Oriental correlation function

The nearest neighbour n -gonal bond orientational order on a flat surface is measured by the bond orientational correlation function [19]

$$g_n(|r - r'|) = \langle \exp in(\theta(\mathbf{r}) - \theta(\mathbf{r}')) \rangle, \quad (2.13)$$

where $\theta(\mathbf{r})$ is the angle made by a bond between a particle at \mathbf{r} and a nearest neighbour with respect to an arbitrary fixed reference axis.

In flat 2D an arbitrary reference axis can be defined straightforwardly. On a sphere this is not possible. If the equator is chosen as a reference instead, for the particles near the poles the angle of a bond to a neighbour with respect to the reference is strongly distorted. In general, on a sphere no single fixed direction can be taken as a reference.

One solution to this problem is the following. A great circle is chosen as a reference and only particles near this circle are used in evaluating the bond order. This can be done for a large number of randomly chosen great circles to obtain a representative statistical average.

Another redefinition of the orientational correlation function for the case of spherical boundary conditions is given by Giarritta *et al.* [45]. In their algorithm all particles are used at once. This is possible since they do not define any global reference direction. Instead, as reference is taken the line joining two particles at distance r . The angles of the nearest neighbours of those particles with that reference line are used. Thus we get the relation

$$h_n(r) = \langle \cos[n(\theta(\mathbf{r}) - \theta(\mathbf{r}'))] \rangle, \quad (2.14)$$

where $r = R \arccos(\mathbf{r} \cdot \mathbf{r}')$ is the geodesic distance (along a great circle) between the particles whose positions on the sphere are defined by the unit vectors \mathbf{r} and \mathbf{r}' . The angle $\theta(\mathbf{r})$ is measured on the plane tangent to the sphere at \mathbf{r} between the projections of the line joining the particle at \mathbf{r} with a nearest neighbour (chosen at random among those identified through the Voronoi construction), and of the line connecting \mathbf{r} and \mathbf{r}' along the geodesic.

Relevant for our studies are the functions $h_n(r)$ for $n=5$ and 6. In a flat 2D triangular lattice the difference $\theta(\mathbf{r}) - \theta(\mathbf{r}')$ is always a multiple of $\pi/3$ so that, while the function $h_5(r)$ vanishes everywhere, $h_6(r)$ has unit value at distances corresponding to the lattice distances and zero otherwise. Deviations from these values reveal distortions in bond orientational order. On a sphere the decay with distance is a result of disorder in the system induced by the equilibrium thermodynamics and by any finite curvature of the space.

Defect density

In a liquid many defects are present. During crystallisation their number greatly diminishes as the order in the system increases. In our case we focus on the number of d-charges present. The average number of d-charges N_d in the system is found as

$$N_d = \langle \sum_{\alpha} \mathcal{N}_{\alpha} \rangle, \quad (2.15)$$

where the sum is taken over all d-charge values α , and where \mathcal{N}_{α} is the amount of d-charges of value α in the system.

For a flat 2D particle system (either infinite or with appropriate boundary conditions) with a circularly symmetric interaction potential in its lowest energy arrangement, no dislocations or disclinations are present. As discussed in Sect. 2.3, on a sphere disclinations are inherently present. In addition to the inherent d-charge (net value -12), introduced by the underlying 3D topology, dipolar pairs with net d-charge zero tend to occur in flat as well as in spherical systems. In the spherical case it is to be expected that the inherent d-charge centres influence the long range orientational order and the positioning of d-charge pairs.

2.7 Temperature $T > 0$

In studies of the thermodynamical behaviour of 2D particle systems, simulations play a dominant role. In some cases physical systems like for example atomic monolayers on a substrate form the actual basis and motivation for the simulation studies. Ideal 2D systems, however, do not exist. In order to develop a full understanding of 2D systems, computer simulations provide an important means to test current 2D thermodynamic theories. Below we discuss important theories that are being validated mainly through computer simulations.

2.7.1 Thermodynamics

The usual theoretical approach [26] to the liquid-solid transition is to treat it from the standpoint of the melting of the solid. This has the advantage that the system can be described using the crystal lattice as a reference frame, and the problem becomes that of analysing the stability of the lattice against thermal agitation. There are theories, however, which treat the fluid phase exactly whereas the solid phase is described approximately. These are theories of freezing rather than melting, which are consequently better tailored to our computer experiments.

The melting and freezing transitions can be first and second order. During a first order transition a discontinuous change of the appropriate thermodynamic

function is found at the critical temperature T_c . In a second order phase transitions continuous changes of the pressure or volume are found. Some theories predict a first order phase transition for the melting of flat 2D particle systems while others find two second order phase transitions with two critical temperatures [19, 24].

2.7.2 Melting

The intuitive idea underlying most melting theories is that only a limited set of collective coordinates is needed to characterise the difference between the solid and the liquid [24]. These theories differ primarily in their choice of relevant degrees of freedom, thus in their modelling of the melting mechanism. Here we only discuss defect mediated melting theories. These are expected to be of special interest for a system with inherent d-charge. In these melting theories the coordinates of the defects (dislocations and disclinations) are the relevant degrees of freedom. The Kosterlitz-Thouless-Halperin-Nelson-Young (KTHNY) theory [50, 51, 52, 53] of 2D melting assumes the occurrence of two subsequent second order phase transitions. As discussed in Sect. 2.1 the low temperature structure is characterised by a quasi-long-range positional order and long-range bond orientational order. The first transition towards melting is associated with a dislocation-unbinding transition. At low temperatures bound dislocation pairs can be locally excited in the solid. The density of bound dislocations increases with temperature, until— at a Kosterlitz-Thouless phase transition (KT) with critical temperature T_{KT} —the free energy for an isolated dislocation becomes negative. Then free dislocations proliferate and the system loses, for example, its resistance to shear stress. It is found [52], however, that the liquid above this dislocation-unbinding transition is still not isotropic. The dissociation of dislocation pairs destroys the quasi-long-range positional order but leaves the fluid characterised by quasi-long-range order in the bond orientations. Halperin and Nelson introduced the term “hexatic” for the resulting phase. A second transition at $T \geq T_{KT}$, in which disclinations unbind, is required to obtain an isotropic fluid characterised by only short range positional and bond orientational order. In this transition the tightly bound disclination pairs (constituting a dislocation) unbind in a way analogous to the dislocation-unbinding transition.

A competing theory [54, 55] assumes that melting begins at grain boundaries. These are collective dislocation excitations; longitudinally stretched clusters (strings) of disclinations with alternating d-charge. A grain boundary in a flat solid can be viewed as the effect of rotating one patch of crystal with respect to another.

Fisher *et al.* [54] conclude for small angle grain boundaries (large spacing between the d-charges), that the melting transition is of the KTHNY type with second order transitions. Chui [55], taking density fluctuations and interactions between grain boundaries and dislocation pairs into account, concluded that the

transition is first order [19].

On a sphere we might expect that grain boundary melting is important since disclinations are inherently present, and tend to form and induce polyhedral arrangements [39, 56]. The faces of the polyhedral distribution of disclinations—projected onto the sphere—consist ideally of a (distorted) piece of triangular lattice. Between the faces typically grain boundaries might exist.

The above mentioned melting theories are instability theories, which predict the point at which the solid phase becomes unstable and defects are excited in large quantities.

The mechanistic picture of two-dimensional melting suggested by the work of Glaser and Clark [24], mentioned in Sect. 2.1, can be summarised as follows: two-dimensional melting is a condensation of localised, thermally generated, geometrical and topological defects. Geometrical defects in their work are considered to be interstitial “holes” with more than three sides (square, pentagonal or hexagonal), formed by rings of particles arranged so that all particles touch their neighbours. In a perfectly regular triangular lattice all holes are triangular. Holes with more than six sides are comparatively rare. It is pointed out that topological defects are responsible for the symmetry change associated with the 2D melting transition, but that this does not imply that this transition and the liquid structure can be understood entirely in terms of the statistical mechanics of topological defects. The condensation of geometrical defects is a phase transition resulting from attractive interactions between the excited geometrical defects. This leads to the identification of the geometrical excitations relevant to 2D melting and provide evidence that the “effective attractive interactions” [24] produce a condensation transition that is identified with melting.

2.7.3 Freezing

The crucial difficulty in describing the freezing transition is to construct a theory in which the onset of long-range ordering is predicted without being “built-in” from the start, by using a crystal lattice as a reference frame [26]. In other words, although we know experimentally that melting is reversible—albeit often with hysteresis effects—crystallisation is much harder to treat theoretically than melting.

The height and position of the first peak in the structure factor play an important role in the highly successful thermodynamic perturbation theories [57, 58, 59] of 3D simple liquids. In addition, it is known that for many 3D liquids studied by computer simulation [60, 61] there is a simple relation between the point where the liquid freezes and the height of the first peak in the structure factor. This can be understood in terms of the observation that the structure of the dense fluid is dominated by the steeply repulsive steric interactions between molecules [18]. This observation is expressed in the melting criterion from Hansen and Verlet [62]: along the melting curve, the value of the first maximum of the structure factor is

constant and approximately equal to the first maximum of the structure factor of a hard-sphere fluid at the freezing density. This finding provides evidence for the geometrical nature of the transition. It suggests that the transition is predominantly structural, occurring when the fluid-phase structural correlations are strong enough. A theoretical basis of the structure-factor-criterion is given by the density functional theory as used by Ramakrishnan and Yussouf [63, 64]. This theory considers the periodic density modulation of the solid near melting as a perturbation of the uniform liquid. The free-energy difference between solid and liquid is determined by short-range correlations. The relevant degrees of freedom are the Fourier components of the density. This mean-field theory leads to a first-order transition.

The freezing parameters do not depend on the particle potential and contain no adjustable parameters because of the geometrical nature of the transition. There are, however, freezing parameters that do depend strongly on the potential used. Examples are melting temperature and volume change at freezing.

2.8 Temperature $T = 0$

Here we discuss some zero temperature properties found by computer simulations of model systems.

Once crystallisation is completed in the simulations and all particles have lost their mobility, a suitable method can be used to precisely minimise the energy. Then we have an arrangement that necessarily is at a minimum in potential energy (apart from saddle points). Because of the approximate nature of practical implementations this does not necessarily have to be the global minimum in potential energy.

During the crystallisation phase the system searches—while minimising the free energy—for deep basins in the potential energy landscape over the state space (see Chapt. 3), the particles settle into a tightly packed arrangement. At this point the system being trapped in a basin, the overall structure of the arrangement remains the same while minor geometric changes are still possible. At the deepest point in the basin a configuration is reached with (locally) minimal potential energy. Indeed, defect excitations with respect to the global minimum may be frozen in. Many different defect excitations of the global minimum may have (nearly) degenerate energies.

2.8.1 Many local minima

There are many distinct (meta)stable arrangements of N particles over the surface of a sphere. These arrangements correspond to (local) minima in the potential energy. For several comparable systems it is found [65, 48] by computer simulations that the number $M(N)$ of minima in the state space increases exponentially

with N , thus

$$M(N) = \mu e^{\nu N}, \quad (2.16)$$

where μ and ν are system dependent. An early study performed by Hoare and McInnes [66], provides 988 metastable arrangements for a 12-particle system with LJ potential. For a soft sphere model with r^{-12} pair potential, La Violette and Stillinger [67] find $\nu \cong 0.07$ and—taking $\mu=1$ —Stillinger and Weber [68] obtain $\nu \simeq 0.16$ for a modified, finite range LJ system with finite volume and 32 particles. Furthermore, they obtain a significant increase of ν with decreasing density for the latter system [69].

2.8.2 Built-in symmetry

The enormous amount of minima expected for simulations on large systems make it difficult to determine the global energy minimum. Therefore besides the unconstrained system we also implemented a method that decreases the number of local minima by constraining the positions of the particles such that the arrangement has a specific symmetry.

In general in our simulations we want to be sure not to exclude the most optimal solution, the true global minimum. Therefore we do not want to apply any *a priori* constraint. On the other hand, however, if we would have indications that the optimal arrangement has a certain symmetry, we might not want to spend our simulation time on irrelevant solutions. For special simulation studies we thus prefer to fix the symmetry of the arrangements. Then the number of degrees of freedom is greatly reduced. Due to a reduced computational time this enables us to find, for a given computational effort, energetically optimal arrangements for much larger N -values than would be possible otherwise. To do this we consider only arrangements with the required symmetry being built-in. The method described below is analogous to the one used by Edmundson [70]. We restrict ourselves here to icosahedral symmetry. The presence of mirror planes is neither imposed nor excluded.

To study all possible arrangements, with the built-in symmetry, we identify five special symmetric sets of particles that can be present or absent (see also Fig. 2.4 for a typical icosahedral face). The numbers A, B, C, D and E signal if the sets are present at all, and how many independent basis particles there are for the sets $\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D}$ and \mathcal{E} respectively:

- \mathcal{A} : 12 particles on the vertices of the icosahedron ($A = 0, 1$),
- \mathcal{B} : 30 particles on the midpoints of the edges between the particles of the icosahedron ($B = 0, 1$),
- \mathcal{C} : 20 particles in the centres of the faces of the icosahedron ($C = 0, 1$),

- \mathcal{D} : a variable number ($D \geq 0$) of basis particles on one half edge between two adjacent faces (only one half is used since these particles are rotated to the other half by the 2-fold rotational axis halfway every edge). The particles in this set do not have a predetermined position along the edge. Each basis particle in this set has two copies on each edge. There are 30 edges, thus in total 60 times the number of particles on one half of an edge are generated by the icosahedral symmetry operator,
- \mathcal{E} : a variable number ($E \geq 0$) of basis particles in one third of a face of the icosahedron (because of the 3-fold rotational axis through the centre of the face each particle is found 3 times in one face). The particles in this set do not have a predetermined position. There are 20 faces, so again 60 copies of each basis particle are generated by the icosahedral symmetry operator.

The positions of the particles in sets \mathcal{D} and \mathcal{E} are optimised by applying special implementations of SA and SD which preserve icosahedral symmetry.

The total number of particles with icosahedral symmetry is $N = 12A + 30B + 20C + 60D + 60E$, where A, B and C are zero or unity, D and E are an element of \mathbb{N} . Not all combinations represent well packed spherical distributions, however, especially the number of particles on an edge and in a face should be balanced. We expect that optimal arrangements have an approximately triangular structure in the faces. The numbers D and E should be adapted to that type of arrangements. As an example, in Fig. 2.4 we show a flat arrangement with $A=1$, $B=1$, $C=1$, $D=2$ and $E=3$, with particles that all should be mapped onto the spherical surface. The number of edge and face particles are in balance. To approximate triangular lattices, the number of particles in a face ($C+3E$) is related to the number of particles on an edge ($B+2D$) by

$$C + 3E \approx \frac{1}{2}(B + 2D)(B + 2D - 1). \quad (2.17)$$

After the system is optimised with SA and SD implementations which preserve the built-in symmetry, it is in an energy minimum while constrained by the symmetry requirement. As soon as the symmetry requirement is lifted, the particles might rearrange themselves. However, we obviously want the found minima with built-in symmetry to be true minima of the N particle system without symmetry restrictions. Therefore we apply the unconstrained SD method to determine if any found solution is indeed also a truly unconstrained stable arrangement.

2.9 Related work

The work of this thesis in many respects is built on—or inspired by—previous work, as found in the literature. In this section we briefly summarise those aspects that are relevant for our studies and the results obtained.




Figure 2.4: An example of a distribution of edge and face particles within one face. In this case $A = 1$, $B = 1$, $C = 1$, $D = 2$, $E = 3$.

2.9.1 Flat 2D space

Most of the published work on simulating flat 2D thermodynamic systems is devoted to the melting phenomenon. Specifically, many researchers try to verify the predictions of the KTHNY theory. In spite of this work still some controversies exist [19]. A main problem is related to the system size N and the number of steps used for equilibration and analysis. These are connected by the amount of computer power and time available. For example, the long-distance behaviour of the bond orientational order is restricted by periodic boundary conditions (PBC). It is found [19] that the system size is very important in testing the predictions made by the KTHNY theory. To resolve conflicting results, simulations are needed on sufficiently large systems with reliable equilibration and analysis. This is only recently possible [71] because of the enormous increase in available computing power. Other controversies arise from the choice of ensemble. Below we discuss work where the LJ potential is used in the canonical and isothermal-isobaric ensemble in flat 2D space with PBC.

In the canonical ensemble it is difficult to discriminate between two-phase coexistence and the hexatic phase as predicted by the KTHNY theory. When the coexistence region consists of patches of fluid and solid rather than a true phase separation, a slow decay of bond-orientational order emerges very similar to characteristics of a hexatic phase [72].

Frenkel and McTague [73] carry out a canonical ensemble MD simulation of

a 256-particle system and observe the existence of a hexatic phase in accordance with the KTHNY theory. In the same ensemble with a 1024-particle system, Tobochnik and Chester [72] find that the nature of the melting transition cannot be unambiguously determined from their data, but later Strandburg, Zollweg, and Chester [74] argue that the transition is first order, contrary to the predictions of the KTHNY theory. A first order transition is reported by Toxvaerd [75] for systems with up to 3600 particles and by Bakker *et al.* [76] for a system of 10864 particles. However, using finite-size scaling to analyse their data on a 12480-particle system, Udink and van der Elsken [49] conclude their results to be in reasonably good agreement with the KTHNY theory.

In the isothermal-isobaric (NPT) ensemble no region of coexistence should occur. Problems occur in implementing this ensemble correctly, and in discriminating between a discontinuous jump of thermodynamic properties at the transition (first order), and a steep but continuous rise (second order). Abraham [77] concludes that the melting transition is first order in his isothermal-isobaric MC simulations of 256- and 529-particle systems and in his isothermal-isobaric MD simulations of a 576-particle system. Chen *et al.* [71] performs MD simulations in the same ensemble on a 102400 particle system with extensive equilibration. Their results show signs of the hexatic phase but fall short of confirming the KTHNY theory.

In a recent study Recently Bladon and Frenkel [78] present unambiguous numerical evidence for the existence of a KT dislocation-unbinding transition in a 2D system with short-range attractive interactions. Contrary to expectation, they find this hexatic region far from the melting line. The solid they consider can exist in a high and a low density solid phase that have the same structure and are separated by a line of first-order phase transitions. The hexatic region forms due to the presence of a nearby critical point that terminates the line of first-order isostructural solid-solid coexistence. They show that in a finite region around the critical point the solid becomes unstable with respect to dislocation unbinding.

Only few studies deal with freezing in 2D systems. The first example of solidification in a 2D fluid with free boundaries is the study by de Wette *et al.* [33] on the crystallisation of a LJ system. In this simulation the temperature of a random array of 400 LJ particles, initially at $T = 0$, is allowed to evolve freely (a friction term is included in the equation of motion to control the temperature). Crystallisation occurs readily. Initially several crystallites emerge, which later merge to form an almost defect-free triangular lattice.

Since the work of Ramakrishnan [63, 64] researchers study the height of the first peak in the structure factor as a function of T . Using perturbative density functional theory Ballone *et al.* [79] study a 2D electron system and find that the first peak in the structure factor has a value of $S(k_m) = 4.73$ at the freezing point (in reasonable agreement with simulation [79]). Udink [49] observes for the LJ system, a value below the theoretically predicted value of 5.75, even below the

value 5.25 from Broughton *et al.* [80] in their simulation on the r^{-12} system.

2.9.2 Spherical 2D space

Distributions of points over the surface of a sphere where the positions are determined by symmetry considerations are studied by for example Föpple [7], Coxeter [81] and Wenninger [3, 82, 83]. In their studies only specific numbers of points can be used, and the generated arrangements are often highly symmetric. In our study we are not only interested in symmetric arrangements with specific N -values.

The optimal distributions of arbitrary numbers of particles with an interaction potential over a spherical surface are studied extensively by several authors. For a given interaction potential and a given number of particles, there are a number of equilibrium arrangements of the particles. In general, these arrangements depend on the interaction potential. Leech [84] finds equilibrium arrangements, which are independent of the interaction. It is shown that only arrangements that have rotational symmetry around every particle have this universal equilibrium property. Special examples of such arrangements are the regular polyhedra.

In the following we discuss simulation experiments of N particles on a spherical surface with Coulomb, hard-disc and logarithmic interactions. In most of the work found in the literature the Coulomb potential is used.

Coulomb

Hansen *et al.* [36] use SBC as an alternative for PBC in their simulation study on self-diffusion of a classical electron gas. They do not carry out crystallisation simulations. Two recent papers on energetically optimal arrangements of particles with Coulomb interaction are from Erber and Hockney [48] and Rakhmanov [56].

The extensive paper from Erber and Hockney [48] deals with optimal distributions of electrons on a spherical surface. They discuss a wide range of properties of this system. One of their main general observations is the emergence of complexity in N -body Coulomb systems for increasing values of N . Many of their findings are qualitatively typical for the general problem of optimal distributions of particles on a sphere with other than Coulomb interaction potentials. They find for increasing N that arrangements with non-vanishing dipole moments appear for $N = 11, 13, 19$ etc.; the first enantiometric, or chiral, arrangement appears at $N=15$; and robust metastable arrangements are encountered with increasing frequency for $N = 16, 22, 32, 35, 37\dots$. When $N \sim 100$, sets of 50-90 metastable arrangements—separated by energy differences of about 0.001%—are the norm. The capture basins of some metastable arrangements are larger than those of the lowest energy (stable) arrangement. Angular comparisons show that energetically similar local minima generally have completely different arrangements. Moreover, they conclude that detailed geometrical patterns of the charge distributions tend

to become increasingly irregular for larger values of N . Nevertheless, this coexists with complex statistical (averaged) order in the overall angular distributions, and isolated regular arrangements appear at a series of N -values extending up to $N = 112$. The dipole moments of the equilibrium arrangements they find (for $N \leq 112$) are bounded by 10^{-2} , whereas the average dipole moments of spherical random N -point distributions grow as $0.92N^{1/2}$.

Rakhmanov *et al.* [56] report on a study of stable arrangements of the Coulomb problem (as well as the logarithmic, see below, and an extremal sum problem, where the sum of all inter particle distances is maximised) for $N \leq 200$ and a few larger N -values. Arrangements with larger N than covered by Erber and Hockney are studied, including their symmetry. They study the general trend of the energy as a function of N and the deviations of the energy of individual N -values from this trend. It turns out that some groups of N -values have lower energies than the general trend while others have higher energies. They also observe that the 12 points whose Voronoi cells are pentagons tend to distribute themselves as far apart from each other as possible, with arrangements quite close to an icosahedron.

Other authors that use the Coulomb interaction in their studies on spherical distributions include: Zhou [31], Calkin *et al.* [85], Frickel and Bronk [86], Glasser and Every [87], Edmundson [70, 88], Weinrach *et al.* [89], Melnyk *et al.* [90], Waal [91] and Berezin [92]

Hard particles

The hard-disc interaction in 2D corresponds to the infinite order generalised Coulomb potential. Giarritta *et al.* [39, 45] and Tobochnik [93] study the emergence of ordering in hard disc systems under spherical boundary conditions (SBC). Giarritta *et al.* give thermodynamic and structural evidence of an ordering phase transition at high densities. They monitor the clustering of defects (in our language d-charges) in a hard-disc system for increasing density. The fraction of particles with zero d-charge first rises linearly with density until $\rho\sigma^2 \simeq 0.83$. Thereafter the fraction rises much more rapidly, indicating the outburst of triangular order in the system. Correspondingly, beyond $\rho\sigma^2 \simeq 0.86$ they observe a drastic drop in the number of particles with d-charge ± 1 . At intermediate densities most of these defects are grouped together, the number of isolated d-charges ± 1 never exceeding 2% of the total number of defects. The excess of d-charge '-1' over d-charge '+1' decreases for increasing densities, eventually tending to the value 12 as demanded by Euler's formula. It is found that d-charges ± 2 are a minor fraction and, at high density, appear only in a few distinct patterns. In particular d-charges -2 usually link to a particles with d-charge +1. The d-charge -2 head of the resulting d-dipole takes the place of a d-charge -1 inside a chain of alternating d-charges ± 1 . Instead, particles with d-charge +2 appear to play the role of "counter-defects" with respect to the d-charges -1 in that they bind

a triplet of d-charges -1 in a sort of starlike arrangement. For increasingly dense packing, d-charges other than ± 1 progressively disappear. It is shown that in the high density phase topological defects are not distributed at random over the surface but segregate into clusters that give rise to an upper level of organisation in the form of a hierarchical superstructure with icosahedral symmetry.

Other authors that using hard particles in their studies include Goldberg [94], Schreiner and Kratky [95] and Tarnai and Gáspár [96].

Logarithmic

The logarithmic—or zero-order generalised Coulomb—potential is used by Rakhmanov *et al.* [56, 97, 31] and Bergersen *et al.* [32]. A special property of the logarithmic potential is that the dipole moment in energy minima is necessarily zero. Bergersen *et al.* use this property as stopping criterion for their SD energy minimisation of the system.

The results of Bergersen *et al.* show that the configurations with minimum energy are in general different from those obtained with the Coulomb potential. However, the deviation of the energies of the solutions from the general trend in the energy as a function of the number of particles show remarkable similarities with the Coulomb case.

Lennard-Jones

The work presented in this thesis and the corresponding preliminary reports [98, 99, 100, 101, 102] forms—as far as we know—the first detailed study dealing with 2D LJ many-particle systems on a sphere. In particular we explore extensively the behaviour of N -particle LJ systems on a sphere for widely different N as a function of temperature as well as at zero temperature.

Chapter 3

Computational methods

In this chapter we discuss the methods used in the computer simulation experiments reported in this thesis. First a general introduction to the available classes of methods is given. We provide selection criteria against which the classes are evaluated. Next, given the selection criteria, we describe the functional aspects, theoretical background and implementation aspects of selected methods. A comparison is given between these methods and their behaviour when applied to the crystallisation process being studied.

3.1 Optimisation methods for crystallisation

In many physical, chemical and biological processes, certain properties of a system are intrinsically optimised. Such processes may be formulated as a simplified—often still complex—optimisation problem for a model system, suited for computer simulation. A majority of these problems involves the determination of the global—possibly degenerate—minimum of an irregular multidimensional (cost) function related to the degrees of freedom (coordinates) in the system. Each point in the multidimensional state space, spanned by the coordinates, represents a state of the system for which the cost function can be evaluated. The cost values form a landscape over the state space. Both the state space and the cost function may be continuous and/or discrete.

Only if the state space is discrete and not prohibitively large, the solution can be found by exhaustive enumeration. Usually the solution must be selected among too large (possibly infinite) a number of candidate states. More precisely: if the amount of work required for an exact solution grows more rapidly than a polynomial function of the problem size (see Chapt. 1) then—for large systems—it is in practice impossible to examine all candidate states. Clearly a sampling method is then required. At the same time the certainty of finding the global minimum is lost. A standard optimisation method can be used to find—from a given initial state—a sequence of states along a downhill path through the land-

scape until a minimum is reached (see Sect. 3.5). However, deeper minima may be present elsewhere in the landscape. Therefore such an optimisation method must be applied many times, with starting points spread over the state space, in order to find with confidence the global minimum. More clever optimisation methods have the possibility to escape from local minima during their search through the landscape towards the global minimum. We summarise the main criteria for choosing our simulation model.

The computer simulation at hand, for a model system of many particles on a sphere, preferably should mimic the corresponding natural system [103]. In natural many-particle systems—and similarly in their model simulations—phase transitions can occur where abrupt changes in the behaviour of the system take place. We hope to find the global minimum with an optimisation method that is closely related to the properties and dynamics of such a natural system. Then we can also study—and experiment with—the behaviour of the system during the optimisation. This can be achieved by modelling and simulating the crystallisation phase transition that for the natural system—in the most perfectly uniform case—would lead to the global minimum. The path taken by the simulation method to bring the system from some initial state to the global minimum provides information on the physics in the problem and can be used to fine-tune the simulation method.

In natural many-particle systems—and similarly in their model simulations—phase transitions can occur where abrupt changes in the behaviour of the system take place. The path taken by the simulation method to bring the system from some initial state to the global minimum provides information on the physics in the problem and can be used to fine-tune the simulation method.

Another criterion that can be kept in mind when choosing a simulation method, is that in the natural system constituent parts (modelled as particles) interact concurrently. Since the desired system sizes are large and the methods complex, we end up with a computationally intensive simulation. Therefore we would like to exploit this naturally implicit parallelism in the crystallisation problem, also as much as possible in the simulation model to obtain optimal computational efficiency. This may require exploiting locality in the system.

The criteria for choosing a simulation method, feasible execution time, quality of the found solution, behavioural resemblance to the natural system suitability for studying the optimisation process in the model itself, and achievable parallel implementations, leads us to the class of “natural solvers” [12].

These solvers are inspired by processes in nature. No complex transformation of the problem to an abstract mathematical model is used. The microscopic dynamics in the simulation model is defined such that it mimics essential parts of the microscopic dynamics of the natural system. Locality, for example provided by short range interactions, is here also explicitly present in the simulation model. Natural solvers in general have the property that the parallelism in the application can be mapped, via a well defined transformation onto a parallel

computer. Important examples of natural solvers are simulated annealing [104], genetic algorithms [105], artificial neural networks [106] and the lattice Boltzmann method [107].

3.2 The choice of simulation methods

In this section we motivate our choices of optimisation methods. The next section describes the chosen methods in detail. Our main goal is to find stable states with minimal potential energy for a classical 2D model system of N particles, interacting by two body Lennard-Jones (LJ) forces, confined to the surface of a sphere.

The total potential energy in the system, to be used as the cost function, depends on the coordinates of all N particles and the interaction potential used. The aim of the optimisation search is to find the global minimum in the landscape of the energy function defined over the continuous $2N + 1$ dimensional state space (two coordinates for each particle and one coordinate for the radius R of the sphere). In Sect. 6.4 it is shown that the number of local minima rises exponentially with N . For the state space to be discretised sufficiently fine, the number of state points should also rise exponentially. The complexity of the optimisation search is given by $S_{\min}^{-1}M(N)$, where $M(N)$ is the number of minima ($M(N) \sim \exp N$) and S_{\min} is related to the size of the capture basin of the global minimum. If all minima have equal size capture basins $S_{\min}=1$. A minimum in potential energy may have a very small capture basin resulting in a large factor S_{\min}^{-1} . In practice this excludes the possibility of examining all states for large systems.

Several natural solvers—as introduced in Sect. 3.1—are available for optimisation of such complex systems. We have chosen mainly for Simulated Annealing (SA) which is (thermodynamically) stochastic in nature and allows the escape from local minima. Moreover this method allows to study system properties at finite temperature, while approaching the zero temperature situation. We use the Monte Carlo (MC) method for equilibration. We also use a Molecular Dynamics (MD) method for equilibration. In addition we apply a non temperature dependent deterministic search method, Steepest Descent (SD), which is faster than SA but gets trapped in the first minimum it finds. This method can be used stand-alone, but—when integrated in a non-deterministic search—also has a great computational potential.

Below we introduce the four algorithms used in our work. First SA and SD are introduced, then MC and MD. A short discussion on the difference between the SA and SD methods and between MC and MD methods is given. Then SA and MC as used in our work are discussed.

Simulated Annealing

The Simulated Annealing (SA) method has been formulated by Kirkpatrick et al., [104], in 1982. They described a detailed analogy between annealing of a solid until it reaches its lowest energy (ground) state and the optimisation of the cost function of a large complex combinatorial optimisation problem.

In their approach an instance of combinatorial optimisation is formalised as a pair (\mathcal{R}, C) , where \mathcal{R} is the finite—or possibly countably infinite—set of states and C a cost function, $C : \mathcal{R} \rightarrow \mathbb{R}$. Our work closely follows the original motivation for SA. The set of states \mathcal{R} is here the state space, the cost function C is the free energy of the system.

The method is inspired by the physics of processes in nature and industry in which large single-crystals are formed through repetitive heating and subsequent cooling of a solid. Annealing of material can be accomplished with different thermodynamic schedules dependent on the properties of the material and the desired final state. An example is to heat the material, keep the temperature T fixed sufficiently to reach thermal equilibrium (TE), followed by cooling. To reach optimal crystallisation T is lowered very slowly.

Usually the heated material is in a heat bath where T is determined by an oven. The energy E of the material shows T -dependent stochastic fluctuations. At high T the particles of the solid move relatively freely with respect to one another and the system explores most of the state space. During cooling the system spends more and more time in deep capture basins, from which, however, it can still escape. When thermal mobility is lost, all particles tend to search individually for the lowest energy, consistent with the physical, geometrical and topological constraints. If the material is quenched, i.e. cooled fast without allowing the material to equilibrate, defects tend to be unnecessarily frozen into the solid. In the latter case the system is trapped in a local minimum rather than in the global minimum.

It can be shown [108] that under well formulated conditions SA is guaranteed to give the most optimal solution. However, it takes an infinite amount of time to satisfy these conditions. Therefore we have to use SA as a method to yield an approximately optimal solution in an acceptable amount of computation time. In our case SA gives a minimum in energy which however is not necessarily a global minimum (i.e. no other minimum being deeper).

In SA the size of computational temperature steps and the allowed equilibration have the same influence on the system as the temperature schedule in physical annealing. If equilibration is insufficient many runs are needed to find with confidence the global minimum. On the other hand, if a slow cooling scheme is used, less runs suffice, but each simulation requires a relatively large amount of computing time.

The SA algorithm is applicable to a wide variety of optimisation problems. One type of problem is—as in this thesis—the determination of the global min-

imum energy configurations of particle systems [109, 110, 111, 112, 113, 114]. Although SA is based on the thermodynamics of atoms or molecules with continuous coordinates, it can be—and often is—applied to combinatorial optimisation problems over a discrete state space, i.e. with discrete states. Examples are image restoration [115], load balancing [116, 117, 118] wiring of chips [119] and others [108].

The average quality of solutions obtained with SA is influenced by the shape of the landscape over the state space. It can be shown [120] that SA works well on a AR(1) landscapes. Then SA can work hierarchically; at each temperature, or energy scale, barriers and basins are present that are relevant to that energy scale and a good choice can be made to which basin the search will be restricted.

To find out if a function has such properties over its state space, Weinberger [121] proposed the autocorrelation function (ACF)

$$Z(d) = \frac{\langle (f(x) - \bar{f})(f(y) - \bar{f}) \rangle_{d(x,y)=d}}{\text{var } f}, \quad (3.1)$$

where f is the value and $\text{var } f$ is the variance of the function over the state space (in our case the potential energy landscape), $d(x, y)$ is the distance measure between the state space points x and y . The function $Z(d)$ can not be applied without adaptations to the state space of particles on a sphere. In that case we deal with a continuous state space, while Eq. 3.1 is used on discrete state space, where $d(x, y)$ is defined for example as the Hamming distance [122]. Here we use the Euclidean distance between the two positions of a particle that is perturbed with the MC method. Another problem is that in the original formulation random moves through state space are used. In our case, however, this mostly results in arrangements where some particles are very close to each other and thus have a large potential energy. This can be solved by putting a maximum to the energy between particles, which does however have a small influence on the shape of the correlation function.

The SA is a time consuming algorithm. For large optimisation problems the computational time involved often becomes prohibitive. Parallelisation of the algorithm is not trivial. The SA algorithm is inherently sequential. Our approach to parallelisation of the SA algorithm and the obtained performance are discussed in Chapt. 4.

Steepest Descent

The steepest descent (SD) method is fast in finding the "nearest" minimum in state space for given initial conditions. In our case it has been adapted for optimising the energy of a many-particle model system basically at zero temperature [123]. The SD algorithm uses the multidimensional derivative of the energy

landscape over the state space, here representing the force components for each of the particles, towards equilibrium. If the system is in an energy minimum all derivatives are zero. The consequence of using only the direction and size of the forces is that the system gets stuck in the first downhill minimum encountered when started in a given initial arrangement.

The SD method is comparable to using the Verlet MD method [25] where the particles are advanced one time step on the basis of the old positions, the velocity and the acceleration. There are two main differences between SD and MD. All velocities are put to zero for all time steps, such that only the present position and the total force acting on a particle are used to propagate the particle to a new position. This is equal to infinite quenching, implying zero temperature. The other difference is that we effectively use a variable time step. The size of the step in the direction of the forces is dynamically adjusted to reach optimal convergence.

Erber and Hockney [48], in a study of Coulomb systems, state that conjugate-gradient methods do not improve this technique because of the highly convoluted—far from quadratic—structure of the energy surface in state space. Near the bottom of the minimum, however, the energy function can be approximated by a quadratic function and then a conjugate-gradient method can be more efficient.

Simulated Annealing versus Steepest Descent

The important difference between SD and SA is that the latter is capable of escaping from local minima. In theory SA is capable of finding the global minimum. The SD method gets stuck in the capture basin to which the initial arrangement belongs.

Since we are dealing with a system that usually has many local energy minima this SD method can only replace SA if repeated many times with starting points spread all over the state space. We must realise, however, that even in high statistics computer searches minima with a very small capture basin can be missed.

In Sect. 3.6 a comparison of the average quality of the found solutions for LJ particles on a sphere between the SA and SD methods is given.

At the end of SA ($T \approx 0$) the system is definitely trapped in a capture basin. Then SA—due to its stochastic nature—it is not well adapted to quickly find the precise minimum. In that case SD is much more efficient. Therefore we apply the SD method to reach the final accuracy, needed for unique identification of the minima obtained with SA. The SD method is discussed in Sect. 3.5. In Sect. 6.4 we estimate the capture basin of “all” (meta)stable arrangements. There we apply SD while randomly sampling the state space. In that case SA would cause a strong bias for low lying states.

Monte Carlo

In a Monte Carlo (MC) simulation of an N particle system only the configurational part of the phase space, the state space, is used. The system performs a biased random walk through state space. In this walk regions of state space are sampled according to their “importance”.

With the Metropolis method [124], where randomly perturbed states are accepted with a probability involving the energy difference and temperature, the sampling probability is proportional to the equilibrium probability distribution of all possible states. Standard MC simulations are used to study the static (time or ensemble averaged) properties of systems at constant N, V and T . The method can easily be adapted for other ensembles [25]. For example, a constant pressure ensemble may be implemented by using changes in the volume of the system as attempted “moves” (perturbations of the system), and using the appropriate isobaric statistical weight to determine whether the moves are accepted or rejected. There is considerable freedom in choosing the moves. Any combination of moves may be used as long as the transition probabilities satisfy detailed balance (see Sect. 3.4).

In the work reported here we apply both the constant density ρ and the constant pressure P MC method at temperature T and constant density or pressure, to an arrangement of N particles confined to the surface of a sphere. In the constant pressure simulations there are two different moves: a particle move in which the position of a particle is perturbed, and a move in which the radius of the sphere is perturbed. Once in every N particle perturbations the radius is perturbed. For the particle move the potential energy difference is used in a Metropolis acceptance scheme. In the effect of the radius move not only the potential energy is involved but also P and V :

$$\Delta W = (U' - U) + P(V' - V) - NkT \log(V'/V), \quad (3.2)$$

where ΔW is the energy change used in the acceptance criterion, U the interaction energy and k Boltzmann’s constant. The values after the perturbation have a prime.

Although the MC method does not explicitly contain particle momenta, and thus leaves the corresponding degrees of freedom unexplored, it can provide a framework to calculate the thermodynamic properties [125]. All of the deviations from ideal gas behaviour are a consequence of interactions between the particles. These are described in the interaction potential function only depending upon the positions of the particles, and thus allow to be fully accounted for by the MC method. The ideal gas part with its distribution dependence can be calculated analytically.

Molecular Dynamics

Molecular Dynamics (MD) simulation amounts to numerical integration of the classical equations of motion for an assembly of N particles [18, 25]. The force on each particle can be evaluated by summing the contributions due to all other particles. For efficiency reasons the summing may sometimes be restricted to particles within a distance r_c from that particle, where r_c is a cutoff distance appropriate to the range of the force (see also [126]). Given the force on each particle the equations of motion can be integrated for a small, but finite time increment. The new positions are then used to evaluate the new forces. The cycle is repeated over and over, thus simulating the time development of the system in the available phase space, given the initial conditions, position and momentum for each particle.

In a standard MD calculation the total energy E , the volume V , and the number of particles N are held constant according to the microcanonical ensemble. If the system is ergodic, taking the time average of a dynamical property will be equivalent to taking the microcanonical ensemble average.

Although the MD method operates most naturally in the microcanonical ensemble [19], extensions of this method to other ensembles have been introduced [25]. For example, a canonical ensemble, in which N , V and the temperature T are constant, may be simulated through coupling to a heat bath. One method of implementing such a coupling is to occasionally randomise the velocity of a given particle according to a Boltzmann distribution. A constant pressure MD simulation method may be implemented by coupling the system to the equivalent of a “piston.” The mass of the piston can be chosen so as to produce volume fluctuations over a physically reasonable time scale [25].

Molecular Dynamics versus Monte Carlo

The important difference between MC and MD methods is that the particle moves in a MC simulation are unrelated to the natural time evolution of the system. Since the standard MC method probes only the state space no dynamical information can be derived from MC simulations. In MC runs “structural relaxation”, the changing of the positions of the particles towards equilibrium, and “nucleation”, the start of freezing, are the result of a succession of non physical moves. Consequently, structural rearrangements that are favoured by the MC method may be impossible to realize in a MD simulation, and vice versa [18].

The MD method is a deterministic method which allows in principle to predict the positions of particles at any future (or past) time. The MC method is stochastic, it does not follow any deterministic set of equations to move the particles, instead it has randomly fluctuating variables.

The MD and MC methods can be adapted to sample various ensembles. Most ensembles have been implemented in terms of both MC and MD, but the MC

implementations are usually simpler. The MC method offers more flexibility in implementation than the MD method [19] since the moves themselves do not have to obey any equations of motion.

In our work the MD method is only applied in a few special cases with the purpose of comparing it with MC.

3.3 The simulated annealing method

Simulated Annealing starts—analogueous to physical annealing—by equilibrating the model system at a high initial temperature T_0 . In Thermal Equilibrium (TE), the states are visited with a probability given by the Boltzmann distribution. At high T the fluctuations in the energy of the system are large compared to the energy variations of interest in the landscape. Large stochastic moves through state space are induced at high T . These are used to globally explore the landscape. After TE has been reached the T is decreased by a cooling step ΔT . Equilibration and a subsequent cooling step are repeated over and over. The system increasingly visits states with low energy more often. It loses the ability to explore the landscape globally and gradually tends to be confined to the low energy regions in the landscape. At low T the small energy variations of the landscape become more important. The SA process is continued until a stop criterion is met, for example reaching a stop temperature T_f .

The value of T_0 , the method of changing the temperature and the stop condition together fix the annealing schedule. In our simulations T_0 is chosen high enough to bring the N -particle system well above the liquid-solid phase transition T_c . The SA run is stopped when escape from the current capture basin has become practically impossible.

Thermal equilibrium is a key element. If the system reaches TE at all temperatures and the cooling is done sufficiently careful, it can be shown [108] that the system converges to the global minimum.

For equilibration two methods are applicable: molecular dynamics [114] and the often used Monte Carlo (Metropolis) method [124]. This latter method has been devised to study TE properties of large systems of interacting particles at a given temperature. It provides a simple algorithm directly applicable to our system. It makes a series of random moves through state space where each new move is only dependent on the current state, i.e. a Markov process. The sequence of states visited is called a Markov chain. In Sect. 3.4 the Metropolis method is discussed.

In the SA implementation we use homogeneous Markov chains. There is also an inhomogeneous Markov chain method where T is decreased after every Metropolis move. We have chosen to use homogeneous Markov chains because then we can study properties of the system as a function of T .

The proof that SA with homogeneous Markov chains, i.e. no parameters that

influence the random moves (e.g. T) are changed during the generation of a chain, is able to find the state with the lowest energy is based on reaching TE for all T . To theoretically ensure that the system is in TE, the Markov chains would have to be of infinite length. Additionally T should be decreased in small cooling steps ΔT such that the change in equilibrium distribution can be considered smooth. In the limit of an infinite number of chains, each of infinite length, while T approaches zero, the system converges [108] to the global minimum: the only state with an appreciable probability at very low T .

In practice we cannot apply an infinite number of moves to get the answer to our optimisation problems. Any practical application of the method has a finite number of Markov chains, each consisting of a finite number of Metropolis moves L , and a criterion to stop. As a consequence it is no longer guaranteed to find the global optimum.

The optimal ΔT depends on L . For instance, if at a certain T the system is in TE a large ΔT brings the system far from the new TE and a long Markov chain is needed. If ΔT is small the system is still close to the new equilibrium distribution and a small Markov chain may suffice.

Therefore the SA has to be repeated a sufficiently large number of times with different initial states such that the global minimum is among the results with acceptable probability.

In order to perform a first study of the dynamics for a spherical LJ system, and for comparison reasons, we have in addition used a MD method for our simulations. The MD implementation used is discussed in Sect. 3.5

3.4 The Metropolis Monte Carlo method

The method of Metropolis *et al.* [124] is used to equilibrate our model system during SA and to study its properties at temperature T .

The equilibrium probability of the system to be in state i , with energy E_i , at temperature T is denoted as $p(i, T)$ and is given by the Boltzmann distribution

$$p(i, T) = \frac{1}{Z(T)} \exp\left(-\frac{E_i}{kT}\right), \quad (3.3)$$

where $Z(T)$ is the partition function, k Boltzmann's constant and $\exp(-E/(kT))$ the Boltzmann factor of the energy. This distribution assures that states with $E_i > kT$ have a low probability compared to states with $E_i < kT$. With T the relative probability of the states can be controlled.

In TE the so-called detailed balance condition, or microscopic reversibility [25], holds

$$\forall i, j : p(i, T)\Gamma_{ij}(T) = p(j, T)\Gamma_{ji}(T), \quad (3.4)$$

where $\Gamma_{ij}(T)$ is the transition probability from state i to state j . This means that in a very large ensemble of systems [124] the amount of transitions from state i to j has to be equal to the amount of transitions going from j to i . Conversely, if this is true for all pairs of states, the system is in TE, where on average all states are sampled according to the probability function Eq. 3.3.

The Metropolis method is an importance sampling technique from a larger class of variance reduction techniques [127]; it samples important (probable) states more often than unimportant (improbable) states. In our case the state space is sampled in accordance with the Boltzmann distribution.

In the Metropolis method a random move to a new state is proposed at temperature T , on the basis of the current state. The probability that this new state is generated is given by $G_{ij}(T)$ which is determined by the implementation of the random move generator. The proposed state is accepted with a probability $A_{ij}(T)$ given by

$$A_{ij}(T) = \min\left[1, \frac{p(j, T)}{p(i, T)}\right], \quad (3.5)$$

where $p(i, T)$ is the Boltzmann probability of the current state of the system, and where $p(j, T)$ is the probability of the proposed state. The acceptance probability $A_{ij}(T)$ makes sure that states with an equal or lower energy than the current state are accepted. Proposed states with a higher energy have an acceptance probability depending on the energy difference ΔE_{ij} between states i and j . The Metropolis method randomly decides on acceptance or rejection, dependent on ΔE_{ij} and T . If the proposed state is rejected the old state is reused. Thus the system spends more time in states with a low energy than in states with a high energy. The random moves are repeatedly applied to generate a Markov chain.

The series of states in the Markov chain represent a random walk through the energy landscape with transition probability $\Gamma_{ij}(T)$ from state i to state j at temperature T . The $\Gamma_{ij}(T)$ probability is given by

$$\Gamma_{ij}(T) = G_{ij}(T)A_{ij}(T). \quad (3.6)$$

The acceptance probability is dependent on T because SA needs to simulate the system at a series of temperatures. In the generation probability, $G_{ij}(T)$, the size of the move through the state space is controlled (see Sect. 3.5 for implementation aspects). For efficiency reasons in a Metropolis simulation the ratio of the number of accepted and rejected moves should be around 0.5 (see for instance the discussion in [25]).

The $\Gamma_{ij}(T)$ transition probability is defined such that if an infinite Markov chain would be generated, the system converges to the TE distribution. To ensure convergence to the desired distribution the chain must be irreducible, i.e. all states must be accessible from all other states in a finite number of moves, and aperiodic, i.e. loops through the state space in which the system gets trapped should be absent (see [108] for mathematical aspects of these conditions).

3.5 Implementation aspects

Simulated Annealing

In the previous sections the SA and Metropolis methods are discussed in general. In this section we discuss their specific implementation for the crystallisation experiments concerning our N -particle LJ system on a spherical surface.

For the SA implementation we need a starting temperature T_0 , a cooling step rule and a stop criterion. In most of the computer experiments we use the following cooling step rule

$$T_{k+1} = cT_k, \quad (3.7)$$

where T_k is the temperature of the current chain, T_{k+1} the temperature of the next chain, c the cool-rate constant (usually taken 0.95, which is a value that gives good results for many different problems [108]). Other cooling step rules can be found in [128], see also Sect. 4.2.

In our simulations we have two possible stop criteria. The first is a stop temperature T_f given at the start of the run. The annealing is stopped when T_f is reached. The second criterion monitors the standard deviation in the energy of the last state in the last m (typically $m=10$) Markov chains. As soon as this deviation falls below a specified value, typically 10^{-6} , it is assumed that the system has been definitely trapped near a—possibly local—minimum. This second stop criterion, used in the next chapter on parallel implementations, is replaced by the first when SD is used in conjunction with SA.

In the simulations discussed in this thesis we use mainly the LJ potential, but also in some studies the Coulomb potential. For both types of potentials we need a perturbation mechanism for the coordinates of the particles on the spherical surface. For the Coulomb potential only the particle positions are essential degrees of freedom in the system, see Sect. 2.3. For the LJ potential we need an additional perturbation mechanism to change the radius of the sphere. We first discuss the particle perturbation mechanism and then the radius perturbation.

After a series of tests we have chosen to perturb one particle per Metropolis move. For this we need only to calculate the energy of that particle before and after the move. Displacement of all particles at once did not give better results (data not shown). We use a 2D displacement in 3D space, where a spherical isotropic randomly generated vector $d\mathbf{k}$ is added to the coordinates of the particle in 3D. The size of the random vector is controlled by a parameter L_g , the current maximum perturbation distance. The new position is rescaled such that the particle remains located on the surface of the sphere. The projection of possible random vectors on the surface of the sphere causes the smaller perturbations to be favoured.

After a Markov chain is generated the acceptance ratio for that chain is calculated. It is used to determine a new value of L_g in order to accept about 50%

of all moves in the next chain. The new value is calculated by

$$L'_g = L_g c(a + 0.5), \quad (3.8)$$

where a is the acceptance ratio and c is the cool-rate. Two mechanisms are combined in this equation. If T decreases it is efficient to make smaller perturbations through state space. Therefore we can expect that—during annealing—the next Markov chain has a smaller L_g . This is achieved by incorporating the cool-rate in the calculation of L_g . The second mechanism uses the acceptance ratio. If in a chain the acceptance ratio is high it means that larger moves through state space can be taken. If the acceptance ratio is low we need to make smaller moves. This is achieved by multiplying L_g with $(a + 0.5)$. Using this method we can keep the acceptance ratio near 50% during the simulation at different temperatures. This method can work down to very low temperatures.

The variable radius R in the simulations with the LJ potential makes the state space slightly larger than in the fixed radius (Coulomb) case. In spite of the slight increase in the number of variables, many more local minima show up (see Sect. 6.2). This is because R is a collective variable influencing the state of the system globally. The larger state space takes longer to explore. Therefore we wish to keep the values of the radius within reasonable bounds.

By experimentation (see Sect. 6.3.2) we find that the radii for the configurations with minimum energy as a function of N are not much different from the value based on the packing on a flat 2D triangular lattice, given by

$$R_0 = \sqrt{\frac{N \cos(\frac{\pi}{6})}{4\pi}}. \quad (3.9)$$

Since a spherical distribution can not be as closely packed as a flat distribution, the radii of the optimal spherical distributions are always larger than R_0 . For $N > 40$ the boundaries $R_0 < R < R_0 + 0.06$ are sufficient, for smaller N a larger upper bound is needed (see Fig. 6.15 in Sect. 6.3.2).

The radius perturbation is implemented by moving R with a random amount between $-\frac{1}{2}\%$ and $+\frac{1}{2}\%$ of its current value while staying in the above given bounds.

During the generation of a Markov chain once every N moves, R is perturbed. In the first series of chains at high T of the annealing run, R is not changed because the system is then in a highly disordered state and particles can be close to each other, giving very high (repulsive core) energies. In that situation it would be favourable for the system to increase R to its maximum. This would not be an optimally efficient start in the search for the global minimum.

For the acceptance test of any perturbation we need to calculate the energy difference due to the move. This is the most time consuming part of the simulation. Therefore we use a method that can speed up the calculations considerably. The LJ potential can be approximated using a cut-off length (see Sect. 2.2.) The

number of interacting particles is then greatly reduced. The computational complexity of calculating the total energy is lowered from $\mathcal{O}(N^2)$ to $\mathcal{O}(N)$, that of the particle energy from $\mathcal{O}(N)$ to $\mathcal{O}(10)$, depending on the cut-off length. A list of particles positioned within each others range is maintained. This list is constructed by taking a slightly larger range such that it remains valid for a number of moves, determined by the distances travelled in the accepted moves. At high T large perturbations are possible and the list has to be updated frequently. In that case it is more efficient to postpone the use of the list until T is lowered sufficiently. At low T , where small moves are taken, the list needs to be updated only occasionally. During SD the exact LJ potential is used to make sure that the final result has full precision. For the Coulomb potential—because of its long range—no cut-off is applied and all the interactions are always taken into account.

To reduce the problem size it is possible to remove some degrees of freedom without excluding any minimum energy configuration. One particle can be fixed on the surface of the sphere, $N - 1$ particles are free to move. The free particles have to arrange themselves with respect to the fixed particle. To reduce the problem size still further it is possible to remove the freedom of rotation around the axis of the fixed particle. This can be achieved by keeping a second particle on a specific great circle around the sphere through the fixed (first) particle.

Molecular Dynamics

In this section we briefly introduce discuss the MD method as implemented [129] for our system.

For particle systems with interactions via continuous potential energy functions, Newtons laws of motion can be integrated. This is achieved by a finite difference approach where (time)step by step the positions of the particles as a function of time are obtained. In our case we deal with the restriction that the particles should remain on the 2D spherical surface. One of the methods suited for solving the equations works by using quaternion parameters [25]. For our system, however, this results in redundancy and therefore is not efficient. A simple available algorithm corresponding the equations given in [25], which describe the behaviour of “linear” molecules where several particles making up a molecule are rigidly bonded in a single line. Our implementation boils down to making a (pseudo) dimer molecule, two particles at a distance equal to the diameter of the sphere. One of the particles is fake and does not take part in the calculation of the interactions. The centre of mass of the molecule is always fixed at the centre of the sphere, taken as the origin. Now N of those dimer molecules are placed with their centres of mass at the origin. The result of this trick is a system consisting of N particles confined to the surface of a sphere.

The rotation of a linear molecule is caused by the torque acting on it. If \mathbf{e} is the unit vector along the axis of the molecule we can write: $\boldsymbol{\tau} = \mathbf{e} \times \mathbf{g}$ where

$\mathbf{g} = \sum_a d_a \mathbf{f}_a$ and d_a are the positions of the particles relative to the centre of mass in the direction of \mathbf{e} . Only the part of \mathbf{g} perpendicular to the molecular axis, \mathbf{g}^\perp , is necessary. The equations to be solved are,

$$\begin{aligned}\dot{\mathbf{e}} &= \mathbf{u} \\ \dot{\mathbf{u}} &= \mathbf{g}^\perp / I + \lambda \mathbf{e},\end{aligned}\tag{3.10}$$

where I is the moment of inertia and the force $\lambda \mathbf{e}$ constrains the bond length to be a constant. An expression for λ is obtained by considering the advancement over half a time step. The equation to advance a full step in the integration then follows as

$$\begin{aligned}\mathbf{u}(t + \frac{1}{2}\delta t) &= \mathbf{u}(t - \frac{1}{2}\delta t) + \delta t \dot{\mathbf{u}}(t) \\ \mathbf{e}(t + \delta t) &= \mathbf{e}(t) + \delta t \mathbf{u}(t + \frac{1}{2}\delta t).\end{aligned}\tag{3.11}$$

The above formulas are applied to each molecule.

The radius R of the sphere has to be able to fluctuate in our implementation. This is achieved by simulating the action of a piston. The force acting on the piston is equal to \mathbf{F}^\perp / M where \mathbf{F}^\perp is the total force perpendicular to the surface of the sphere and M is the mass of the sphere, N times the mass of one particle. In our case zero external pressure is used.

Steepest Descent

A fast minimisation method is Steepest Descent (SD). We have developed an efficient implementation based on an adaptation of the method described in [123]. This method can be applied to functions for which the (multi-dimensional) derivative—here of the potential energy—can be calculated. In our case these derivatives are the forces that the particles exert on each other. The total force acting on the particles, points in the direction of the steepest downhill slope at the current position of the system in the energy landscape. The trial position \mathbf{r}'_i for particle i is based on its old position \mathbf{r}_i and the total force \mathbf{F}_i acting on it, like

$$\mathbf{r}'_i = \frac{\mathbf{r}_i + \gamma \mathbf{F}_i}{|\mathbf{r}_i + \gamma \mathbf{F}_i|} R.\tag{3.12}$$

If γ is large the original position is not important and the trial position is effectively the force vector rescaled such that the trial position is a point on the surface of the sphere. If γ is small the trial position is the current position plus a small step in the direction of the total force. First a large γ is used. Simulation experiments have shown that typically $\gamma = 20$ is a good starting value, higher values for γ frequently result in trial positions that have a high potential energy. If the move to the trial positions of the particles does not lower the total energy,

the scale of the step is lowered by decreasing γ . This is repeated until either a) a configuration is generated that does have a lower energy or b) the value of γ has become lower than a specified stop value (typically 10^{-10}). If this method is used to optimise functions with more than one minimum, the resulting arrangement depends on the starting point of the minimisation and on the sizes of the steps that are taken (the method might step over a minimum into the next if the forces are large in the initial configuration).

For the LJ potential the radius R of the sphere is optimised as well. This is achieved by summing, for all particles i , of the inner product of their net force with the unit vector in the direction of the position:

$$S = \sum_i \frac{\mathbf{r}_i}{|\mathbf{r}_i|} \cdot \mathbf{F}_i. \quad (3.13)$$

The inner product gives the part of the force that points perpendicular to the surface. If this inner product is negative the particle favours smaller R . The sum S indicates whether the system can decrease its energy E by enlargement or shrinkage of the sphere. This can be used to control the size of the sphere.

In our SD implementation the optimisation moves for the particles and for the radius are repeated until both fail to further optimise the energy. The final energies of the configurations are accurate to 12 digits. This is necessary to discriminate between nearly degenerate local minima. For the LJ potential the final R -values are less accurate (up to 8 digits). This is because in a stable arrangement, where most neighbouring particles are in the LJ minimum, a small change in radius has a relatively small effect on the energy. In the SD method the system can only move in the direction given by the forces. No uphill move is possible. Therefore the system is stuck in the first downhill minimum along the path given by the accepted optimisation moves.

3.6 Tuning and comparison

In this section we discuss how we determine the settings of important parameters used in the SA method. The parameters are the length L of the Markov chains and the start- and stop temperatures T_0 and T_f respectively. A comparison of SA with SD is given and their behaviour on the potential energy landscape of our model is discussed.

Simulated Annealing

The chain length L needed to obtain a certain quality of solution, increases with N , thus $L=L(N)$. We study the results for N ranging from 20 to 100 over many experiments for different L at one fixed cool-rate. Next we determine which

L corresponds to a specified quality for each tested N value. We observe that the length increases at least as N^3 but probably exponential (data not shown).

The values T_0 and T_f are chosen independent of N . Specifically we suppose that the phase transition temperature T_c is independent of N . For small N this might not be fully justified (see Chapt. 5), but there the landscape over the state space is still simple enough such that the used temperature trajectory results in the global minimum.

The annealing process should start well above T_c . Then the system equilibrates to a situation where the whole state space can be explored. The value of T_c is determined a.o. from system properties such as the average potential energy as a function of temperature (see Sect. 5.2), and the radial distribution functions (see Sect. 5.3). Our tuning simulations start with T_0 between 5 and 10. In experiments discussed below to find an appropriate T_f , we observe that T_0 the system is not trapped in one capture basin.

The T_f is determined for many independent experiments as follows. During a SA run, at each T , we take the last configuration of the Markov chain. Then, separately from the SA run, we apply the SD method to find the current capture basin. Then, for each T , we determine how often the system at the next lower temperature resettles in another capture basin. At high T the system frequently changes capture basin, at low T the arrangement is stuck in one basin. From Fig. 3.1 we can see—as an example—that for one run with $N = 14$ at high temperatures, transitions take place from one basin to another ($N=14$ has three minima near the radius of the global minimum), while at low T almost no transitions take place. If for each T the transition probability is plotted, we see (Fig. 3.2a) that below $T = 0.01$ few transitions take place. In fact we take $T_f=0.01$ as the lower bound for our annealing schedule. The distribution of T -values at which the last transition has taken place is given in Fig. 3.2b. It can be concluded that some runs, or at least the last configurations of the Markov chains, stay in the capture basin of one and the same minimum from $T = 10$. Most of the runs decide around $T = 0.3$ what minimum they will end in. From a similar experiment for $N=75$ the same conclusions can be drawn.

The quality of SA simulations is influenced by the anneal schedule parameters. Here “quality” is defined as the fraction of runs that end in the global minimum, in fact the deepest found. In order to give a comparison between SA and SD we keep the cool-rate fixed to 0.9 and study the quality as a function of the chain length L . On one hand SA should give results of higher quality than SD since it is able to escape from local minima especially for larger L . On the other hand it takes a longer execution time to obtain the solutions. To find which method has higher quality as a function of computational effort we performed a series of 10 sets of 100 simulations with SD and SA for $N=14$.

In Table 3.1 the results of the quality comparison are listed. The first row is a SD run without SA, therefore the number of moves is zero. The first column is L , the second and third column list the quality and average execution time per run.




Figure 3.1: Transitions between capture basins of the three minima for $N=14$ during one SA simulation. The horizontal lines indicate the energies of the minima.




Figure 3.2: (a) The fraction of runs that showed transitions to another capture basin as a function of temperature during SA and (b) the distribution of temperatures where the last transition took place. Below $T = 0.01$ only few runs escape from a capture basin.

The measurements are performed on a Sun Ultra 1 workstation. The last column gives the time needed to achieve the same quality for all listed simulations. The numbers in this column are obtained by dividing the time per run by the quality.

L	quality	time (seconds)	time/quality
0	0.37 ± 0.04	0.81 ± 0.01	2.2 ± 0.2
14	0.58 ± 0.05	0.98 ± 0.02	1.7 ± 0.2
28	0.58 ± 0.05	1.05 ± 0.02	1.8 ± 0.2
42	0.60 ± 0.06	1.21 ± 0.02	2.0 ± 0.2
70	0.53 ± 0.04	1.43 ± 0.01	2.7 ± 0.2
140	0.48 ± 0.05	2.17 ± 0.01	4.6 ± 0.5

Table 3.1: Comparison for $N=14$ between SA and SD ($L=0$).

Naturally the average execution time rises with L . The quality for SA runs is higher than for SD runs. However, for the SA runs—in this simple example—the quality does not rise as a function of the number of Metropolis moves per chain. A minimum in time/quality ratio at $L=14$ appears, indicating that this small L for $N=14$ already provide sufficient equilibration for SA to indeed reach a level of quality in a shorter time than SD.

Since the necessary L probably grows exponential with N , we can expect that for large N the minimum L with the best quality v.s. execution time also grows.

Next we try to understand how SA and SD behave when applied to our simulations, in terms of the energy landscape.

It can be shown that SA does not give good results for any shape of the landscape over the state space. Here we study this for a simple case. We take a simple landscape that consists of two parabolas, one with its minimum at -1 and one with its minimum at 1. The parabola with its minimum at -1 can be displaced in height compared to the other and the height of the barrier between the minima can be controlled:

$$\begin{aligned} H(x+1)^2 + D & \text{ if } x < 0. \\ H(x-1)^2 & \text{ if } x \geq 0. \end{aligned} \tag{3.14}$$

The height of the barrier to escape from the, for $D > 0$, locally minimal parabola centred at $x = -1$, is equal to H while the difference in height between the minima is equal to D . We study the influence of H and D . Since we can expect that the quality improves if D increases and if H decreases, we study the behaviour of the annealing as a function of the ratio D/H .

The results shown are determined by averaging over 250 experiments. In Fig. 3.3 we can see that at a D/H ratio above 0.1, the annealing parameters




Figure 3.3: The fraction of runs that ended in the global minimum as a function of the energy ratio D/H . For the cool-rate the values 0.85 and 0.95 and for L the values 10 and 80 are used, respectively.

start having influence. In the figure two different L -values and cool-rates are used. If the allowed maximal perturbation is large enough to bring the system from one minimum to the other in one move, we find a much higher quality (data not shown). However in our crystallisation experiments we do not use such “enormous” moves. Besides, in a multi-dimensional state space with a maximum perturbation size large enough to go from one minimum to another, the chance on generating such a move is very small.

These experiments indicate that the shape of the landscape is important to explain the quality as a function of the annealing schedule. For some landscapes it is in practice not possible to obtain higher quality if the anneal schedule is improved within reasonable limits. Of course, if an infinite number of steps in the Markov chain are taken and the temperature decreases to zero infinitely slow, every barrier is crossed and the quality should be unity.

Sorkin [120] has shown that annealing on hierarchical or fractal landscapes works fine. Stadler has shown [130] that the autocorrelation function (ACF), defined in Eq. 3.1, of the landscape for combinatorial optimisation problems that can be tackled reasonably well with SA, are exponentially decaying functions of the “distance” travelled through the state space. A short correlation length corresponds to many minima. In our crystallisation problem we work with continuous variables. This makes a direct comparison of the crystallisation ACF with those of combinatorial optimisation functions troublesome.

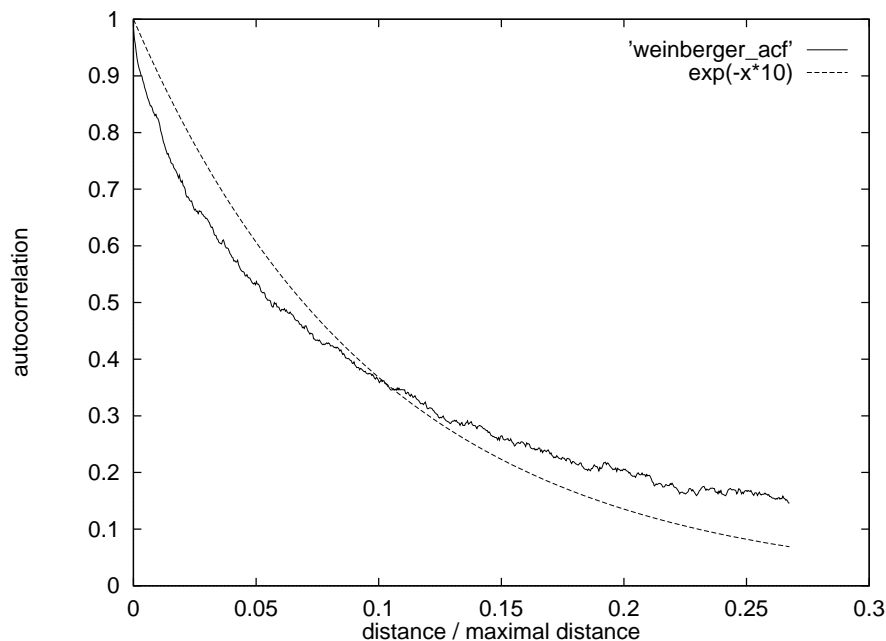


Figure 3.4: The autocorrelation function for the crystallisation on a sphere problem ($N=20$).

In a random walk through a landscape with infinitely high peaks the walk now and then climbs up the peaks and ruin the computer calculation of the ACF. Therefore we introduce a minimum distance, which determines a maximum energy, between the particles. This flattening of the peaks does however influence the behaviour of the ACF (data not shown). From the ACF in Fig. 3.4 it can be seen that at short distances the ACF falls of more rapidly than an exponential function while at larger distances the correlation falls off slower. The same behaviour is observed if the ACF is calculated on the basis of a Markov chain generated by the Metropolis method. This means that the state space landscape for our crystallisation problem is not AR(1) as in the combinatorial optimisation problems. This is also in accordance with the results shown in Chapt. 6 where find many nearly degenerate minima in a landscape between infinitely high peaks. In a AR(1) landscape it can be expected that also minima near relatively high energies are present.

The speed with which the SD method finds a minimum in the energy landscape is also dependent on the shape of the landscape. One of the problems with this method is that it tends to crawl slowly through narrow channels in the landscape [131]. Although Erber and Hockney [48] argue that a conjugate-gradient method does not improve the implementation as discussed in Sect. 3.5, we find that in some cases the SD method generates an unreasonable number of steps in order to find the minimum in the landscape. An example of this is taken from


the research discussed in Sect. 6.2. There, starting with an initially nearly flat configuration on a very large sphere, the curvature is increased in steps. After every step the system is brought to a minimum by applying the SD method. If we do this for $N=6$ we find the behaviour as shown in Fig. 3.5 and Fig. 3.6. There we show the energy, the number of necessary SD steps, the Euclidean distance (see Eq. 6.1) travelled through state space and the energies of the six particles as a function of the curvature. For a discussion on the results of the experiments see Sect. 6.2, here we only discuss the behaviour of the SD method.

In Fig. 3.5b it is shown that around $R \approx 1.75$ the SD method needs a large number of steps to reach a minimum in energy. This is remarkable since for energy as a function of the radius, Fig. 3.5a, and for the energies of the particles, Fig. 3.6b, no obvious changes appear. If the Euclidean distance, Fig. 3.6a, is compared with the number of SD steps we see that no large distance in state space is travelled around $R \approx 1.75$. The conclusion is that there the SD method is trapped in a narrow channel for some time.

We also studied the effects of the initial arrangement with which the SD method starts. In Fig. 3.7 we show the energy gained in the SD steps. If SD is applied to a random arrangement, Fig. 3.7a, we see an erratic behaviour. If first SA is used to relax the initial random arrangement we see, Fig. 3.7b, that the behaviour is simpler. These experiments concern a particle system at a radius larger than that for the optimal arrangement, see Sect. 6.3. If the same experiment is conducted on the same system but now at the density for which the global minimum in the landscape is found, we find that SD behaves much better there, see Fig. 3.7c.

Thus the quality of the SD method depends strongly on the type of landscape and the initial configuration that determines which path through phase space is taken (and thus how many narrow channels are explored).

(a)



(b)





Figure 3.5: The energy (a), the number of SD steps (b), the Euclidean distance travelled (c) and the energy per particle (d) following a finite increase in curvature (1%) for $N=6$ with SD relaxation. The system is started from an effectively flat arrangement towards smaller radius.

(a)



(b)






Figure 3.6: The same as Fig. 3.5 for the Euclidean distance travelled (a) and the energy per particle (b).

(a)



(b)



(c)




Figure 3.7: An example for $N=50$ of the behaviour of the SD method on a random arrangement (a) and when first SA is used (b). Both for a system at a large radius (2.88). If the radius of the global minimum (1.90) is used a simpler behaviour is found (c).

Chapter 4

Parallel algorithms

In this chapter we first discuss the various forms of parallelism that can be applied to Simulated Annealing with the Metropolis method and which architecture paradigms are used. Then it is discussed how these forms can be incorporated in the inherently sequential simulation method. The time complexities and scaling properties for the different types of parallel methods are determined, and a comparison between the methods is given.

4.1 Possibilities for parallelism

We use the Simulated Annealing (SA) method to find configurations with minimal potential energy for a system of N particles on a spherical surface, and to study the crystallisation process towards these configurations.

The SA method, in a practical implementation, does not guarantee to find these configurations (see Sect. 3.3). Therefore, for each N value, we have to perform a large number of independent stochastically initiated simulations. The amount of simulations required to find the lowest energy configuration with an acceptable probability is a function of N . For small N few simulations generally suffice. From literature [48] we know that for the Coulomb potential the number of nearly degenerate minima in the system rises exponentially with N . In Sect. 6.4 we show that this is also the case for the LJ potential. The required number of simulations therefore can be expected to rise also exponentially.

Not only are more simulations required for increasing N , they also require more computing time each. The effectiveness of the SA method depends mainly on how well the thermal equilibrium (TE) distribution is approached in each Markov chain. The amount of Metropolis moves needed to come close to the equilibrium distribution takes for large N a much longer chain than for small N since the dimensionality and irregularity of the landscape over the state space explored grows rapidly with N .

There is a balance between the number of independent simulations needed

and the length of the Markov chains (number of Metropolis moves). If a short chain length is used (insufficient equilibration of the system) the SA quality is low, and many simulations with different starting conditions are required. For longer chain lengths the chance of finding the global minimum is higher and less simulations suffice, but each takes more computing cycles.

The amount of computational effort needed to achieve the aims is enormous (for a detailed time complexity analysis see Sect. 4.3). Executing this programme sequentially on one machine, is out of the question.

In parallel execution the work is split into tasks which are distributed over a number of processors. At the cost of some programming effort the “wall clock” time needed for the total execution of the tasks can be much lower than for sequential execution.

Different paradigms of parallel computing exist. For example, a main distinction can be made between distributed- and shared memory parallel machines. In distributed memory machines each processor has its own (local) memory. High speed communication links using a message passing paradigm facilitates communication between processes on different processors. On a shared memory machine the communication between processors takes place by leaving messages in main memory, common for all processors. If workstations are connected together via a network they are effectively a distributed memory parallel machine with a slow communication network. In this work we also consider the usefulness of a vector supercomputer, where individual but powerful machine instructions are executed in parallel.

In our work we choose to use distributed computing. With this more general approach it becomes possible to use distributed memory parallel computers as well as general purpose—of the shelf—hardware such as clusters of workstations in a network.

We use two types of computing environments for our work. In a High Performance Computing (HPC) environment the focus is on processing a single job as fast as possible. This is achieved by using parallel computers or vector supercomputers. We also use High Throughput Computing (HTC) where the focus is on processing as many jobs as possible in a given time span. In the section on Distributed SA (Sect. 4.2) these two environments are discussed further.

The simulation method used allows for several levels of parallelism. On a coarse level we have job parallelism. Within one job we can apply algorithmic parallelism. On a finer level of parallelism within one job we can use data parallelism. On the finest level we have instruction parallelism such as used in vector supercomputers. In general, the execution of a program may involve a combination of these levels, dependent on for example the application, algorithm and hardware limitations. Below we discuss in some detail these four [132, 133] levels of parallelism in relation to the simulation problem at hand.

Job parallelism

Job parallelism amounts to running the same program, with different input parameters, concurrently on a parallel multiprocessor machine or on a cluster of independent computers. We show that for our work it provides a 100% efficiency and is therefore of fundamental importance. But also in general it becomes more and more practical to use this form of parallelism due to improved networks between computers, standardisation of soft- and hardware and an improvement in resource management [134, 135]. This form of parallelism is especially interesting for usage in a High Throughput Computing (HTC) environment where the turn around time for one job is not of main concern but rather how many jobs can be completed over a possibly large time span. Job parallelism can be combined with the other forms of parallelism discussed below. However, as discussed in Sect. 4.2 we use the sequential implementation.

Algorithmic parallelism

In algorithmic parallelism the algorithm is rewritten to allow for parallel execution. Usually communication of results calculated on different processors in parallel, becomes necessary. The ratio of time spend on communication and on calculation becomes a central issue. This ratio is dependent on the chosen grain size and on the bandwidth of the communication network between the processors. If communication is intensive, a distributed memory parallel machine with high bandwidth connections between the processors or a shared memory parallel machine has to be used. Besides communication time there may also be a synchronisation time. This is the time processors have to wait until the slowest process has finished before being able to process further data. The synchronisation time is mainly dependent on how well the load is balanced [132].

Data parallelism

In data parallelism a large number of similar and independent calculations is split into parts (decomposed), and distributed over the available processors. Usually the results have to be combined afterwards. This computational scheme is known as farming. A master processor controls the slave processors in the processor farm.

Typically this type of parallelism shows frequent communications between relatively short calculations, and therefore is best suited for a distributed memory parallel computer with high bandwidth communication channels or for a shared memory parallel computer. Besides the time spend on communication a synchronisation time may also be of importance if the load is not well balanced.

Instruction parallelism

On the finest level of parallelism the instructions given to a processor are partitioned in steps which are executed in parallel. Often a series of calculations have to be performed on a set of data. Whenever these data can be casted into a vector it can be processed by special hardware present in vector supercomputers. Vector instructions are then issued to functional units which operate on the vector concurrently, thus very fast.

4.2 Parallelisation of the simulation code

In this paragraph we show how the four parallelisation strategies, distributed SA (job parallelism), parallel SA (algorithmic parallelism), Metropolis with parallel energy calculations (data parallelism) and vectorised SA (instruction parallelism) are incorporated in the simulation method and used for our experiments [100].

The SA method is inherently sequential. It consists of a series of Markov chains where each chain is based on the previous. Therefore parallelisation of this method is not straightforward. Functional differences have to be introduced. As a consequence the parallel algorithm differs from the sequential algorithm. A number of different parallel SA implementations exist [136, 137, 138].

The Metropolis method, used to generate the Markov chains, is also inherently sequential. It generates a sequence of points in state space, each move to a new point is based on the current. Still a number of parallel Metropolis algorithms exist [139, 140, 141].

Several authors [142, 143, 144, 136] claim that they have constructed a parallel SA implementation while in fact they only provide a parallel Metropolis implementation for equilibration of the system during annealing. We make a clear distinction between parallel SA and parallel Metropolis. A similar distinction on a deeper level is applicable to the Metropolis method itself. In parallel Metropolis moves are performed concurrently. It is also possible to use sequential Metropolis moves with parallel executed importance difference calculations.

Distributed SA

In distributed SA the method itself is not parallelised but a substantial number of sequential jobs is executed at the same time on as many different processors. This is job parallelism where in fact we parallelise the parameter space. Most of the results reported in Chapt. 5 and 6 were obtained using this type of parallelism. In Sect. 4.3.1 the time complexity analysis of the sequential implementation is discussed.

For our simulations each job differs at least in the combination of the value of N and the value of the random seed. These jobs individually are not time

critical and are each processed sequentially on one processor, though using many processors at the same time.

Then the focus is not on High Performance Computing (HPC), where a large number of floating-point operations per second is of primary concern, but rather on High Throughput Computing (HTC), where sustained production (i.e. floating-point operations per year) is important. In HTC environments large amount of processing capacity can be delivered over long periods of time. The Condor distributed batch system [145] developed at the University of Wisconsin-Madison is an HTC environment that can manage very large collections of distributively owned workstations.

With this system hundreds of jobs can be submitted on a cluster of workstations. One function of Condor is to monitor the machines in the cluster. If a machine is idle a waiting job is allocated to it and started. By using “checkpoints” of running jobs it is possible, as soon as the “owner” reclaims the machine, to stop and migrate the job to another, still idle, machine. It is found that workstations are idle typically 60% [145] of the time. The Condor system thus turns otherwise wasted computing cycles into productive work.

As discussed in Sect. 3.2 for our work we have to perform an enormous amount of independent jobs. We of course would like to have these jobs processed as fast as possible. There was no chance, however, to get sufficient access to parallel machines (HPC environments) to do all the work. The parallel machines accessible to us have to be shared with many other users. Claiming a substantial part of these machines for a total time of possibly years is not desirable. However, for time critical jobs, involving software design, parameter tuning or very large N -values, the parallel machine is used.

Since the processing of a large number of jobs in a few years is in our case more important than processing single jobs very fast, HTC is better suited for our needs. As stated above the Condor system can deliver a large amount of CPU-time without annoying other users. There is, however, competition from other Condor jobs.

To obtain the very large amount of required computer resources we participated in a number of initiatives on distributed computing development. This resulted in a cooperation on Condor Flocking where several Condor clusters are connected to form one large cluster of computers. Jobs can migrate from heavily loaded clusters to clusters with still idle machines. The flocking design [146] is transparent to the users of the workstations, to the users of Condor and to Condor itself.

In the summer of 1994, a production flock of nine clusters in five different countries was set-up. The clusters are owned by the University of Wisconsin-Madison in the USA, the NIKHEF institute, the University of Amsterdam (UvA) and the Delft University of Technology all three in The Netherlands, CERN in Switzerland, JINR in Dubna, near Moscow and the Warsaw University in Poland. Our code served as a test case for this “World Flock” and at the same time we

site	# CPU-years	Architecture	#machines
UvA	10	sun4c	10
NIKHEF	20	sun4c, sun4m	25
Madison (CAE)	25	HP, sun4m, Intel	200
Madison (CS)	80	HP, sun4m, Intel	200
INFN-Bologna	10	DEC-alpha, HP, Intel	30
flock	6	sun4c, sun4m	250

Table 4.1: Estimated CPU-time delivered to this work through the Condor system.

profited from the computational resources that were disclosed to us worldwide.

In Table 4.1 we list statistics of our usage of the Condor system over the period 1993-1998.

Algorithmic (systolic) parallel SA

In our parallel SA, algorithmic parallelism is imposed on the generation of Markov chains. Markov chains of equal length but simulating different temperatures are generated concurrently in a systolic parallelisation scheme [147, 148]. The chains are generated in parallel and during their generation information is transferred from a given chain to its successor.

Each Markov chain is divided into a number of sub-chains equal to the number of available processors. The generation of chain $k+1$ is started as soon as the first sub-chain of chain k is completed. Equilibrium is not yet established by then. Quasi-equilibrium of the system is preserved by adopting intermediate results of previous Markov chains.

Let P be the number of processors, L the length of the Markov chains, $S = L/P$ the length of the sub-chains, T_k the temperature in the k^{th} Markov chain (denoted as M_k), and $X_{k,m,i}$ the i -th configuration of $M_{k,m}$, the m^{th} sub-chain of chain M_k . Fig. 4.1 shows the organisation of the systolic annealing algorithm for a three processor implementation. The first configuration of a new sub-chain in chain number k , $X_{k,m,1}$, $m > 1$, is either the last configuration of the previous sub-chain, $X_{k,m-1,S}$, or the final configuration of the last generated sub-chain of M_{k-1} , $X_{k-1,m,S}$. The choice is based on the Boltzmann distributions for two successive anneal temperatures:

$$X_{k,m,1} = \begin{cases} X_{k,m-1,S} & \text{if } \exp(E(X_{k,m-1,S})/T_k) \geq \exp(E(X_{k-1,m,S})/T_{k-1}) \\ X_{k-1,m,S} & \text{otherwise} \end{cases}$$

In the method above a new chain is started with the last configuration of the first subchain of the previous chain, $X_{k-1,1,S}$. The overlap enhancement developed by Kim and Kim [138] allows a new chain to start with either $X_{k-1,1,S}$ or the final




Figure 4.1: Schematic representation of a three processor implementation of the systolic algorithm. The arrows with large heads are the overlap enhancement.

configuration of chain $k - P$, $X_{k-P,P,S}$. This choice is also based on Eq. 4.1. The enhancement is explicit in Fig. 4.1 as the arrows with large heads between $M_{1,3}$ and $M_{4,1}$ and between $M_{2,3}$ and the right side of the figure. This method allows a more uniform communication pattern and therefore better systolic implementation. It also allows for a new cooling schedule where the standard deviation of the energy during a Markov chain is used to determine the size of the cooling step. The systolic implementation promises a large speedup over the whole anneal temperature range.

In our implementation we experiment with different cooling schedules: the anneal schedule based on a cool-rate discussed in Sect. 3.3, as well as an adaptive schedule. The latter is used in Sect. 4.3.4 in a study on the scaling properties of systolic SA. The adaptive schedule uses the following relation to calculate the temperature of the next Markov chain:

$$T_{k+1} = T_k \left(1 + \frac{\ln(1 + \delta) T_k}{3\sigma(M_{k-P})} \right)^{-1}, \quad (4.1)$$

where $\sigma(M_{k-P})$ is the standard deviation in the energy of chain $k - P$, which is mapped on the same processor as the calculation of chain k . For the first P chains $\sigma(M_{k-P})$ is not known and a constant is used instead. The parameter δ controls the difference between the probability distributions of two successive chains. Due to the overlap enhancement the information of a full Markov chain is used to determine the new temperature. The data provided by $\sigma(M_{k-P})$ is outdated (the more processors used the more outdated the information is). The effects of this is the following: since $\sigma(M_{k-P})$ is larger than $\sigma(M_{k-1})$ the temperature decrement is small for large P . This allows for a better exploration of the state space than for the fixed cooling schedule.

For the fixed cooling schedule, if the number of processors is chosen too large, the sub chains become too short to maintain quasi-equilibrium. For increasing

number of processors the new chains are started with worse and worse starting configurations and a deterioration of the quality is expected. At the same time more Markov chains need to be generated before a stable minimum is reached. For the adaptive cooling schedule effectively a slower cooling is used for larger numbers of processors. This slower cooling counterbalances the negative effects on the quality of the results for shorter sub-chains. But, more chains have to be generated to reach a stable minimum in energy.

The systolic algorithm has a simple communication pattern that can be efficiently implemented on a ring topology. The communication overhead is small since each processor contains a complete independent database for the optimisation problem. To interchange information a processor only has to send and receive information at the end of each sub-chain. This computational scheme can be implemented efficiently on a distributed memory Multiple Instruction stream Multiple Data stream (MIMD) architecture. In Sect. 4.3.2 the time complexity analysis of the systolic SA method is given.

Data parallel Metropolis

Although there are several algorithmic parallelisation methods for the Metropolis method [139, 140, 141] we have chosen a data parallel method [148]. Data parallelism is applied to the calculation of the relative importance within one Metropolis move. The calculation of the importance of the current and the proposed configuration involves calculation of the energy. In our case this is the most time consuming part of a Metropolis move. We use particles with two-body interactions and therefore the total energy is a sum over the interaction energy for all pairs of particles. These two-particle interactions are independent for each pair. The calculation of these interactions can be distributed over the processors by data decomposition (farming) such that each processor can operate independently. Afterward the partial results are combined. This is called farming.

In our implementation a master processor generates the Markov chain. At each Metropolis move the energy calculations are distributed over a processor farm while the master itself also takes part in the calculations.

The number of processors in the farm depends on the calculation and communication time required. The calculation time needed is dependent on the number of particles. The communication time is dependent on the number of farm processors and on the way these are connected. For reasons of efficient communication the farm processors are connected to the master as a binary tree. The number of farm processors is chosen such that the communication/calculation time ratio is optimal for the range of particles of interest. Section 4.3.3 contains the time complexity analysis for the data parallel Metropolis implementation.

Hybrid parallel SA

The parallel simulation methods discussed above, systolic SA and data parallel Metropolis, are chosen such that combinations of these parallelisation methods is possible [148, 99]. This results in a novel hybrid parallel implementation where several Markov chains are generated in parallel and importance calculation for Metropolis moves are distributed over a farm of processors.

Since the processors for the systolic implementation all generate independently a Markov chain, a farm is attached to each of these systolic processors (see Fig. 4.2). The total number of processors used in the simulation is P times the number of farm processors per systolic processor. In Sect. 4.3.4 the number of processors in the systolic ring is varied while the number of farm processors per systolic processor is kept constant. Both parallelisation strategies are completely




Figure 4.2: Diagram of hybrid implementation. Systolic processors (shaded circles) connected in a ring perform the annealing process, the tree processors (open circles) together with the systolic processors perform energy function calculations.

independent. Therefore both contribute fully to the overall speedup.

For the studies on the parallel and hybrid implementations we use a Parsytec GC with 512 T805 transputers with distributed memory. The large number of processors allow for the simultaneous implementation of both parallelisation schemes and scaling studies. We use the native communication primitives, PARIX, to send and receive information between processors. In Sect. 4.3.4 the time complexity analysis of the hybrid SA method is discussed. Although the T805 transputers are not fast by nowadays standard, the time complexity analysis gives us the possibility of extrapolating to modern and even future parallel machines.

Vectorised SA

We use a vector supercomputer to investigate how well the SA code can be vectorised and how well this compares to the above discussed hybrid parallel implementation.

The target machine in this case, a CRAY Y-MP 4/464, has four vector processors. However, since we aim at a comparison of the vector implementation versus the parallel implementation we have not exploited the parallelisation provided by the four processors.

The source code for the vector version is the same as for the sequential version discussed above. Apart from some minor changes in the code to improve vectorised processing, the vectorisation is carried out automatically by the highly optimised compiler.

Vector registers are used to hold the data before processing in a functional unit. These registers are also used to “pipeline” the vector operations (chaining of functional units). This means that the result of an arithmetic action is not written to memory but stored in one of the vector registers, in order to be used by a different arithmetic unit working in parallel with the first.

The CRAY Y-MP has eight scalar registers and eight vector registers of length 64. The number of available vector registers sets a limit to the amount of chaining of pipelines that can be obtained.

The theoretical maximum performance is reached when the two functional units, dedicated to addition/subtraction calculations and multiplication calculations, are constantly active. In this case during a clock cycle the results of both an addition/subtraction and a multiplication can be completed. The clock cycle of the CRAY Y-MP is 6 ns, which corresponds to a peak-rate of 166.7 MFlop/s, so the theoretical peak performance would be 333.3 MFlop/s.

The parts of the algorithm best suited for vector processing are the calculations of the energy of one particle and the calculation of the energy of the total system. The part in which the energy of one particle is calculated is the most time consuming part of the sequential algorithm. Vector processing of these parts promises a significant speedup. The time complexity analysis for the vectorised implementation is given in Sect. 4.3.5.

4.3 Time complexities

The time complexity formulas derived in this section (see also [99, 102]) give a trend indication for the execution time of the application. The trends are studied as a function of N , the number of particles, and P , the number of processors. The actual performance reached for a distributed memory machine is dependent on the type of nodes (e.g. the instruction issue-rate) and the bandwidth of the communication channels between the nodes. The performance of a program on

a vector computer depends on many elements of the hardware design and how well these are used by the compiler [132]. Among the hardware elements are: the size and number of vector registers (where the data vectors are stored for processing), number of concurrent paths to memory, instruction issue-rate and number of duplicate functional units (add/subtract and multiplication units.)

In this section we present and compare the time complexities for the sequential, hybrid and vectorised implementations.

4.3.1 Sequential SA

The core of the simulation program consists of Metropolis moves. These moves are repeated over and over to create a Markov chain. They consist of a displacement of a particle (a perturbation), a calculation of the energy difference (ΔE_p) and an update if the move is accepted. The radius is only perturbed in intervals. The energy difference ΔE_r that has to be calculated for each radius move, once per aN particle moves, where $a \in \mathbb{N}$ indicates how often the particles are moved between one radius move. In the following we multiply the time for a radius perturbation with $\alpha=(aN)^{-1}$ to account for this.

The time needed to displace a particle, t_m , is independent of N . The calculations of ΔE_p and ΔE_r , however, are dependent on N . The time needed to sequentially calculate ΔE_p is denoted as $t_{e,p}^{\text{seq}}(N)$, the time needed to perturb the radius and calculate ΔE_r as $t_{e,r}^{\text{seq}}(N)$.

For two body interactions, $t_{e,p}^{\text{seq}}(N)$ is linearly dependent on N . However, if the LJ potential is used, an interaction list can be implemented (see Sect. 3.5) to economise the calculation. For sake of generality, in the time complexity analysis discussed, we do not use such a list. The time $t_{e,r}^{\text{seq}}(N)$ includes the calculation of the total energy in the system since all distances are influenced by a change in radius. This is a $\mathcal{O}(N^2)$ calculation. If the factor α is chosen as $1/N$ ($m=1$), the radius is updated once every N moves (which gives reasonable results, data not shown). Then the contribution to the time complexity of a radius perturbation is also $\mathcal{O}(N)$. In that case the time complexity for one step in the Markov chain, t_1^{seq} , can be written as

$$t_1^{\text{seq}} = t_m + t_{e,p}^{\text{seq}}(N) + \alpha t_{e,r}^{\text{seq}}(N) = c_1 + c_2 N, \quad (4.2)$$

where c_1 and c_2 depend on the processor speed.

To obtain the time complexity of the complete annealing algorithm, the time complexity for one Metropolis move has to be multiplied by the length L of the Markov chains and the number of chains M needed for reaching a stable minimum:

$$t^{\text{seq}} = t_1^{\text{seq}} L M. \quad (4.3)$$

If for the stop criterion the deviation in the energies of the last configuration of the last 10 chains is used (see Sect. 3.5), M has an almost constant value as a function of N . This can be understood as follows. Using the normalised energy (energy per particle), the deviation in the energy as a function of temperature is independent of N . Since we keep the initial temperature T_i fixed (see 3.6) and the process is stopped at effectively the same temperature T_f , M is approximately constant.

The constants appearing in the time complexity formula (Eq. 4.2 with Eq. 4.3) are determined by taking two N -values and measuring their execution times. Then—as a check—for other N values the execution time is measured and compared to the execution time given by the time complexity formula. The trend in the execution time as a function of N is found correct (data not shown).

4.3.2 Systolic SA

In systolic SA, P chains are generated concurrently for a systolic ring of P processors. This means that in the time that in the sequential version one chain is generated, P chains are generated in the systolic version. But now communication time is associated with the implementation.

The number of chains generated in systolic SA is effectively somewhat larger than in the sequential implementation since the algorithm has to start up the processors in the ring one by one (see Fig. 4.1) and finishes likewise. This means that the total number of chains is equal to the number of chains in the sequential version plus the number of processors in the ring minus one.

Thus the time complexity of the systolic version is given by

$$t^{\text{sys}} = (t_1^{\text{seq}} \frac{L}{P} + t_c^{\text{sys}})(M + P - 1), \quad (4.4)$$

where t_c^{sys} is the communication time for transfer of information between sub-chains.

If a parallel computer system or a cluster of workstations in a reasonably fast network is used, the communication time is small compared with the calculation time of the Markov chain for long sub-chain lengths $S = L/P$. The information for the next processor in the ring is sent using asynchronous communication. However, the communication time in Eq. 4.4 is much longer than estimated from the length of the message and the start-up time [149]. Actually it is the time between finishing one sub-chain and starting the next one. We have measured for $N = 550$ that

$$t_{c,\text{calc}}^{\text{sys}} = 12 * 10^3 \mu s \quad (4.5)$$

$$t_{c,\text{timed}}^{\text{sys}} = 75 * 10^3 \mu s, \quad (4.6)$$

where $t_{c,\text{calc}}^{\text{sys}}$ is the time calculated from the setup time and the length of the message (see [149]) and $t_{c,\text{timed}}^{\text{sys}}$ is determined in an experiment. The difference is

due to the fact that processors have to wait for their predecessor to finish. Some processors have to wait a long time before they receive their information. The calculation of the Markov chain involves a stochastic process, therefore not all processors finish their work at the same time.

The complete time complexity of the systolic implementation including measured constants, is

$$t^{\text{sys}}(N) = ((1.4 * 10^3 + 45N) \frac{L}{P} + 75 * 10^3)(M + P - 1) \mu s, \quad (4.7)$$

To show the influence of the number of processors on the systolic SA algorithm we list the results of 10 experiments with $N = 50$ ($L=2500$) in Table 4.2. In this table we give the average number of chains, the average minimum energy found, the average execution time of the simulations and the time calculated from the time complexity. These simulations were performed with fixed radius, keeping the density equal to the optimum density for a flat surface.

P	M	average energy	execution time(s)	time complexity (s)
1	157 ± 1	-2.761 ± 0.002	$1.31 * 10^3 \pm 20$	-
4	161 ± 2	-2.754 ± 0.007	359 ± 4	390
8	175 ± 10	-2.74 ± 0.02	196 ± 4	220
16	230 ± 30	-2.70 ± 0.03	140 ± 20	160
32	360 ± 50	-2.68 ± 0.06	120 ± 15	141

Table 4.2: Results on the systolic SA algorithm for $N = 50$ (based on 5 experiments). One processor ($P = 1$) means that the sequential implementation is used.

The execution time compares reasonably well with the time estimated from the time complexity. Although there seems to be a systematic overestimation, the trend, however, is correct modelled.

The speedup (the time for sequential execution divided by the time for parallel execution) is larger than unity for all P . The efficiency (the speedup divided P) decreases from 0.91 ($P=4$) to 0.34 ($P=32$). More importantly, from the Table it is clear that for large P the sub-chains are too short to allow for quasi-equilibrium. This results in a larger number of chains in order to achieve a stable minimum. Therefore we have a dependency $M = M(P)$.

Also a deterioration in the quality is observed for increasing numbers of systolic processors. To keep the quality of the solution the same as in the sequential case we can put $L = L(N, P)$. If we use longer chain lengths for larger numbers of processors the sub-chains do not become too short and the obtained solution and the number of chains are approximately those in the sequential case.

If we determine what the chain length $L(N, P)$ should be to keep the quality of the solutions the same as in the sequential case we find, as an example, for

$N = 100$, the values listed in Table 4.3. The execution times for these new chain lengths are such that a major part of the speedup is lost. This behaviour is further discussed in Sect. 4.3.4 on the hybrid implementation where the effects of the fixed and adaptive cooling schedules are compared.

P	L/N
4	90
8	140
16	275
32	700

Table 4.3: Chain length per particle at some values for the number of systolic processors to keep the quality equal to the sequential case.

4.3.3 Data parallel Metropolis

In the Metropolis implementation we have chosen to use data decomposition by farming. The calculation of the energy of one particle and the calculation of the total energy, needed for a radius perturbation, are both performed in the farm.

For the data decomposition communication time is necessary to inform the farm processors when to calculate the energy for a given particle with proposed coordinates. Then all the processors (including the master) do their part of the calculation. When the calculations are finished the farm processors send their partial results upwards in the tree which again costs some communication time.

In the data parallel Metropolis implementation all the farm processors have the complete set of coordinates for all the particles. So the messages can be kept short. In the message to start the next calculation a flag signals if the previous move was accepted. In this case the coordinates of the particle from the previous move are updated in the farm using the already locally present data.

For the implementation we use a farm connected as a binary tree of depth two, i.e. seven processors including the master. This connection scheme results in a small communication overhead compared to a directly connected farm.

The time complexities for the calculations of ΔE_p and ΔE_r are

$$t_{e,p}^{\text{fun}} = c_1 \frac{N}{P_{\text{tree}}} + t_c^{\text{tree}} \quad (4.8)$$

$$t_{e,r}^{\text{fun}} = c_2(N) \frac{N}{P_{\text{tree}}} + t_c^{\text{tree}}, \quad (4.9)$$

where t_c^{tree} is the total communication time associated with the tree implementation, and where c_1 and $c_2(N)$ depend on the processor speed. They are different from the constants appearing in Eq. 4.2. The constant c_2 is of $\mathcal{O}(N)$, if we choose $\alpha = 1/N$ this dependency is removed in the total time complexity, $\tilde{c}_2 = c_2(N)/N$.

The time complexity for one step in the Markov chains is analogous to the sequential time complexity given in Eq. 4.2:

$$t_1^{\text{fun}} = t_m + (c_1 + \tilde{c}_2) \frac{N}{P_{\text{tree}}} + t_c^{\text{tree}} \left(1 + \frac{1}{N}\right). \quad (4.10)$$

The speedup gained by data decomposition, t_1/t_1^{fun} , can be close to P_{tree} for large N .

If we determine the constants c_1 and c_2 by time measurements for different N we find

$$t_{e,p}^{\text{fun}} = 4.4N + 960 \mu s \quad (4.11)$$

$$\frac{t_{e,r}^{\text{fun}}}{N} = 2.0N + \frac{960}{N} \mu s. \quad (4.12)$$

$$(4.13)$$

As comparison we measured $t_{e,p}^{\text{seq}}$ for the sequential case and found $31N \mu s$. At large N the speedup for $t_{e,p}^{\text{fun}}$ is close to P_{tree} , 7 in this case. Then the efficiency is near unity, i.e. the processors are fully utilised and contribute fully to the decrement of execution time. Our functional decomposition outperforms the sequential version if we use more than $N \approx 40$ particles.

The constant t_m is found to be $1.4 * 10^2 \mu s$

4.3.4 Hybrid SA

Since the algorithmic and functional parallelism introduced are independent, the time complexity for the hybrid implementation (systolic and tree decomposition together) is equal to that of the systolic version as given in Eq. 4.4 where t_1^{seq} is replaced by t_1^{fun} as given by Eq. 4.10. If N is large and P is small this implementation can give a speedup close to Pt_1/t_1^{fun} .

The complete time complexity of the hybrid implementation is

$$t^{\text{hyb}}(N) = \left((2.4 * 10^3 + \frac{9.6 * 10^2}{N} + 6.4N) \frac{L(N, P)}{P} + 75 * 10^3 \right) (M + P - 1) \mu s \quad (4.14)$$

In sequential annealing the chain length and number of chains are dependent on the anneal schedule. In Sect. 4.3.2 we used the cool-rate based anneal schedule. We also investigated the adaptive anneal schedule, based on information gathered in previous chains, and studied its influence on the behaviour of the systolic implementation. This adaptive rule is based on the principle that the stationary distribution of configurations of two successive chains should be close to each other. This strategy uses additional information from a previous chain to determine the temperature of the new chain. This additional information comes

from $\sigma(M_{k-P})$ which is the standard deviation in the energy values of chain $k-P$. The parameter δ controls how much the stationary distributions of the successive chains differ; it has a small value (here 0.1, see also [138]) which is kept constant for the present experiments.

It is observed for the adaptive anneal schedule that the quality is not decreasing with increasing number of systolic processors as is observed in Sect. 4.3.2 for the cool-rate based anneal schedule (cool-rate=0.95). This is a consequence of using $\sigma(M_{k-P})$ in the determination of the temperature T_{k+1} . When looking P chains back the deviation in the energy is large, which has as a consequence that the temperature is only slightly lowered for larger P . This results in an increasing number of Markov chains needed with increasing P .

For both cooling schedules the speed-up gained by parallelising the code is lost for larger numbers of systolic processors. In order to give a fair comparison between the two cooling strategies we look for the chain lengths that give the same quality of solutions. For the cool-rate based anneal schedule this means that we use a longer chain length for larger numbers of processors and for the dynamically adjusted temperature, where the cooling is effectively a function of P , we can in some cases use slightly smaller chain lengths with increasing number of processors. In Fig. 4.3 we give the execution times for both strategies for $N=20$. We observe that the execution times for both anneal schedules increase after a certain number of systolic processors. For the cool-rate based cooling the execution time increases because of the longer chain length that has to be used. In the dynamically adjusted temperature case the increase is due to the larger number of Markov chains to be generated before a stable configuration is found. For this particular case (Fig. 4.3) the dynamic adjustment has a minimum in the execution time at four systolic processors, while the temperature adjustment with the cool-rate gives a minimum at six processors. The execution time that can be achieved is lowest for the cool-rate temperature adjustment strategy.

For the studies presented above we use the Parsytec GCel machine with 512 processors. A much more up to date machine is the Parsytec CC40. Although it does not have a large number of processors and scaling studies can thus not be performed as well as on the 512 processor machine, we can adjust—based on machine parameters [150]—the constants appearing in the time complexities in Eq. 4.14. For the CC40 the calculation speed is increased by a factor ≈ 30 while communication speed is increased by a factor ≈ 2.5 . We then find

$$t^{\text{hyb}}(N) = \left((4.3 * 10^2 + \frac{3.8 * 10^2}{N} + 0.21N) \frac{L(N, P)}{P} + 6.9 * 10^3 \right) (M + P - 1) \mu s \quad (4.15)$$

Since the calculation time to communication time has changed, the tree decomposition, with the most frequent communications, we see that only for large N is the calculation time longer than the communication time. Thus for the new machine a tree decomposition over a smaller number of processors is advisable. Then,

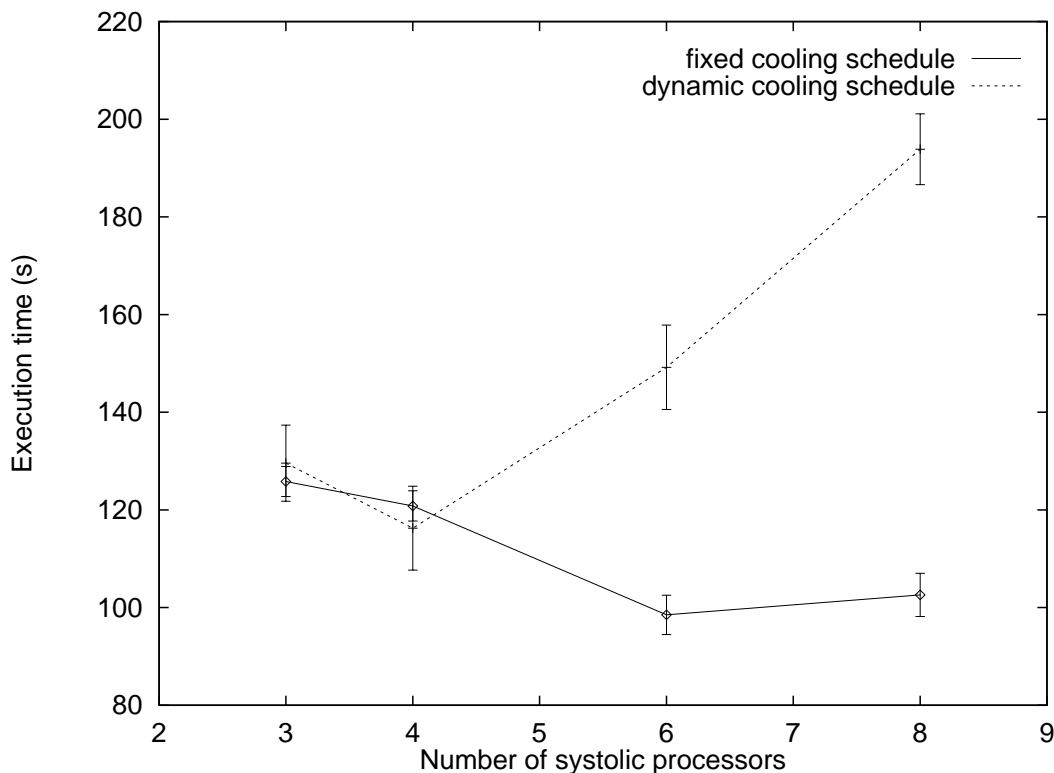


Figure 4.3: Execution times of the parallel implementation with different anneal schedules for $N=20$.

since only infrequent communication between systolic processors takes place, we can expect a speedup close to the factor for increased calculation speed.

4.3.5 Vectorised SA

The time complexity of the vector version consists of a part that is sequentially processed plus a part that is vector processed.

The execution time for any vector instruction consists of the set up time to route the operands to the functional units, and a processing time that is proportional to the vector length (number of elements of the vector).

The perturbation part of a Metropolis move has no dependencies on N and does not contain parts that can be vectorised. The calculation of ΔE_p takes by far the most time. First the distances are calculated and then for the LJ potential, before calculation of the interaction potential, the cut-off condition is applied. Then a vector mask operation is used to substantially reduce the vector length in the second part of the calculation: the calculation of the interaction potential. The full vector length is equal to $N - 1$. The vector of reduced length is equal to the number of particles that fall within the cut-off distance of the potential, typically 20.

The time complexity for the calculation of ΔE_r , necessary for a radius update, is given by

$$t_{e,r}^{\text{vec}} \approx \frac{1}{2} N t_{e,p}^{\text{vec}}(N). \quad (4.16)$$

We assume, for long vector lengths (large N), that we can again write

$$t_{e,p}^{\text{vec}} = c_1 + c_2 N. \quad (4.17)$$

By measuring execution times for N ranging up to 4000 we measured for large N (i.e. removing a small constant overhead)

$$t_{e,p}^{\text{vec}}(N) = (7.30 \pm 1.30) + (0.148 \pm 3.4 * 10^{-4}) N \mu s. \quad (4.18)$$

And thus, using Eq. 4.16, we find

$$t_{e,r}^{\text{vec}} = N(3.65 + 0.074N) \mu s. \quad (4.19)$$

By measuring the execution time for the sequentially processed perturbation step we find $t_m^{\text{seq}} = 125.0 \mu s$.

The time complexity for one Metropolis move of the vector implementation, analogous to the sequential time complexity Eq. 4.2, is given by

$$t_1^{\text{vec}}(N) = t_m^{\text{seq}} + t_{e,p}^{\text{vec}}(N) + \alpha t_{e,r}^{\text{vec}}(N). \quad (4.20)$$

If we choose $\alpha = \frac{1}{N}$ and fill in the values and expressions for t_m , $t_{e,p}^{\text{vec}}(N)$ and $t_{e,r}^{\text{vec}}(N)$ in Eq. 4.20, we obtain for the time complexity of one step in the Markov chain for the vector version for large number of particles:

$$t_1^{\text{vec}}(N) = 136 + 0.222N \mu s. \quad (4.21)$$

The time complexity for the complete vector version with cut-off restriction is now given by

$$t^{\text{vec}}(N) = t_1^{\text{vec}}(N) LM = (136 + 0.222N) LM \mu s. \quad (4.22)$$

We now compare the execution times measured on the CRAY Y-MP with those calculated from the time complexity. The results are given in Table 4.4. The data indicates that the predicted results are in reasonable correspondence with the timed values. The trend in N is modelled correctly.

4.4 Comparison of the sequential, hybrid and vector implementations

In this section we compare the three implementations: sequential, vector and hybrid [98].

N	L	M	$t_{\text{timed}}^{\text{vec}}(N)$ (sec)	$t_{\text{calc}}^{\text{vec}}(N)$ (sec)
50	3000	393	156	173
100	11000	407	652	708

Table 4.4: A comparison of the timed and calculated (from the time complexity) execution times on the CRAY Y-MP of the vectorised version of the SA program with applied cut-off condition.

The algorithms for the sequential and vector implementations are effectively the same. The much higher Flop/s rate of the CRAY and the vector capability makes its CPU's much faster than any common off the shelf CPU. If, for example for a $N=50$ system, the two implementations, run on the CRAY and the Parsytec machines, are compared we find 55 seconds on the CRAY and $1.3 * 10^2$ seconds on a T805. Thus a speedup of 24 is found, which is less than expected from the Flop/s rate (up to 333 Mflop/s for the CRAY and 1 Mflop/s for the T805). The vector capability of the CRAY is not fully exploited because our implementation does not solely exist of vector instructions that can be optimally chained together.

Vector computing does not have the ability to be adapted to a specific problem. Its power lies in the fact that you can easily adapt the sequential code. With a minimum of trouble high speedups can be achieved. In contrast, a massively parallel machine has a high degree of freedom in connectivity, the processor configuration can be adapted to the parallelism present in the algorithm. This is applied in the processor connections for the hybrid implementation.

The systolic part of the hybrid implementation is algorithmically different from the sequential code. For the two cooling rules which were used, we had to experimentally adapt chain lengths in order to ensure that the quality of the solutions is independent of the number of processors. For the fixed cool-rate cooling schedule this caused much larger chain lengths, for the adaptive schedule about the same length could be used for all P .

It turns out that the systolic part of the hybrid implementation in the present form does not scale properly with P . The reason for this is that not many configurations from a chain get accepted in the next chain at lower temperature. So the first configuration of a chain (which is the configuration of the first sub-chain of the previous Markov chain) is often the only configuration that is used for the calculation of the whole chain. The ratio of the number of accepted configurations of a previous chain divided by the total number of tried configurations from a previous chain is approximately constant and very low (≈ 0.04). Comparing the obtained results for the two anneal schedules, we find that the cool-rate based cooling rule has its minimum in the execution time at a larger number of processors than the dynamic temperature adjustment. Therefore the cool-rate step rule implementation gives shorter execution times.

We now compare the execution time of the vector implementation with that

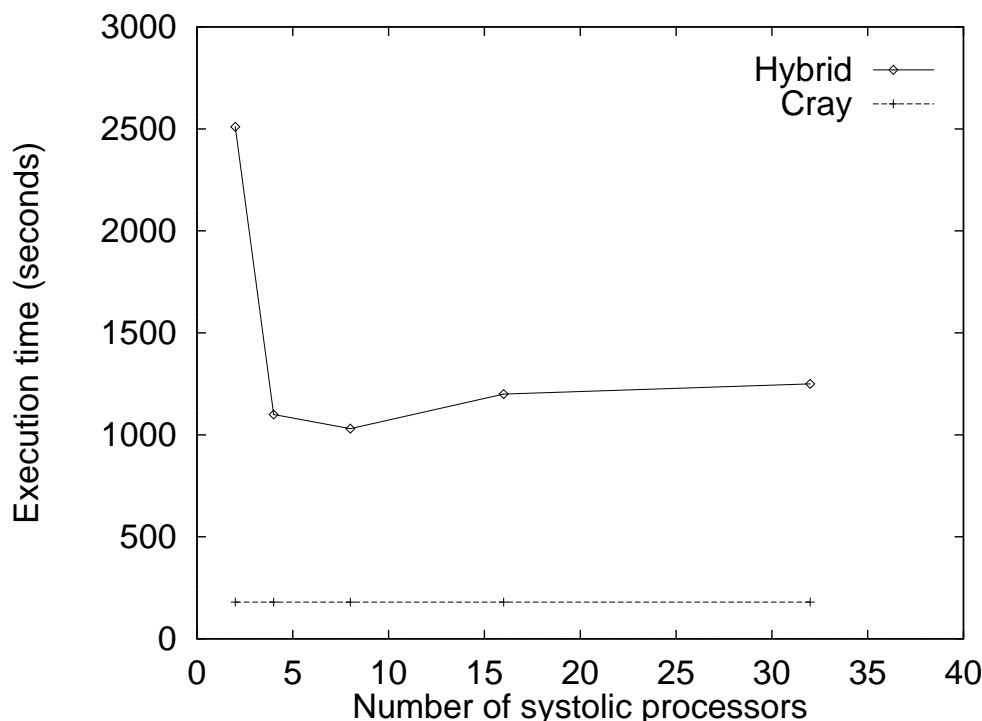


Figure 4.4: The execution times ($N = 100$) for the vector (CRAY) and the hybrid (PARSYTEC) implementation, for the same quality of results, as a function of the number of systolic processors P . The total number of processors is $7P$.

of the best performing parallel version (the hybrid implementation) for a system with $N = 100$ where the chain lengths are tuned to give the same quality for all data points. The results are given in Fig.5 where the data for the CRAY is shown as a horizontal line. These results show that the systolic part of the hybrid implementation for large P does not have good scaling properties and therefore it is not possible to outrun the CRAY computer.

The data in Fig.4.4 at 8 systolic processors has the lowest execution time for the hybrid implementation. The 8×7 processors are capable of processing 56 MFlop/s which is about 6 times slower than the capability of the CRAY. The execution time is about 5 times slower. This means that the parallel implementation has a high efficiency at that number of processors. For the 16 systolic processors case, capable of processing 112 MFlop/s which is one third of that of the CRAY, we see that the execution time is six times longer than for the CRAY. This means that a factor two is spent on communication and on the loss of precision in the iterative scheme of the parallel version.

R. Azencott [137] has shown that parallel SA implementations, where one Markov chain is generated in parallel by dividing the chain in sub-chains (simulating the same temperature) which periodically interact to determine one system state which is used for a further search, are not as efficient as independent

searches.

In the systolic part of our hybrid implementation the Markov chains interact periodically but—in contrast to the methods described by Azencott—these simulate slightly different temperatures and there is only a choice between two system states. Due to the dependence of the chain length on P to mask the consequences of the introduced algorithmic changes, we observe a reasonable efficiency for small P but inefficiencies for larger P . We conclude that independent searches are also more efficient for our implementation.

4.5 Future expectations

The processor and communication speed of coupled CPU's increase at an incredible rate. The processing power of the T805 transputers in the Parsytec GC 512 machine have—compared to processors in nowadays parallel machines—a relatively low processing power. In the time complexity analysis given in this chapter we can find the influence of machine parameters. These can then be altered to match a different machine. For example, the efficiency of the data parallel metropolis implementation strongly depends on the communication versus calculation time ratio. If a very fast communication network is available it becomes possible to use more processors in the farm. On the other hand if more powerful processors are available it is more efficient to use a smaller farm.

It has recently become possible to build machines that consists of a network with heterogeneous processors. One example, which is very well suited for our work, is the machine currently being built partly at the University of Amsterdam in cooperation with the research school ASCI (Advanced School for Computing and Imaging). This machine consists of 16 Pentium Pro 200 Mhz processors of which six are equipped with GRAPE4 [151] boards (48 processors each). These latter processors are special purpose, suited for fast evaluation of potential energy functions of the generalised Coulomb type [152].

The hybrid implementation, for the shown examples, has a minimum in execution time around six systolic processors. Each of the systolic processors has a processor farm attached for the energy calculations. The general purpose processors in the farm could be replaced by a single GRAPE processor. Then we would have optimally suited processor architectures for each of the two different parallelisation strategies.

If we change the timed constants in the time complexities to match the described heterogeneous network (compared to Parsytec GC: an average workstation is ≈ 100 times faster, the GRAPE boards are $\approx 3 \cdot 10^4$ times faster). We assume that the communication speed is 10 times faster. Then we find that the processors farm can easily be replaced by one GRAPE board, even if we use large N . The sequential and systolic implementation both running on workstations, do not change in relative speedup. The time complexity for the energy calcu-

lations in the hybrid implementation is greatly reduced by the use of GRAPE processors. For large N this can give a large speedup over the sequential version. For example, from the time complexities, we find for $N=1000$ that the systolic version gives a speedup of ≈ 5 over the sequential version while the hybrid version gives a speedup of ≈ 20 .

Not only from hardware but also from software developments a large increase of processing power can be expected. The hierarchical methods [153] are for this work an interesting approach. There remote particles are grouped together and their total interaction is calculated at once via multi-pole expansions of the fields. With this method the complexity of the energy calculations is reduced. For example for the long range Coulomb potential the direct energy calculation is $\mathcal{O}(N^2)$, with the use of hierarchical methods this can be reduced to $\mathcal{O}(N \log N)$ or even $\mathcal{O}(N)$.

4.6 Conclusions

We find that the parallel High Performance Computing (HPC) (vector and parallel) implementations give a speedup over the sequential implementation. Our simulation code, however, does not exploit the vector processor fully. The data parallel implementation is capable of achieving a high efficiency. The systolic implementation, is not scalable to large P due to the introduced functional differences. The hybrid implementation is an efficient parallel simulation tool for small numbers of systolic processors. As shown in our examples, 56 processors can be exploited at reasonable efficiency.

The vector implementation is easily obtained from the sequential code but vector supercomputers are scarce and CPU time costly. The parallel implementation has a relative long development time and the usage of the for us available parallel machines is limited by many other users. We use the HPC implementations for tuning of parameters and to find solutions for relatively large optimisation problems where only a few solutions suffice but which have to be obtained in a short time span. It is, however, not practical for the bulk of the simulations we have to perform.

The time complexities derived in this chapter can be used to give an approximate idea how our code would run on other computer architectures. Changing calculation to communication speed has mainly influence on the functional decomposition of the energy calculations. If calculation speed increases more than communication speed, a smaller number of tree processors is more efficient. For the systolic implementation mainly the calculation speed is of importance.

In our case hundreds of thousands of independent jobs were processed in a time span of a few years. Concerning the availability of computers, the High Throughput Computing (HTC) paradigm is much more practical than the HPC paradigm. The Condor distributed batch system, opening the possibility for HTC

access allows to harvest otherwise wasted computing cycles, turning them into productive work without any negative impact on interactive users.

Chapter 5

Studies of crystallisation on a spherical surface

In this chapter we discuss several aspects of the behaviour of a N -particle Lennard-Jones (LJ) system on a spherical surface during crystallisation. First we look for signatures of the phase transition and the temperature T_c at which it takes place. Around that temperature we study the behaviour of positional and bond orientational order and the behaviour of defects, in particular disclinations (d-charges).

5.1 Introduction

With the Metropolis method described in Sect.3.4, we perform simulations in the NVT ensemble using a Monte Carlo (MC) Simulated Annealing (SA) method. The chosen system size, $N=100$, is typical for the results in Chapt.6 on crystalline arrangements. This number of particles, however, is too small for studies concerning coexistence and testing of the KTHNY theory (discussed in Chapt. 2).

The simulations are started at $T_i=64$, where the system is far above the freezing transition. After the system is brought to thermal equilibrium (TE) it is used for the “measurement” of TE properties. While the simulation continues to provide TE data, a configuration of the equilibrated system is used to start a simulation at a slightly lower temperature. After equilibration, simulation data are extracted and a new simulation, again at a slightly lower temperature, is started. We continue this process until $T = 10^{-4}$. This scheme is optimally suited for job parallelism (see Sect.4.2). We used the Condor distributed batch system (see Sect.4.2) to obtain the necessary computer power. Additional to the $N=100$ system we performed simulations on systems with $N=32$, 33 and 48 .

The TE data are gathered as follows. At given T we perform S_e Metropolis equilibration steps (per particle) before starting data extraction. The system is then monitored for S_m steps where data are collected at each step. Averages over these data are stored; for any property being measured, at the given T ,

they provide one data point. Then again an equilibration cycle to separate the measurement cycles is made. Of all the collected averages the first S_c data points are only used to check if the system has indeed reached TE sufficiently. The next $10S_c$ data points are used in the results of this chapter.

Additional to the above described simulations we perform a small number of simulations on larger systems in a NVE ensemble with the Molecular Dynamics (MD) method described in Sect. 3.5. For example, in a system of $N=4000$ the potential energies are calculated as a function of T . The system is started with a random configuration and then allowed to equilibrate over S_e time steps. At TE the average kinetic energy in the system is equal to the desired temperature. We also show a single result for a $N=10242$ system at $T=0.7$. The start configuration for this simulation is an icosahedral distribution.

For the MC calculations we set $S_e=100$, $S_m=50$, $S_c=100$ in the T range 64 to 10^{-4} . For the MD simulations we use $S_e=5000$, $S_m=21000$ ($\delta t = 0.0005$) in the T range from 0.5 to 1.2

5.2 Identification of the phase transition

In studying the crystallisation process we explore various thermodynamics properties over the temperature range in which the transition from liquid to solid takes place. In particular we try to find the freezing transition for the spherical LJ system by studying the energy and pressure as a function of T .

A rough way to locate the freezing transition is to plot the potential energy versus temperature $E(T)$ and look for a change of slope [76]. In this chapter the energy E is taken as the average per particle. In the plot of pressure versus temperature $P(T)$ the transition can also be visible [49]. The pressure is calculated as discussed in Sect. 2.6.2. The T range in which the transition temperature T_c is expected should be comparable with a similar system on a flat surface, at least for sufficiently large system size N .

In the chapter on zero temperature arrangement the energy as a function of N shows fluctuations. In this chapter we study in which T -range the difference in energy remains present. For these studies we calculate $E(T)$ for $N=31, \dots, 37$.

Results

In Fig. 5.1 we plot $E(T)$ for $N=100$ and 4000. The $P(T)$ curve is shown in Fig. 5.2 for $N=100$. In the figures the data from Udink for a flat 2D LJ system for $N=12480$ are also plotted (Fig 1. on page 60 in Ref. [49]).

At low temperatures $E(T)$ is an approximately straight line. No steep changes in slope are visible for the $N=100$ and 4000 systems, while for Udink we do see such changes below $T \approx 1$. The behaviour of $E(T)$ for $N=4000$, however, is more like that of Udink than the $N=100$ system.

The $P(T)$ curve crosses zero at $T \approx 0.10$. Below that temperature the density would increase if the radius is a free parameter (see Sect. 6.2 and 6.3.2). For higher temperatures the system favours a larger sphere. Also here only a gradual change in slope is apparent in the $N=100$ data, while Udink's data clearly show a transition effect.

If we compare the results for the spherical system and the results for Udink's large flat system, we find that both the potential energy and the pressure are higher for the spherical system.




Figure 5.1: The average potential energy at fixed density as a function of temperature for $N=100$ and 4000 , compared with data from Udink [49] for a flat 2D LJ system with $N=12480$.

In Fig. 5.3 we plot $E(T)$ obtained for $N = 31$ to 37 . The deviation in the data is of the same size as the markers. The ordering of the energies at $T=0$ (see Chapt. 6) is reflected in the ordering at finite T . The differences in energy remain significant at least up to $T \approx 0.3$. For $N=32$ $E(T)$ is significantly lower than for the other N -values up to $T \approx 1$, i.e. over the whole crystallisation range.

Analysis and discussion

That in Fig. 5.1 no abrupt transition is visible may be due to the representation and the noise in the data. If a sharp phase transition is present, it should be found in the constant volume specific heat C_v . The obtained $C_v(T)$ is shown in Fig. 5.4. Noise in the $E(T)$ data is here suppressed by a $\kappa=9$ (see Sect. 2.6.1) points derivative with intrinsic smoothing. At low T , as expected $C_v \rightarrow 1$.




Figure 5.2: The same as Fig. 5.1 for the average pressure for $N=100$ and $N=12480$.




Figure 5.3: The average potential energy as a function of temperature for $N=31$ to 37.

Around $T=0.35$ a clear change in slope is found. A subtle but remarkable change is visible around $T=0.8$. Indicating that at both temperatures a change in system behaviour takes place.




Figure 5.4: The specific heat C_v as a function of temperature for $N=100$.

The energy, like the pressure, is expected—as observed in Fig. 5.1 and 5.3—to depend on N . Two effects play a role. Firstly, at zero temperature, the packing efficiency is a function of N (see Sect. 6.3.2). The general trend in the packing efficiency indicates that systems with small N need on average for each particle more space due to the relatively higher number of d-charges. This may also play a role at $T > 0$. Secondly, we evaluate the LJ interaction using the interparticle distance through the sphere. Then, compared with a flat surface, distances to far particles are shorter, and—because of the LJ attractive tail—the energy contribution becomes lower. For very large N both effects gradually disappear. In Sect. 5.3 we show that the effects of evaluating distances *through* compared with *over* the sphere are rather small. The packing efficiency is the main effect concerning the difference between a spherical and a flat LJ system. The $E(T)$ curves as a function of N (for $N=31, \dots, 37$) indicate that the energy differences due to packing constraints remain clearly visible for $T < 0.3$. For the especially optimal zero-temperature arrangement, here $N=32$ (see Chapt. 6), the difference in energy compared to the other N values shows up already at $T \approx 1$. This indicates that packing efficiencies start playing an important role in initial freezing.

From the data given above we infer that at sufficiently large N , most likely there is a rather sharp transition at a temperature T_c which is near that for a

flat surface. For $N=4000$ a change in slope, giving a peak in C_v below $T \approx 1$ is observed in the same T -range as for Udink's flat system; however, not as pronounced because of the smaller system size (data not shown). Specifically both by Udink [49] and Bakker *et al.* [76] where the change in slope is more pronounced, much larger systems were used (12480 and 10864 respectively). But also on general grounds, the width of a phase transition is known to be broader at small system size [154, 155]. However, in $C_v(T)$ two significant effects occur, one around $T = 0.8$, the other around $T = 0.35$. It looks as if part of the $N=100$ system freezes at the same T_c as the large N systems (the small peak around $T=0.8$) and the rest of the system freezing gradually between $T=0.8$ and $T=0.35$.

5.3 Positional order

In this section we study the positional order as a function of T . Sensitive to the positional order are the radial distribution function $g(r)$, and the related structure factor $S(k)$. The definition and methods of calculation on a spherical surface of these functions is discussed in Sect. 2.6.5. Especially the positions and heights of peaks in $g(r)$ and $S(k)$ are used in evaluating the positional order.

We show results of the MC simulations with $N=100$ and one MD result with $N=10242$. The simulation set-ups are discussed in the introduction of this chapter.

For the calculation of $g(r)$ and $S(k)$ we use the distance *over* the sphere. Then the frequency of peaks does not increase for larger distances due to the spherical geometry. The value of r ranges from 0 to πR , which is the largest distance possible over a sphere. For $N=100$ $R = 2.689668$ and for $N=10242$ $R=29.94221$. In the figures we plot the data versus the distance over the sphere expressed as arc-angle or simply angle (R being known and fixed). The value of k in the structure factor ranges from 5 to 25. Within this range the peaks of the first few reciprocal lattice vectors of a flat triangular lattice are located.

On a flat surface, the area of a ring of width dr and radius r increases linearly with r . On the surface of a sphere this area changes as a sine, with its maximum at the largest circle, an "equator". Interparticle distances to an equator (arc-angle π) have a large probability. In evaluating $g(r)$ on a sphere this can be simply taken into account. In the evaluation of $S(k)$, however, this is not taken into account. Therefore comparison with a flat surface is not straightforward. Near the peak corresponding to the first reciprocal lattice vector, however, the distances are not yet severely distorted by the curved surface and only minor differences with a flat surface are there expected.

The number, and the parameters (height, width and position) of peaks in $g(r)$ and $S(k)$ are of importance in studying crystallisation. We obtain the peaks in the following way. First the functions are smoothed with a Savitzky-Golay

filter [131] where 7 points to the left and right with a smoothing polynomial of order 4 are used. Then the most prominent noise is removed. Peaks are detected by finding points that have on both sides two points with progressively lower values. Finally a quadratic least squares fit [131] is made on 5 points of the original (non-smoothed) data around the highest point of the peaks. The maximum of the fit and the position of this maximum are then used as the height and position of the peak.

Results

In Fig. 5.5 we plot $g(r)$ for $N=100$ at three temperatures: $T=4.0$, 0.70 and 0.10. At $T = 4.0$ the order in the system for larger distances becomes small and thus




Figure 5.5: The radial distribution function $g(r)$ at three temperatures for $N=100$, where the distance is expressed as arc-angle over the sphere in radians.

$g(r)$ approximates unity. The first peak is then also relatively low and broad compared to those at lower T . At $T = 0.70$ fluctuations around unity are visible up to the largest distance. At $T = 0.10$ the split second peak and the many small peaks indicate strong ordering over the entire system.

In Fig. 5.6 $S(k)$ is plotted for the same three temperatures as used for $g(r)$. As T decreases the peaks in the structure factor become more pronounced and they often split-up into several peaks, indicating the formation of neighbour rings with sharply peaked distances to a central particle.




Figure 5.6: The structure factor $S(k)$ at three temperatures for $N=100$.

Analysis and discussion

An estimate of the distance over which the system is ordered can be deduced from the height of the peaks in $g(r)$ as a function of distance. By fitting a decaying exponential or power function to the peaks, the parameter in the fit function controlling the decay, can be determined. This parameter can be studied as a function of T .

If we only take account for the peaks in $g(r)$ higher than the unit value, we find that at $T < 0.3$ a large variation in height is present. The results of fitting an exponential or power function are not clear in this region (data not shown). For $T > 0.3$ the peaks fall on a smooth curve. Then fitting an exponential function on the third and further peaks gives better results than fitting a power function. It is found that the ordering decreases for higher T . Below $T \approx 0.3$ the parameter as a function of T has a different slope (data not shown).

As discussed in Sect. 2.7.3 concerning crystallisation the first peak in the structure factor $S(k)$ plays an important role. In Fig. 5.7 the height of this peak is shown as a function of T . The distance corresponding to this peak is not different very much on a flat surface. We find that around $T=0.8$ a small broad peak is present in $S(k)$. Around $T=0.35$ a fluctuation is observed.

That we find a different behaviour for the peaks in $g(r)$ for $T > 0.3$ compared to $T < 0.3$ suggests that a drastic change in system behaviour take place around that temperature. Because of the presence of d-charges an exponential decrease in




Figure 5.7: The height of the first peak in $S(k)$ as a function of temperature ($N=100$).

the height of the peaks is expected [19]. Above $T = 0.3$ we can fit an exponential function through the peaks (data not shown).

The functions $g(r)$ and $S(k)$ are essentially related through a Fourier transform. Actually $S(k)$ describes the Fourier components of density fluctuations (see Sect. 2.6.5. The height of the first peak of $S(k)$ describes the strength of the fluctuation in $g(r)$ with length corresponding to the first reciprocal “lattice vector”. In our case it corresponds to the average nearest neighbour distance. Therefore the peaks at high temperature for a spherical surface are found at approximately the same k -values as for a flat surface.

Ramakrishnan and Yussouf [63][64] predict that for a flat 2D system (with hard core repulsion) the height of the first peak in $S(k)$ is 5.75 at the freezing transition. If we use that value for our data we obtain $T_c \approx 0.3$. According to the data of Udink, who finds a value somewhat below 5.2, we expect there $T_c \approx 0.5$.

The value of our $S(k)$ for $T = 0.70$ is different from that found by Udink for flat 2D at $T = 0.69$. Especially the second peak just starts to split in our case, while on a flat surface the second peak is then already clearly split. This can again be ascribed to the larger N . In the MD run with $N=10242$ at $T = 0.7$ we find a structure factor, shown in Fig. 5.8, much alike the flat situation. The positions of the peaks occur at the same k -values. Somewhat remarkably, the first peak is much higher (sharper) in our case.

It thus seems that T_c has shifted downward for our $N=100$ system. The second peak does not split as clearly as for large N . In the work of Udink the




Figure 5.8: The structure factor $S(k)$ for a spherical $N=10242$ system at $T=0.7$, calculated with MD. The first peak is located at $k=7.1$.

second peak already splits at $T = 1.01$ while our second peak for $N=100$ splits at $T = 0.43$. Little below that temperature the peak splits a second time.


As a check, we can study the influence of evaluating interparticle distances used in potential energy evaluations *over* instead of *through* (default) the sphere. The distances used for calculation of $g(r)$ are in both cases taken over the sphere. In Fig. 5.9a the function $g(r)$ is plotted for two simulations with $N=100$, $T=0.70$ and $\rho = 1.10$. In one run the particle interactions are evaluated through the sphere, in the other over the sphere. In Fig. 5.9b the difference of the two $g(r)$ curves is given. This difference increases with distance. Over the range where the LJ potential has an appreciable value, not much difference between the two $g(r)$ is found. Therefore the difference caused by the two metrics is not of importance. However, since distances to next nearest neighbours are larger when calculated over the sphere, a higher energy is found; $E = -2.389$ over the sphere and $E = -2.436$ through the sphere.

5.4 Bond orientational order

In this section we consider the bond orientational order. This ordering indicates how the relative orientations of particle pairs in the system become correlated.

The simulation set-up is described in the introduction of this chapter. We use

(a)



(b)

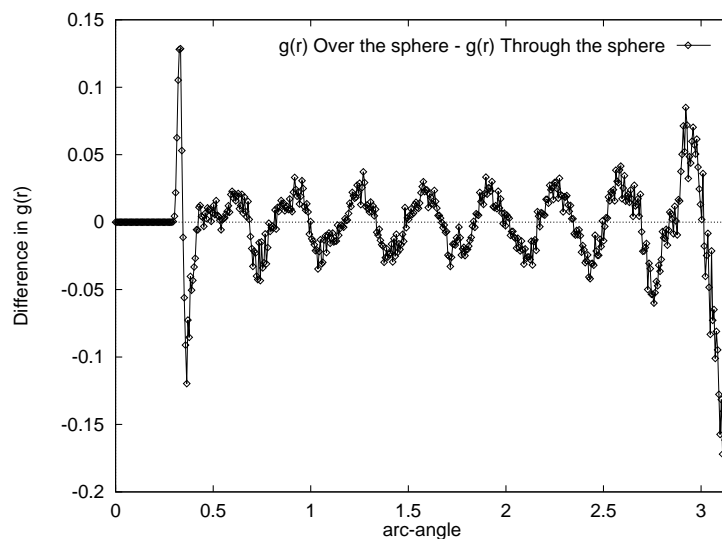


Figure 5.9: a) The function $g(r)$ for interactions evaluated *over* the sphere (line) and interactions evaluated *through* the sphere (diamonds); b) the difference between the two $g(r)$ distributions. The function $g(r)$ itself is in both cases evaluated with distances over the sphere.

the bond order $h_n(r)$ as defined by Giarritta *et al.* [45] (see Sect. 2.6.5). For n here the values 5 and 6 are considered only. The value of r ranges from 0 to πR . We show mainly results of simulations with $N=100$ and a few results for $N=32$, 33, 48, 4000 and 10242.

Results

In Fig. 5.10 the correlation functions $h_n(r)$ are shown for $n=5$ and 6 at $T=4.0$, 0.70 and 0.10, for $N=100$. In comparing the figures we notice that at longer distance, and at higher T the functions $h_5(r)$ approach a non-zero value. In fact, due to the comparatively small number of particles and the curved surface we find a strongly fluctuating behaviour for $h_n(r)$. Moreover, $h_5(r)$ does not approach zero at large distance.

For $N=4000$ and 10242, at $T \approx 0.7$ we deduce $h_6(r)$ as shown in Fig. 5.11. There the function goes to zero halfway the sphere and then rises again in magnitude. The difference in sign at large distance reflects a rotation in average bond angle of $\approx \pi/6$.

Analysis and discussion

It is somewhat remarkable that $h_5(r)$ does not go to zero at larger distances at high T while the calculations for the $h_n(r)$ are, apart from the value of n , the same. This would indicate that there is a correlation in bond order, but not in position. This is not completely understood.

Compared with $h_6(r)$ for $N=10242$ it is also remarkable that for $N=100$ and 4000 a change in sign takes place at large r . To understand this, we consider the symmetries favoured by the system, as revealed by the $T = 0$ arrangements. For the $N=10242$ system, where the start arrangement at low T is a generated frequency breakdown of an icosahedron, remnants of icosahedral symmetry are still present. Then, on two opposite sides of the sphere, the orientation of the bonds is the same. For the $N=100$ system tetrahedral symmetry is found at low T . With the latter system it is possible that the orientation of the bonds changes. The fact that $h_6(r)$ for $N=4000$ and 10242 is near zero halfway the sphere comes from the contribution of many particles near the equatorial region. The orientations of the bonds in this region change circumnavigating around the sphere and in total tend to cancel.

As shown in Sect. 6.3.4, the symmetry of the arrangements at zero temperature is strongly N dependent. If the change of sign found in $h_6(r)$ for $N=100$ is due to the symmetry, then for other N and thus for other symmetries, also differences in positions and sign of the peaks are expected to occur. To test this we performed some more simulations at $T = 0.1$ for $N=32$ with icosahedral symmetry, $N=33$ with only a mirror plane (C_s), and $N=48$ with octahedral symmetry. The functions $h_5(r)$ and $h_6(r)$ for these three cases are shown in Fig. 5.12.




Figure 5.10: The bond orientational order $h_5(r)$ and $h_6(r)$ for $N=100$ at different temperatures.




Figure 5.11: The function $h_6(r)$ for $N=10242$ and $N=4000$ at $T=0.7$. The curve for $N=10242$ has positive correlation at large distance while $N=4000$ has a negative correlation.

The data show that for C_s and octahedral symmetry an inversion of the correlation takes place. For the icosahedral symmetry the correlation is positive near the largest distance. The bond order is strongly dependent on the symmetry as shown by the differences for $N=32$ and 33 . For $N=33$ the peaks are less high and drop off faster than for $N=32$. Additionally $h_6(r)$ becomes negatively correlated for r larger than $\pi/2$. For $N=48$ at large distance also a negative correlation is observed. The peaks do not fall-off as fast as for $N=33$ but faster than for $N=32$. Conversely, from the positions of the peaks and the length over which appreciable correlation is observed in the bond angle correlation functions in principle one can deduce information on the symmetries of the averaged arrangement, even at finite temperature.

The measurements of $h_n(r)$ show that at low T , bond order is present throughout the system. For the icosahedral symmetry we find a positive correlation. For other cases, however, we find that in $h_6(r)$ also a negative correlation can occur at large distance. This means that there on average opposite hemispheres have an additional rotation of $\pi/6$.

Besides the change in sign of the correlation we find that for higher symmetry the correlation is stronger in magnitude. This is shown in Fig. 5.12 for $T = 0.10$, but the same conclusion can be drawn at $T = 0.5$ and 1.0 (data not shown). Thus T_c is presumably not a continuous function of N , but also a function of




Figure 5.12: The functions $h_5(r)$ and $h_6(r)$ at $T = 0.1$ for $N=32$ (top), 33 (middle) and 48 (bottom).

symmetry, shifting upward for increasing symmetry. There the ordering is set-in at a higher temperature as is also indicated in $E(T)$ as shown in Fig. 5.3. The structure of the energy landscape over the state space is probably simpler for systems with highly symmetric arrangements, allowing crystallisation at higher temperatures.

5.5 Disclination ordering

In this section we study the radial distribution functions $g_{\alpha,\beta}(r)$ for specific combinations of disclinations, d-charges, with α and β as the d-charge values. The definition of $g_{\alpha,\beta}(r)$ is given in Sect. 2.6.5.

Since d-charges are inherent to the spherical system we are particularly interested to study their behaviour as a function of T with $g_{\alpha,\beta}(r)$ for the appropriate α and β . At low T no particles with d-charge other than ± 1 are found.

In theories on the (indirect, lattice mediated) interaction between d-charges (see Chapt. 2), it is found that a repulsive interaction is present between like-sign d-charges and attractive interaction between unlike-sign d-charges. The interaction between different d-charges can be deduced by examining the favoured distances between them.

In this section all results are obtained for the $N=100$ system. The simulation set-up is described in the introduction of the chapter. For the results shown we use (combinations of) d-charges -2, -1, 0, 1, 2. In the figures the $g_{\alpha,\beta}(r)$ -plots are made again as a function of arc-distance, measured as arc-angle over the sphere.

Results

To study the influence of d-charges on the system properties we first calculate their densities, this indicates for which α and β it is interesting to evaluate $g_{\alpha,\beta}(r)$.

In Fig. 5.13 the relative density of d-charges as a function of T is plotted. It can be expected that at very high temperatures (much higher than shown in the figure) the concentrations of the various d-charges is approximately equal, although a net d-charge of -12 remains present. Lowering the temperature, the d-charges ± 2 are the first to disappear. At $T \approx 7$ the concentration becomes less than 1%, i.e. it becomes negligible. Below $T = 0.2$ no d-charges ± 2 are observed in our simulations. At $T \approx 0.03$ the concentration of d-charges +1 reaches the 1% level, and below $T = 0.006$ of this type no d-charge at all is found anymore. Then the density of d-charges -1 and 0 remain constant at a finite value. At zero temperature a relative density of 0.12, as expected from the Euler theorem for the spherical topology of the system remains present (see Chapt. 2). In Fig. 5.13 we observe discontinuities in slope. Most prominently at $T \approx 0.3$ but—less visible—also around $T=0.6$. The standard deviation on the density curves for d-charges -1, 0 and +1, is relatively large in this T -range indicating that in this range a




Figure 5.13: The concentration of d-charges as a function of temperature.

change in behaviour takes place.

We now study the behaviour of the defect radial distribution functions where we show results for appropriate d-charge values for the used temperatures.

In Fig. 5.14 we plot $g_{\alpha,\beta}(r)$ at $T = 4.0$, in the liquid phase far above T_c , for several combinations of α and β . Since the correlation length is short we restrict the arc-angle range to 2. The curves for $g_{-2,-2}(r)$ and $g_{+2,+2}(r)$ have been smoothed to remove some of the noise. A Savitzky-Golay filter [131] with 7 points to the right and left and a smoothing polynomial of order 4 is used. The peaks in $g_{0,0}(r)$ are associated to rings of neighbouring particles for the zero d-charge particles. The following arc-angles for the first two peaks are found: 0.340 and 0.668. These are used for comparison with the positions of the peaks in the $g_{\alpha,\beta}(r)$ for d-charges $\neq 0$. The function $g_{-1,-1}$ peaks at the nearest neighbour distance, $g_{-2,-2}$ has its first peak around the second neighbour rings. The d-charge pairs $-2,-1$ and $+1,+2$ differ on a sphere mainly because d-charges -1 are in excess. The corresponding functions $g(r)$ appear to be negatively correlated.

In Fig. 5.15 we plot $g_{\alpha,\beta}(r)$ at $T = 0.70$. The d-charges ± 2 have (almost) disappeared, therefore we only plot combinations of d-charges $-1,0$ and $+1$. If $g(r)$ is compared with $g_{0,0}(r)$ we find that the ordering of particles with d-charge zero is softened by the particles with d-charge different from zero.




Figure 5.14: Various $g_{\alpha,\beta}(r)$ functions at $T = 4.0$ (liquid).

Analysis and discussion

We find that the density of d-charges ± 2 disappear well above the freezing transition and thus do not influence crystallisation in a significant way. This does not hold for the d-charges $+1$. Below $T = 0.006$ we find that—as expected—only 12 d-charges -1 are present. Around $T = 0.3$ a change in slope of the densities for d-charges -1 , 0 and $+1$ is observed in Fig. 5.13. This indicates that some change in system behaviour takes place.

The position and height of the first peak in $g_{\alpha,\beta}$ gives information about how favourable the nearest neighbour distance is for d-charges α and β . This can be used to understand the clustering of d-charges at higher temperature and their distribution at low temperatures. We see that like-sign d-charges tend to be relatively far away from each other compared to $g_{0,0}$ (see for example $g_{+1,+2}$ and $g_{-2,-1}$). On the other hand unlike-sign d-charges tend to be found relatively close to each other compared to $g_{0,0}$ (see for example $g_{-2,+2}$ and $g_{-1,+1}$).

In Table 5.1 the position of the first main peak is given at three temperatures. For the d-charge combinations $(0,0)$, $(-1,0)$ and $(-1,1)$ the favoured distance becomes shorter at higher T . With more kinetic energy present at constant density the LJ core can be climbed higher, giving shorter average distances. The favoured distances for $(-1,1)$ are also somewhat smaller than for $(0,0)$. This indicates that




Figure 5.15: Various $g_{\alpha,\beta}(r)$ functions at $T = 0.70$ (beginning of crystallisation).

α,β	$T = 0.1$	$T = 0.7$	$T = 4.0$
0,0	0.37	0.36	0.34
-1,-1	0.53	0.55	0.32
-1,0	0.36	0.35	0.33
-1,1	0.36	0.35	0.34
1,1	1.3	0.50	0.49

Table 5.1: The arc-angle of the first main peak in $g_{\alpha,\beta}(r)$ for combinations of d-charges -1, 0 and +1. For all cases the deviation in the values is less than 0.01.

d-charges -1 and $+1$ are on average closer together than d-charges 0 . The same is true for the combination $(-1,0)$. The favoured distance is large for the combinations $(-1,-1)$ and $(1,1)$. At $T = 0.1$ the combination $(1,1)$ shows several high peaks. The first main peak has position 1.3 . We note that this is another peak than the ones at higher T .

An especially interesting case is $(-1,1)$. At both $T = 4.0$ and 0.70 the first peak is very high, while further peaks are either absent or small. This indicates that these d-charges are mainly found as neighbours, and that no significant ordering is present at larger distances. Locally excited d-charge pairs (thermally or by d-charge polarisation, see Sect. 2.3) are mainly—and at lower T exclusively—of type ± 1 .

Around $T = 0.3$ a change in the system behaviour seems to take place. There the d-charges ± 2 completely disappear in our simulations and the density of d-charges ± 1 shows an abrupt change in slope.

The d-charge pairs ± 1 remain present down to very low T where the entire system is in a solid state. This reflects that the system is in an excited state (with frozen-in thermally excited d-charge pairs), or that it is in the ground state with d-charge pairs induced by d-charges polarisation due to the inherent d-charges.

The qualitatively expected behaviour of the interaction between d-charges is in good correspondence with our results. Unlike-sign d-charges are found as close neighbours while like-sign d-charges at low T show up at relatively large distances.

5.6 Disclination clustering

In this section we examine the composition of the clusters formed by d-charges. As shown in the previous section, unlike-sign d-charges tend to be close to each other. Therefore we expect these d-charges to form clusters.

The results in this section are for the $N=100$ system as described in Sect. 5.2. The method to identify d-charge clusters is discussed in Sect. 2.6.3.

Results

In Fig. 5.16 we plot the number of clusters as a function of T . At high T only few clusters are found while many d-charges are present in the system (see Fig. 5.13). Thus the clusters are large. Between $T=0.8$ and 0.3 the number of clusters appears to have a plateau. There the number of clusters does not rise with decreasing T , but—since the number of d-charges diminishes—the clusters must become smaller. At lower temperatures the number of clusters rises again saturating at 12 for $T \approx 0.01$.

In Fig. 5.17 the concentration of different types of clusters as a function of T is plotted. At high T many exotic cluster types are present. Besides the shown




Figure 5.16: The average number of clusters as a function of temperature.

cluster types others are present as well. Their concentration, however, is very low. At low T only two types of clusters remain: clusters with only d-charges -1 and clusters with d-charges ± 1 , of which the latter disappears completely at $T \approx 0.006$. Between $T \approx 1$ and 0.2 several types of clusters disappear.

The type of the clusters only shows which d-charges are present within the clusters. It does—for example—not show how many of the different d-charges are present. The clusters with only d-charges -1 essentially always consist of one single d-charge -1. It rises from 1 per cluster at low temperatures to 1.06 per cluster at $T=64$. In Fig. 5.18 we plot the composition of two important cluster types $(-1,+1)$ and $(-2,-1,+1)$ as a function of T .

Analysis and discussion

In Fig. 5.16 it is shown that the number of clusters is relatively small at high T . By visual inspection it is found that then indeed large clusters are present. Although in our cluster detection scheme only particles with non-zero d-charge are taken into account, it is observed that particles with zero d-charge can be trapped inside a large cluster. For the applied definition of a cluster, the number of clusters rises from 2.5 at $T \approx 64$ down to 12 at low temperatures. Then, in combination with Fig. 5.13, we find that each of these clusters consists of a single d-charge -1. The number of clusters and their composition as a function of temperature indicate that freezing occurs in stages.




Figure 5.17: The fraction of clusters of different types.

As shown in Fig. 5.18 the composition of clusters changes significantly with T . In general the number of d-charges present—and thus the size of the clusters—diminishes for lower T . However, as can be seen for example for cluster type $(+1,-1)$, far above the freezing transition the number of d-charges first rises as T decreases, and then drops again. The more exotic types of clusters show—for the temperature range considered here—only decreasing numbers of d-charges for decreasing T . The exotic clusters diminish in number for two reasons, first the number of d-charges decreases as a function of temperature and second, the exotic clusters with d-charges higher than ± 1 cost more energy and are thus not favoured at lower temperatures. These clusters break up into smaller clusters. The clusters with d-charges -1 and $+1$ are the most stable. These consist, at low temperatures, of 3 d-charges -1 and a single d-charges $+1$, with a total d-charge -2 .

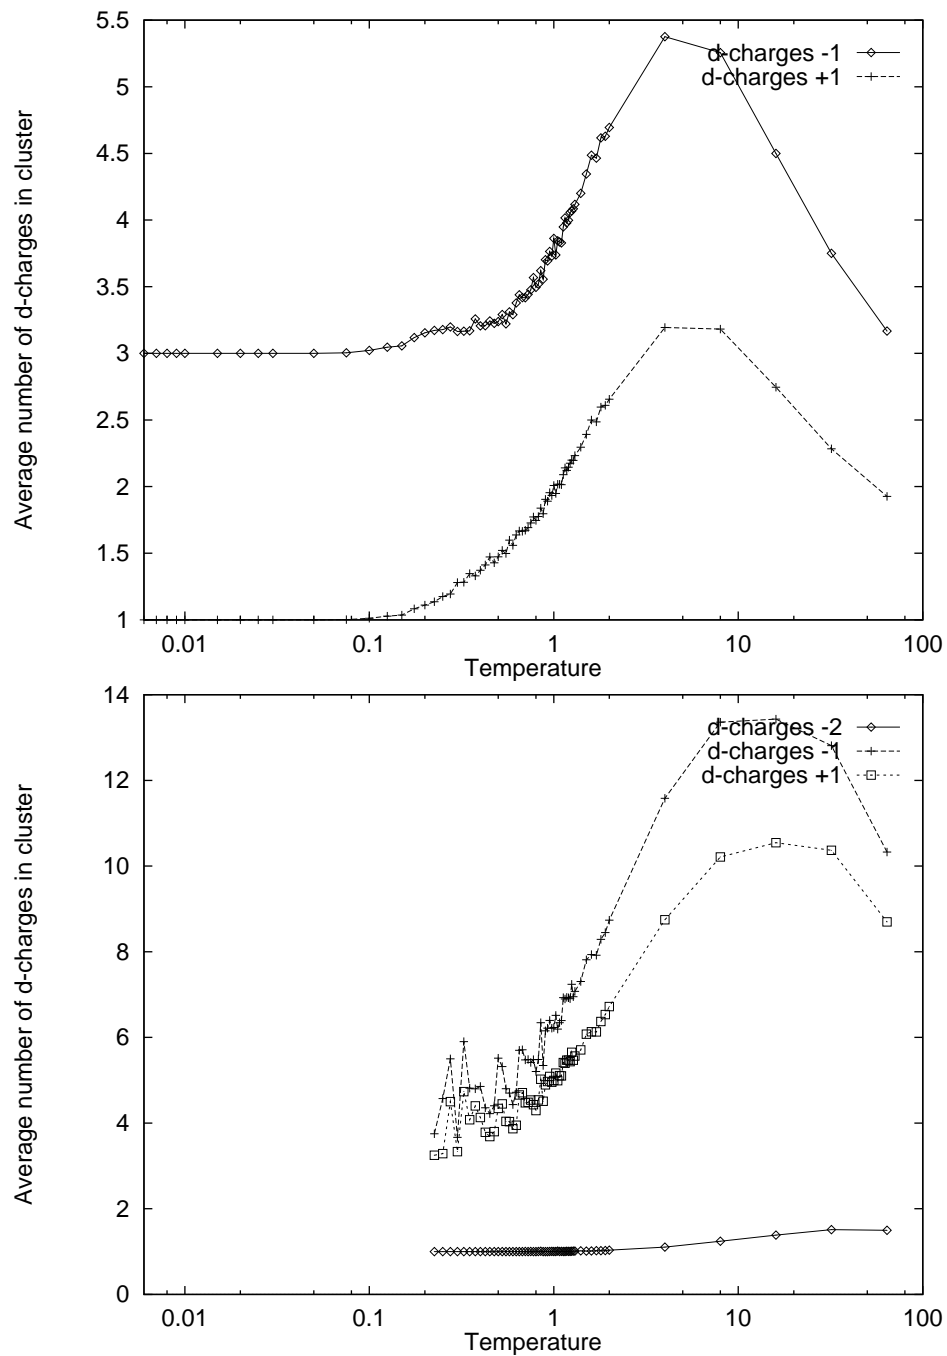


Figure 5.18: The average number of d-charges present in a cluster of the cluster type $(-1,+1)$ in the upper figure, and $(-2,-1,1)$ in the lower figure, in both cases as a function of temperature.

Chapter 6

Crystalline arrangements ($T=0$)

In this chapter we discuss zero temperature crystalline arrangements, formed by a 2D many-particle system with LJ interaction on the surface of a sphere. We investigate this system as a function of system size N at zero external pressure and at density ρ (by varying the radius of the sphere).

6.1 Introduction

In the simulations with fixed density (fixed radius) the Helmholtz free energy is reduced to the potential energy. For crystallisation processes where the radius is a free variable, the Gibbs free energy drives the system. Since we only apply this ensemble at zero external pressure only the potential energy E remains.

We start this part of the research by examining arrangements while the surface is changed from flat to spherical. We explore—in different ways—the energy landscape of some typical systems as a function of radius R . Subsequently, we describe the crystalline arrangements corresponding to the global minimum in potential energy with free radius at zero pressure, as a function of N . We show the behaviour of the energy, radius, defect distribution, symmetry and other properties of (meta)stable arrangements as a function of N . Next we examine the symmetries of arrangements formed at high density and compare them with the zero pressure results. Finally we study global energy minima (GEM) at zero external pressure for arrangements where a specific symmetry is built-in.

For completeness, unless otherwise indicated, we use the term “energy” for the average energy per particle. When we speak of details of the geometry of the structure formed by the N particles (for example their exact coordinates) we use the term “configuration”. Where we mean the 2D topology of the nearest neighbour connections, determined by the Voronoi analysis, we use the term “arrangement”. The arrangement remains unchanged when, for example due to a change in radius, the configuration changes in small details leaving the nearest neighbour connections essentially intact.

6.2 From flat to curved 2D space

In this section we study arrangements corresponding to minima in potential energy as a function of the radius [101]. This is achieved in two ways. Firstly by starting a system of N particles on a sphere with very large R , such that the surface is essentially flat. Then the curvature is increased while continuously relaxing the arrangement of particles to a (meta)stable arrangement. Secondly, by minimising the energy E of a system at a random—but during the minimisation fixed—radius R .

Since we are interested in minima in the potential energy, our results for both types of simulations correspond to zero temperature. The LJ potential has a minimum at unit distance; the—as a consequence—energetically favoured distance for neighbouring particles. If the surface of the sphere is considerably larger than typically N surface units (a surface unit is the area of the Wigner-Seitz cell [20] in a flat triangular lattice), the particles group together because of their mutual attractions, leaving part of the sphere open. Thus at radii significantly larger than the equilibrium radius R_{eq} , a patch of particles is formed covering only part of the sphere. For large N a major fraction of the particles is in the interior of the patch; these have low energies. For smaller N a larger fraction is near the edge of the patch and consequently the average energy is higher. The energy also depends on the shape of the patch and the internal arrangement.

At zero temperature, all configurations we find are in an energy minimum; all forces are balanced such that the particles do not move over the surface, and thus all particles are in minima of their local potential energy. As a function of R , for each (meta)stable configuration, the positions and depths of the local energy minima change.

The two types of simulation experiments are used to explore the energy landscape in relation to various properties of the arrangements, in particular, we want to find out how inherent d-charges (and others) are gradually being built-in (see Sect. 2.3).

Changing the radius may enforce a structural transition to a new (2D) topological arrangement if one or more particles are no longer in an energy minimum, the arrangement thus becoming locally unstable. Allowing relaxation at those radii, the arrangement “breaks up” and at least some particles undergo a reordering in their relative positions until an arrangement emerges that is (meta)stable at the new curvature.

On a flat surface no d-charges have to be present in the bulk; the net d-charge is zero. At the radius where the arrangement spans the sphere, a net d-charge of -12 has to be present (see Sect. 2.3). While increasing the curvature, d-charges thus have to be gradually incorporated. This can happen during the reordering of the particles at the transitions. In this first type of simulation experiment only one path in the E - R plane is traced.

A geometrical measure of how much a configuration changes after a step in

radius and the accompanying application of the Steepest Descent (SD) method, is the Euclidean distance D_e . This distance between two sets of coordinates (\mathbf{x}_1 and \mathbf{x}_2), is defined by

$$D_e^2 = \sum_{i=1}^{i=N} (\mathbf{x}_{1,i} - \mathbf{x}_{2,i})^2. \quad (6.1)$$

In the first experiment, to study the incorporation of d-charges, we start with a nearly circular patch of triangular lattice on a sphere with a very large radius R_i , an essentially flat surface. Because of the minuscule curvature the arrangement is slightly distorted, but not so much that d-charges or points with high stress are present. For an initial patch of N particles we take, from a flat triangular lattice, the closest $N-1$ particles around a central particle. If several particles are at the same distance while less are needed, an arbitrary choice is made. The initial arrangement is relaxed using the SD method, minimising the potential energy. Next we decrease the radius of the sphere successively in small steps until a given final radius R_f , while at each step the positions of the particles are again relaxed with SD, until the energy is minimal with a precision up to 12 digits. At each step the behaviour of the system is monitored.

The simulation experiment is repeated for a series of N -values up to $N=1000$. We show here typical results of simulations with $N=6$, 50, 350 and 500. For larger N only partly completed runs are available to check the validity of the general results.

In the second type of experiment we search for (meta)stable arrangements in the E - R plane. Random arrangements at randomly chosen—but during the minimisation fixed—radii, are minimised in energy. At any radius a large number of different arrangements can be metastable. In each simulation the potential energy of an initially randomly chosen arrangement is minimised using the Simulated Annealing (SA) and SD methods, while the randomly chosen radius is not altered. We show results of a 25-particle system. This number of particles is sufficiently small to allow a large number of simulation runs and it does not have a built-in tendency for an obvious global symmetry, like for example $N=24$ or 32 (see Sect. 6.3.4).

Experiment 1: stepwise increasing curvature.

When the curvature is increased in 1% steps, as discussed above, from an initially nearly flat nearly circular patch of triangular lattice to a spherically closed arrangement (completely covering the sphere), we find an energy behaviour as shown in Figs. 6.1a ($N=6$), 6.2a ($N=50$) and 6.3a ($N=350$). For $N=6$ and 50 we start with $R_i=10$, for $N > 100$ with $R_i = 100$.

During the build-up of stress, due to increasing curvature, the rising energy curve shows a series of peaks and sharp downward edges. These represent the

structural transitions described above. During the transitions stress is released by making changes in the arrangement. This takes the arrangement further from that on a flat surface, and more specifically, the packing becomes less efficient, causing a higher energy.

For $N=50$ we see a smooth rising curve for $R=10$ down to $R \approx 3.9$ (Fig. 6.2a). For smaller R , we see a series of transitions. At $R \approx 1.87$, the arrangement closes around the sphere, implying a low energy. This arrangement corresponds to that of the global minimum (see Sect. 6.3). Below that radius the system is compressed and—due to the LJ repulsive core—the energy rises steeply. For $N=350$ (Fig. 6.3a) we see a smooth rising curve from $R = 100$ down to $R \approx 12.3$. At $R \approx 5$, the arrangement closes around the sphere.

The smallest number of particles where, besides the global minimum arrangement, another stable arrangement is observed, is $N=6$. In Fig. 6.1 the potential energy for $N=6$ as a function of the radius is plotted. Only a few transitions (the first at $R \approx 1.36$) suffice to make a spherically closed arrangement. Two minima are observed in the E - R path. The minimum at $R \approx 0.93$ has an “open” (only partially covering the sphere) regular pentagonal structure with the sixth particle in the centre. Due to the curvature the sixth particle is not in the same plane as the pentagon. The minimum at $R_{\text{eq}} \approx 0.70$ is the global minimum, a “closed” configuration, which is reached if the particle positions and the radius together are optimised (see Sect. 6.3). This configuration is the regular octahedron.


In the E - R plot for $N=350$ (Fig. 6.3) many minima are observed. The lowest energy in this curve ($E \approx -3.13$) is found at $R_{\text{eq}} \approx 5$. Probably, the truly optimum arrangement, corresponding to the global minimum in E , is not found in this simulation (see Sect. 6.3). The number of minima with radius near R_{eq} rises exponentially with the number of particles (see Sect. 6.4), and in one simulation only one single route through the E - R plane is traced.

The potential energy of a 6-particle system as a function of R turns out to have a positive derivative for $R \gg R_{\text{eq}}$. Therefore the start arrangement is not stable if the radius is a free parameter. The system can decrease its energy by forming a non-flat arrangement. For the $N=50$ system we see that the energy increases when R is decreased. This suggests that the regular lattice is stable for large R . For $N = 350$ we also find a slight decrease of energy for large R (not visible in Fig. 6.3). Therefore this system also does prefer initially a slightly curved surface.

During transitions between arrangements the relative positions of particles change significantly. A large Euclidean distance (Eq. 6.1) is travelled before a new (meta)stable arrangement is reached. Between two transitions the arrangement is not significantly altered, the particles only shift slightly and smoothly in position. Thus a small Euclidean distance between the configurations of successive steps is found. In Figs. 6.1b, 6.2b and 6.3b (for $N=6$, 50 respectively 350) we plot the Euclidean distance travelled before reaching a minimum in potential energy after a decrease in radius. We see that more transitions are present for larger N and

they start at larger R .

(a)



(b)






Figure 6.1: The energy (a) and the Euclidean distance travelled in the relaxation after the steps (b) for $N=6$ as a function of the radius, where the system is started from a nearly flat regular arrangement towards smaller radii.

The structural transitions that take place at $R > R_{\text{eq}}$ are usually not just a shuffle of particles near the edge of the patch, but rather a global—topological—rearrangement. In Fig. 6.4 we show two configurations of $N=500$ at two different radii that differ only one step (1%), $R=14.794$ and $R=14.646$. In the dark parts of the arrangement the particles have a low energy, in the gray parts the particles have a relatively high energy. This figure illustrates the large change in 2D topological packing that can occur in a structural transition. The particularly well


(a)




(b)

Figure 6.2: Same as Fig. 6.1 for $N=50$.

(a)



(b)

Figure 6.3: Same as Fig. 6.1 for $N=350$.

packed parts differ indeed very much between the two successive arrangements. The first transition for $N=500$ takes place at $R \approx 15.6$. Therefore it is not surprising that the arrangement with $R = 14.794$ in Fig. 6.4 already shows a defect at the bottom left. At $R = 14.646$ we see that groups of well packed particles are formed, surrounded by less well packed regions.

Experiment 2: minimisation at fixed random radii.

To exclude any bias due to the history of the system, independent simulations are performed with random, but during the simulation fixed, radius. Thus, instead of a single (meta)stable arrangements per R , we find many of them, corresponding to local minima in the E - R plane.

During the SA phase, 20 Monte Carlo Metropolis moves per particles are used to equilibrate the system at each temperature (T). Then T is decreased with cool rate 0.95. The annealing is stopped when the deviation in the energies of the last generated arrangement of the last 10 T -values falls below 10^{-6} . On average 350 chains need to be generated before a minimum is reached. Then the SD method is used to increase the precision up to 12 digits. In total we here present 1300 simulations at randomly chosen radii between $R=1.3$ and $R=2$.

In Fig. 6.5 we plot the found minima as a function of R for $N=25$. In this figure we see the global energy minimum (GEM) appearing at $R_{\text{eq}} \approx 1.37$, as confirmed by our minimisation simulations discussed in Sect. 6.3.

For $R < R_{\text{eq}}$ the system is compressed and the energy rises steeply. If R is significantly greater than R_{eq} the energy is higher than around $R = R_{\text{eq}}$ because of the earlier mentioned edge effect. For R near or below R_{eq} only few different minima are present (see also Sect. 6.4). For larger R a large number of possible arrangements shows up. The difference in energy between the highest and lowest found energy, at a particular radius, is also much higher in this R -range.

Another clearly visible and interesting second minimum is found at $R \approx 1.52$ with energy $E \approx -2.66$. Also here the system cannot close around the sphere and forms a semi-regular open arrangement, as shown in Fig. 6.6. Due to the specific choice of $N=25$, the free radius equilibrium configuration has no symmetry (see Sect. 6.3). The open arrangement of Fig 6.6, however, has a C_{3v} symmetry (see also Sect. 6.4).

In Fig. 6.5 clearly many points fall along distinct curves. These curves appear to correspond to a single (meta)stable arrangement with only secondary differences in configuration as a function of R . The lowest points correspond to deepest minimum energy arrangements as a function of R . These form a series of arrangements, each stable within a specific range of R . Remarkably, judging from the density of points in Fig. 6.5, the deepest minimum as a function of R is not necessarily the most probable arrangement found, indicating that it often has a smaller capture basin than some local minima.




Figure 6.4: Two configurations (top $R=14.794$ and bottom $R=14.646$) differing only one step (1%). Clearly a structural transition to another arrangement occurs. The dark coloured particles have low energy, the gray a relatively high energy.




Figure 6.5: The energies of arrangements with $N=25$ obtained by minimisation of random initial configurations, while keeping the radius fixed to a random value.




Figure 6.6: The open configuration for $N=25$ at $R = 1.52$, discussed in the text. Top view (left) and side view (right). The vertices are positions of particles.

Combining the two experiments.

Starting from any of the $N=25$ configurations in Fig. 6.5 we can perform a simulation where the radius is changed in small steps with relaxation, applying the SD method after each step. As an example we use the open configuration shown in Fig. 6.6. Starting from this configuration we let R run up and down between $R=2.1$ and $R=1.3$ while applying SD until the system is back at the starting radius. An interesting curve is traced in this way, see Fig. 6.7. The arrows indicate whether the radius is increased or decreased. As in Figs. 6.1a, 6.2a and 6.3a we obviously trace a path in the E - R plane with transitions between different arrangements. Each of the arrangements is (meta)stable in some range of the radius. Since we now go up and down in R , for some arrangements we can trace the whole of the range of stability. When this procedure is applied to a number of arrangements we find that many of them exhibit individually a minimum in E as a function of R . Prominent minima obtained in this way are shown in Fig. 6.8.

In Fig. 6.7 the path traced in the E - R plane reveals hysteresis effects, demonstrating the relevance of combining the two experimental approaches. The inset in Fig. 6.7 shows a simple case of hysteresis.




Figure 6.7: Energy of an $N=25$ system if the radius is changed while relaxing with the SD method. The inset shows a simple case where hysteresis occurs.




Figure 6.8: The energies of a number of arrangements with $N=25$ showing a minimum in energy as a function of the radius.

Analysis

In the simulations where an initially circular patch of triangular lattice is curved in steps while relaxing with SD after every step, we can distinguish several points of interest. First, the energy $E(R)$ and its derivative for the arrangement at large radii. Secondly, the radius where the arrangement has its first structural transition. Finally the radius where the arrangement closes around the sphere. For these three points a theoretical analysis is given below.

The energy and its derivative at large radii

The energy of an arrangement at large R , in the first type of simulations, is largely determined by the amount of particles in the bulk compared to those on the edge. To estimate energy at large R in a simple model, we suppose that the arrangement is a circular patch of triangular lattice for which the edge is of unit width, the lattice distance for LJ particles. By using the surface area we can calculate the number of particles in the bulk N_{bulk} and on the edge N_{edge} . The energy can now be estimated by determining two typical energy calibration constants, one for particles in the bulk E_{bulk} and one for particles on the edge E_{edge} . The energy E is equal to $(N_{\text{bulk}}E_{\text{bulk}} + N_{\text{edge}}E_{\text{edge}})/N$. The value of E_{bulk} can be calculated for an infinitely large flat triangular lattice and is equal to $-3.382\dots$. The particles on the edge have less neighbours and therefore have

N	E_{model}	E_{data}
6	-2.162	-1.560 ($R=10.0$)
25	-2.591	-2.446 ($R=2.00$)
50	-2.784	-2.712 ($R=10.0$)
350	-3.134	-3.134 ($R=100.0$)
500	-3.173	-3.178 ($R=100.0$)

Table 6.1: Comparison between the simple model and simulation data on the energy at large radii.

higher energy. The value can be estimated by a single calibration. If we use $N=350$ for calibration ($E=-3.134$ at $R=100.$) it follows that $E_{edge} = -2.099$. In Table 6.1 we give the simulation results and the simple model predictions. For $N=25$ the lowest point in Fig. 6.5 for $R=2.00$ is taken, although the figure strongly suggests that an arrangement with lower energy $E \approx 2.47$ (closer to the model estimate) is possible at that radius. We see in Table 6.1 that for the larger N the model fits reasonably well while for small N it is much too low. Clearly—and understandably—the edge is effectively much more pronounced in small systems than assumed in the simple model.

For $N=6$ the derivative of the energy at large R is positive while for the other shown cases it appears to be negative. There are several influences on the behaviour of the derivative of the energy at large R . Firstly, if the radius is changed—in spite of relaxation—stress can build up, which costs potential energy. Secondly, if the surface is curved the distances between non-neighbouring particles decreases which—due to the LJ attractive tail—lowers the energy. Thirdly, with increasing curvature the edge (circumference) of the patch decreases and so does the unfavourable edge energy. With increasing curvature the build-up of stress in the bulk is opposed by the increased binding in the bulk and reduced edge contribution. For example for the tiny $N=6$ system enough room is present to take full advantage of the edge effect. For $N=50$ we observe that the build-up of stress is more important than the energy gained. For $N=350$ the energy first decreases slightly due to the attractive part of the potential and then rises much faster due to the build-up of stress.

The radius of the first structural transition

The radius where—at increasing curvature—for the first time the arrangement undergoes a significant rearrangement (R_u), can be approximated quantitatively. For this we assume first that we start our simulation with a hexagonally shaped (rather than circular) flat triangular patch with unit lattice constant, as for example in Fig. 6.9. The hexagonal rings are numbered by n , the largest complete ring has number n_{max} . We consider the distance between the radial lines originating from the central particles and having angle $\pi/3$, as indicated in the figure. If the




Figure 6.9: A patch of flat triangular lattice with $n_{\max}=3$ hexagonal rings.

patch is placed on a large sphere, such that no defects are included, the distance between the lines at the n -th ring is less than n . With increasing curvature the rings are being compressed and the arrangement can become unstable, implying that a global rearrangement is triggered.

In our case we have circularly rather than hexagonally shaped patches. Therefore we have to find a relation between n_{\max} and N . By considering the surface area per particle we find that

$$n_{\max} = \left\lfloor \sqrt{\frac{N\sqrt{3}}{2\pi}} - \frac{1}{2} \right\rfloor, \quad (6.2)$$

where $\lfloor a \rfloor$ means the largest integer below the value a (floor).

To find at which compression the system rearranges, we use $N = 350$ for calibration. There the first structural transition takes place if the distances in the $n_{\max} = 9$ ring are smaller than 0.913. Using this value we can find a relation between n_{\max} and R_u :

$$0.913 * 2n_{\max} \arcsin(1/(2R_u)) = \sin(2n_{\max} \arcsin(1/(2R_u))), \quad (6.3)$$

where the sine and arcsine enter from measuring distances in 3D through the sphere. The equation is well behaved in the range of interest and has only one root there. It can be solved easily with a simple half step refinement method.

Having a relation between N and R_u we can predict at which radius a patch of N particles becomes unstable. In Fig. 6.10 we show R_u as a function of N . Since N and R_u are related through n , discrete steps appear in Fig. 6.10.




Figure 6.10: The predicted radius R_u where a circular patch of triangular lattice becomes unstable as a function of N .

N	n_{\max}	R_u (data)	R_u (Eq. 6.3)
6	1	1.37	1.34
50	3	3.95	4.11
350	9	12.3	12.3
500	11	15.6	15.0

Table 6.2: The radius R_u where the arrangement becomes unstable as derived from our simple model and from the data given in this section for $N=6, 50, 350$ and 500 .

The value of R_u for $N=6, 50, 350$ and 500 , from the model and data are given in Table 6.2. The model predictions are in good agreement with the data.

When the arrangement has become unstable at $R \approx R_u(N)$ a transition to a new (meta)stable arrangement is made. During the relaxation of the arrangement at the transition radius, d-charges can be incorporated into the arrangement and local “stress” that has built-up can be released and/or spread over the system. By comparing the $E-R$ curves for $N=6$ with those for $N=50$ and 350 we observe that for larger N more (meta)stable arrangements become possible along the radius axis. From the data on the Euclidean distance, we find only changes in geometrical details of the arrangement, while at the the sharp peaks the arrangement moves to a structurally (i.e. 2D topologically) different arrangement.

The radius where arrangements completely cover the sphere

The radius R_c at which the arrangement closes around the sphere can be approximated by requiring the density to be essentially that of an optimum flat

lattice:

$$R_c(N) = \sqrt{\frac{N}{4\pi} \cos\left(\frac{\pi}{6}\right)}. \quad (6.4)$$

For $N=6, 25, 50$ and 350 it follows from Eq. 6.4 that $R_c(N)=0.64, 1.3, 1.9$ en 4.9 respectively. Especially for larger N these values are in reasonable agreement with the data shown in Figs. 6.1a, 6.2a, 6.3a, and 6.5. For small N , however, the packing is far less optimal than that of a flat triangular lattice and the estimated R_c is too small (see also Sect. 6.3.2).

At large radii the Euclidean distance between successively relaxed configurations with changes in radius, 1% for each step, decreases smoothly, as for example shown in Fig. 6.11 for a system with $N=500$. If only the 1% change in radius—without any reshuffle—is taken into account, the Euclidean distance $D_{e,r}$ is

$$D_{e,r} = \left(\sum_{i=1}^{i=N} (0.01R)^2\right)^{1/2} = \sqrt{N} * 0.01R. \quad (6.5)$$

In Fig. 6.11 we plot this curve and the experimental Euclidean distance for $N=500$ as a function of radius. For asymptotically large R we have $D_e = D_{e,r}$. Below $R \approx R_c$ the change in Euclidean distance is mainly due to the change in radius, the particles hardly move over the surface, except during an occasional transition visible in that regime.

There exists an analogy between the increase of curvature, where stress is introduced externally, and the sand-pile model of self-organised criticality [156] where instabilities are created by adding sand grains. Similarly, stress is built-up due to tectonic movement causing sudden structural transitions, earthquakes. For a sand pile the amount of avalanche as a function of size is distributed as a power law. Examining the analogy, as a size measure for the “avalanches” in our system we use $D = D_e - D_{e,r}$ between successive arrangements. We indeed find a power law behaviour for $D > 0.1$. In Fig. 6.12 we show the size distribution of D in 20 bins for $N=350$ where the data from Fig. 6.3b for $R=14$ down to 4 are used. The line plotted to guide the eye ($8.96D^{-1.29}$) is determined by least square fitting where it is assumed that a \sqrt{N} deviation is present in the data (statistical noise). A second property of systems showing self-organised criticality is that arbitrary small perturbations can trigger a large change. This is also the case in our simulations. If the system during increasing curvature is in a state just where the stability of the current arrangement ends, an arbitrary small change in the position of one or more particles triggers a structural transition.

Discussion

In the studies above the system traces only a single path through the E - R plane. If, however, we would use different start arrangements we might trace




Figure 6.11: The Euclidean distance between successive arrangements at increasing curvature for $N=500$. The dashed curve represents the analytically calculated contribution to the Euclidean distance simply due to the change in radius. The configuration completely covers the sphere for $R < 5.9$.




Figure 6.12: The size distribution of the Euclidean distance travelled by the particles during the transitions for $N=350$ between $R=14$ and $R=4$.

other paths. This could be studied by starting with for example an arrangement that is not a circular patch of triangular lattice but elliptic or arbitrarily shaped, although many of these tend to reshape themselves to a nearly circular arrangement. Energetically it is favourable at low density to form a roughly circular patch on the curved surface. Then the ratio of bulk to edge particles is optimised.

In the simulations where a random start arrangement is used with either the SD method alone or with the combination of SA and SD, at a random but during minimisation fixed radius, we can find in principle all of the possible arrangements.

Examining a number of open configurations at large R we find that they are indeed nearly circular. However, this circular patch can be in many different (meta)stable configurations. This is reflected by the difference in energy observed in Fig. 6.5 for large R . For smaller R we observe that there is less variation in the energy. We find that the number of different minima is small near R_{eq} (see Sect. 6.4). This might be due to the increased constraints resulting from a closed spherical configuration.

Many arrangements are (meta)stable only within a certain range of the radius and show a minimum at a specific radius.

Although we only have shown results for systems with $N=6, 25, 50, 350$ and 500 , qualitatively most of the observed phenomena are representative also for systems with other—including much larger— N -values.

We find that the incorporation of topological defects, in particular d-charges, is accomplished by structural transitions, substantial rearrangements of the particles. In such cases barriers can exist between arrangements with different numbers of d-charges. Also, when only slightly changing the curvature, the system can be forced to rearrange and to distribute “stress”, without necessarily incorporating new d-charges. Hence not all structural transitions are related to d-charge incorporation. There can be more than 12 (the maximum number of inherently needed d-charges) peaks in the Euclidean distance (see Fig. 6.3a) over the whole curvature variation.

Some of the arrangements, even with $R \neq R_{\text{eq}}$, can exhibit interesting symmetry. Examples are the second minimum for $N=6$ in Fig. 6.1 and the minimum at $R = 1.52$ (shown in Fig. 6.6) for the $N=25$ system.

The data show that several (meta)stable arrangements, corresponding to minima in potential energy, exist as a function of curvature. When R is a free parameter, the system may be trapped into any of the minima as shown in Fig. 6.8.

Therefore, the deepest minimum as a function of R does not correspond to one single topological homogeneous arrangement. In contrast to for example the Coulomb potential, we find that arrangements exist that are stable only in specific ranges of R . From the concentrations of the points in the E - R plane (see Fig. 6.5) for the $N=25$ system, we infer that an arrangement with the lowest energy is not always the one with the largest population probability. This means

that the capture basin of the deepest minimum is often not the largest in state space.

6.3 Global potential energy minima

In this section we present our results on potential energy minima at zero external pressure for a system of N LJ particles on a spherical surface with freely variable radius. These minima are assumed to correspond to the global energy minimum (GEM) in the E - R plane. Results for $2 \leq N \leq 200$ are shown. The results on global *and* local minima are presented in Sect. 6.4.

To study the influence of curvature on the crystalline state of the system, we calculate the energy, radius, symmetry group, d-charge distribution, dipole moment and enantiometry of the global minima (see Sect. 2.6).

Properties of configurations corresponding to the global minima have a strong dependence on N . The global minimum configuration can be predicted for some small N -values permitting special symmetry properties. In general, however, it is impossible to prove—or even anticipate—analytically which configuration corresponds to the global energy minimum.

In searching for the GEM arrangements in the energy landscape over the state space we employ two methods: Simulated Annealing (SA) and Steepest Descent (SD), both discussed in Chapt. 3. The system is crystallised using the SA method. At the temperature where the system is supposed to be definitely trapped inside a minimum, a switch from SA to SD is made.

For the LJ potential the radius of the configuration corresponding to the precise GEM arrangements is not known in advance. Therefore, in addition to the positions of the particles, the radius has to be optimised in a NPT ensemble with zero external pressure (see Sect. 2.5). The SA method starts at a temperature T_i where the average radius of the constant zero pressure system in equilibrium is significantly larger than that of the GEM arrangement. For reasons of computational efficiency we restrict the range in which the radius can change. Over the first 20 Markov chains, the high temperature part of the minimisation with SA, the radius is kept constant. The simulation conditions then correspond to a NVT ensemble. In the low temperature part of SA the radius is free to change however within a judiciously chosen range. Initially the radius remains near the largest allowed value. At a certain temperature the radius starts converging to its final value. Thus in our simulations a transition from NVT to NPT ensembles is enforced.

The results shown in this chapter (see Sect. 6.3.2) indicate that the range of the radius for $N < 40$ can be taken as $R_0 \leq R \leq R_0 + 0.15$ and for $N \geq 40$ we can use $R_0 \leq R \leq R_0 + 0.06$ with R_0 the radius where the density is equal to the

density of an optimally packed flat surface, given by

$$R_0(N) = \sqrt{\frac{N}{4\pi} \cos\left(\frac{\pi}{6}\right)}. \quad (6.6)$$

The number of Metropolis steps per particle (S_e), necessary to equilibrate the system at each temperature, is a function of N (see Sect. 3.6). For $N < 30$ we use a fixed number of steps, typically $S_e=20$. For $30 < N < 100$ we typically use $S_e=200$ and for $N \geq 100$ we use up to $S_e=2000$. For our system we use starting temperatures $T_i = 10$. Then—with a cool rate of 0.95—we anneal until $T_f = 0.01$ (see Sect. 3.6). These values are not very sensitive to the number of particles used and not varied as a function of N .

During SA we use a cut-off and interaction list method (see Sect. 3.5). During SD we do not use the cut-off method, in order to obtain at this stage high precision energies. After SD the precision of the energy is 12 digits. This precision is necessary to discriminate between nearly degenerate local minima. The radius of the system is then precise to 8 digits.

The parameters of SA are such that there is no guarantee that the GEM arrangement is found in a single run. Therefore we have to do many independent searches through the state space. The amount of independent searches necessarily depends heavily on N . The number of different minima that exist for a specific N value is an indication for the complexity of the problem (see Sect. 6.4).

We find that about 200 independent runs suffice for $N \leq 100$ and about 800 for $100 \leq N \leq 120$, but for larger N thousands are necessary to be confident that the GEM arrangement is found. If the lowest energy minimum found occurs several times within a series of independent runs we gain some trust in having actually found the GEM arrangement.

Even with many independent searches still one can not be 100% sure that the GEM arrangement is found. It is possible that it has a very small capture basin and only is found when an extremely large number of searches is made.

6.3.1 Energy as a function of N .

In Table A in the Appendix, properties of the found GEM (or rather the best candidate GEM) arrangements for $2 \leq N \leq 200$ are listed. We list the energy E_g , radius R_g , d-charge distribution and symmetry group S. In this section we discuss these (and several other) properties of the data. For clarity we repeat here that with the term energy we mean the normalised energy (energy of the system divided by the number of particles), unless stated otherwise.

At first sight there appears to be no general rule for packing an arbitrary number N of particles in an energetically optimum way on a spherical surface. Therefore each number of particles can show-up in highly specific arrangements with associated differences in the energy. In general, if N particles can be packed

optimally efficiently (homogeneously and compactly) on the spherical surface, the whole system has a relatively low energy.

For $N=2,3$ and 4 the configurations can be found by reasoning and their energy calculated analytically. For these trivial cases we find two particles at unit distance, a triangle with unit distance and a tetrahedron with unit distance, and with energies $E_g = -0.5, -1.0$ and -1.5 , respectively. The fast decrease of the trend in energy at very small N is due to the number of interactions per particle. For $N > 4$ with each additional particle less than N new interactions at unit distance are formed, and the energy drops less than 0.5. For $N > 12$ particles with d-charge 0 are formed. From then on the ratio of particles with d-charge -1 to particles with d-charge 0 can be expected to cause a general trend in the energy as a function of N .

Results

In Fig. 6.13 we show the candidate GEM $E_g(N)$ -values. The main trend for increasing N indeed is towards lower energies. For relatively small N there are large jumps in energy. Some N -values clearly can form arrangements with lower energy than neighbouring N -values. For large N , where a constant energy is approached, variations diminish on an absolute scale but on a relative scale they do persist.

Analysis

The differences in $E_g(N)$ due to changes in the packing are emphasised by modelling the above described general trend by a smooth function and subtracting that from $E_g(N)$. Our model is based on the following arguments:

1. In principle most particles in a large spherical configuration have relative positions comparable to those in an optimum flat 2D triangular lattice. In zero order approximation at large N the energy per particle thus is a constant, here denoted a
2. On a sphere a net d-charge of -12 is present. In the ground state this is usually achieved by having 12 “spots” in the spherical arrangement each with d-charge -1 (see also Sect. 6.3.3). These “unfavourable” spots represent extra energy. In first approximation their contribution to the model function is assumed constant, contributing b/N per particle.
3. The influence of the curvature on the energy (bending energy), as a deviation from the flat lattice contribution (point 1), is inversely proportional to the square of the radius [157]. Integrated over the surface of the sphere and




Figure 6.13: The energy of the GEM arrangements as a function of N . The inset shows an enlargement of the range $40 \leq N \leq 200$. The individual data points represent the lowest energy found in at least 1000 simulation runs per N -value.




Figure 6.14: Non-normalised difference between the fitted function and data.

divided by N , this contribution to the model function is also of the form b/N .

4. If the spots with net d-charge -1 are in fact grain boundaries, i.e. strings of alternating d-charges (see Sect 2.7.2), they have an excess energy per unit length [54] (for small angle grain boundaries). If the length of these defects scales with the size of the sphere, they contribute to the model function a term of the form c/\sqrt{N} .

Thus our model $E_m(N)$ for the general trend of the ground state energy is

$$E_m(N) = a + \frac{c}{\sqrt{N}} + \frac{b}{N}. \quad (6.7)$$

This model contains 3 constants that have to be determined by fitting. The data can be assumed to be known to a sufficiently high precision, to interpret any local deviations from the general trend as due to a specific N -dependence—in particular the effect of packing constraints—not included in our model.

We use a least-squares fit of the model on the data $E_g(N)$ in Table A to determine the constants. We find:

$$a = -3.24, b = 4.80, c = 1.11$$

In Fig. 6.14 we plot the total deviation energy $E_d(N) = (E_g(N) - E_m(N))N$. Configurations with $E_d > 0$ are less optimal than expected from the general trend. In this case, by exception, we plot the total (rather than normalised) energy to compare deviations at small N with those at large N . A remarkable change of general trend is observed around $N=110$, in addition to fluctuations at specific N -values.

Discussion

In modelling the general trend we find several interesting aspects. The value of a in the model could be expected to correspond to the lowest energy of LJ systems for $N \rightarrow \infty$. Based on the data we observe that the value of a is slightly above the value found for a flat infinite triangular lattice, which is -3.382... . Constants b and c are positive, indicating that the properties they stand for cost energy. Except at small N , the $1/\sqrt{N}$ term is more important than the $1/N$ term. The relatively high value of c indicates that line defects play an important role.

That a in the general trend $E_m(N)$ does not approach the energy of an optimum flat 2D lattice might be due to the influence of d-charges growing with N . This possibility is supported by the indication that line defects are important.

After removing the general trend from our data (see Fig. 6.14) we find that for $N < 100$ some specific values of N give considerably lower ground state energies

than those of the general trend. Extreme examples of these are $N=12$, 24, 32, 44 and 48.

For $N < 100$ we observe a large variation in energy. This variation must be due to the 2D packing efficiency on a spherical surface, which has a strong N -dependence, and which is not taken into account in our model. For $N > 100$ most particles have no d-charge and are positioned in a locally nearly triangular lattice. Then the packing efficiency does no longer change erratically as a function of N and thus the variations in energy between nearby N -values are smaller. Around $N=110$, however, we find an unexpected behaviour. A clear downward trend is observed. We can expect that something interesting happens in the configurations around $N=110$.

The above described behaviour of the total deviation of the system energy from the model fit suggests that some significant phenomenon is not included. In later sections it will become clear that it is the contribution of the “spots”, the d-charge clusters, that are modelled continuous while in our data they appear to change abruptly, and to be connected to the overall 2D packing efficiency of N particles on a sphere (see Sect. 6.3.3).

6.3.2 The radius as a function of N .

In this section we examine the radius $R_g(N)$ ($= R_{eq}(N)$) of the obtained GEM arrangements as a function of N as listed in Table A.

At constant pressure the radius of the sphere is dependent on the number of particles because of the area associated with a LJ particle. Here we use zero external pressure. The radius is also dependent on how efficiently the particles can be packed on the sphere.

Results

According to general expectation for 2D spherical arrangements, and to the data in Table A, the radius increases approximately as the square root of the number of particles. For all N we find that $R_g(N_1) < R_g(N_2)$ for $N_1 < N_2$. If the general trend is removed by subtracting R_0 (Eq. 6.6) from $R_g(N)$, giving $R_d(N) = R_g(N) - R_0$, we find the behaviour shown in Fig. 6.15. For small N the radii are on average considerable larger than R_0 and they display also large variations in R_d for nearby N -values. For larger N the density, $\rho = N/(4\pi R^2(N))$, comes closer to the density of the optimum triangular lattice on a flat surface ($\rho_0 = \cos^{-1}(\pi/6) \approx 1.15$) and the variation becomes smaller. For $N > 100$ there is a tendency for larger radii than expected from the trend in $R_d(N)$. For $N > 160$, $R_d(N)$ shows a tendency to decrease.




Figure 6.15: Difference between the found optimum radius and the radius derived from a flat lattice (R_0) as a function of N .

Discussion

The region $N < 80$ can be characterised by an average decreasing trend for $R_d(N)$ from 0.1 to 0.04 with large fluctuations for individual N -values. In particular $N=12, 32, 44, 48$ but also $22, 24, 38, 60, 72, 78$ approach the optimum flat packing conditions to within $\approx 4\%$, while neighbouring N -values have mostly much larger R_d . In the region above $N=80$ fluctuations diminish, broaden and the average trend becomes constant at the level of $\approx 4.5\%$.

By comparing the data on $E_g(N)$ with $R_g(N)$ we find that configurations with energies lower than E_m usually also have relatively small radii while configurations with energy higher than E_m have less dense packing.

6.3.3 Number and distribution of d-charges

In this section we discuss the d-charges present as a function of N in the GEM configurations as listed in Table A, for d-charges -2, -1, 0, 1 and 2. The d-charges are determined by the Voronoi analysis (see Sect. 2.6.3). For particle distributions over a spherical surface the net d-charge must be -12, but this allows many different combinations including an arbitrary number of d-charge neutral pairs.

Results

In Table A we can check that the net d-charge is -12 (except for $N=2, 3, 4, 5$ and 7 with d-charges larger than ± 2 , which are not listed). Although the configuration with $N=8$ has 8 particles in equal positions (it forms a twisted cube), Table A shows that there are 4 d-charges -2 and 4 d-charges -1. The arrangement is “degenerate”, meaning that it depends on an infinitesimal distance whether a particle is considered a neighbour or not. Because the coordinates are determined with finite precision the Voronoi construction is not hindered by this degeneracy.

Although 12 d-charges -1 would suffice to close an arrangement around the sphere, we observe that the total number of d-charges $\neq 0$ grows with N . In




Figure 6.16: The number of d-charges $\neq 0$ as a function of N .

Fig. 6.16 we show the total number of d-charges present in the configurations, and in Fig. 6.17 the number of d-charges -1 and +1. The other d-charges (larger than ± 1) hardly play any role in the energetically optimum distributions with $N > 12$. Roughly four regions in N can be distinguished for which the number of d-charges shows a different trend. For $N < 12$ the number of d-charges -1 grows towards the value 12, for $12 \leq N \leq 105$ it is mostly 12 with excursions to a maximum of 24, for $105 \leq N \leq 160$ it rises to 24, and finally for $160 \leq N \leq 200$ it is approximately 24. Of course the excess d-charge caused by having more than 12 d-charges -1 is compensated by d-charges +1.




Figure 6.17: The number of d-charges -1 and +1 as a function of N .

Analysis

For arrangements with more than 12 d-charges -1, clusters are formed. For example in the GEM arrangement found for $N=100$, in total 12 single d-charges -1 are present, while for $N=200$, 12 d-charge clusters with net d-charge -1 are present, each consisting of 2 d-charges -1 and one d-charge +1 (see Fig. 6.18), thus in total 36 single d-charges $\neq 0$. In the region $105 \leq N \leq 160$ different types of clusters occur. There are free d-charges -1 and there are clusters of d-charges -1 and +1.

The question then arises how the d-charges are spread over the spherical arrangement. Since like-sign d-charges tend to repel each other we can expect that the necessary 12 disclinations are distributed approximately as an icosahedron. By inspection of the configurations we find that indeed this is in general the case.

We have seen that above a certain value of N , clusters of d-charges seem to be energetically favourable. There is a transition range in N where the number of d-charges changes from 12 ($N \lesssim 105$) to 36 ($N \gtrsim 160$). The system appears to strive towards 12 equal clusters of d-charges; therefore the d-charge distribution shows a discrete behaviour. The number of d-charges within one cluster (L) is thus odd. First one (inherent) d-charge -1 is present, then a d-charge pair ± 1 is added. In general it can be expected that the 12 clusters consist of n neutral d-charges pairs with $n = 0, 1, 2, \dots$ and one inherent d-charge, thus the d-charge cluster size $L = 2n + 1$.




Figure 6.18: The Voronoi polygons for configurations with $N=100$ (left) and 200 (right). Dark shaded polygons are d-charges 0 , lightly shaded polygons are d-charges $+1$ and white polygons are d-charges -1 .

We can derive a formula that approximately gives the N -value at which the transition as a function of L take place if we concern ourselves with the clusters of d-charges and the d-charges inside, treating zero d-charge particles (the lattice) as a “medium” in which the d-charges are embedded. It is assumed that 12 clusters, each with a net d-charge -1 are distributed approximately icosahedrally. We base our derivation on the surface taken by the clusters where all curvature stress is supposed to be concentrated. If these clusters consist of single d-charges -1 , the amount of surface where the curvature is hidden is only a small fraction of the total surface, giving a high stress and costing much energy. If clusters of d-charges are present the curvature stress is spread over a larger fraction of the total surface, lowering the concentration of stress and thus energy. On the other hand to produce a d-charge pair, energy is needed.

We consider the energy costs of having clusters E_{cluster} and the energy costs E_{area} related to the area where the stress is concentrated. We assume the cluster energy to be proportional to the linear cluster size [19, 54], thus

$$E_{\text{cluster}} = \alpha\sqrt{L}. \quad (6.8)$$

If we assume that the energy is inversely proportional to the fraction of the surface where in our model all stress is concentrated, we find

$$E_{\text{area}} = \beta \frac{N\sqrt{3}/2}{12L}. \quad (6.9)$$

L	$L + 2$	N
1	3	130
3	5	448
5	7	849

Table 6.3: N -values where, given our simple model, a transition in the total number of d-charges per cluster is expected.

If we put the two terms together we have

$$E(N, L) = \alpha\sqrt{L} + \beta\frac{N\sqrt{3}/2}{12L}. \quad (6.10)$$

From our data we already know that the transition from $L = 1$ to $L = 3$ occurs between $N=100$ and $N=160$. To remove one of the parameters from the above given formula we use $N=130$ as the transition point (see Fig. 6.16). If we put $E(N = 130, L = 1) = E(N = 130, L = 3)$ we derive $\alpha = 8.54\beta$. Then

$$E(N, L) = 8.54\beta\sqrt{L} + 0.0722\beta\frac{N}{L}. \quad (6.11)$$

The N -value where the transition in cluster size L to $L + 2$ takes place is determined by $E(N, L) = E(N, L + 2)$. The constant β drops out and we find

$$N = 118.4\frac{L(\sqrt{L+2} - \sqrt{L})}{1 - \frac{L}{L+2}}. \quad (6.12)$$

In Table 6.3 we show some N -value where transitions are expected given our assumptions in the model. These N -values, however, are too large to be validated with our present results.

Discussion

The regions of N in terms of d-charges (Fig. 6.16) can be correlated with the $R_d(N)$ data (Fig. 6.15) and with the energy deviation from the fit (Fig. 6.14).

For $N < 105$ the packing is quite efficient at distinct N -values and quite inefficient at others. Few extra d-charges are present. In some of the cases where extra d-charges are present we find in Sect. 6.4 that these are actually stable arrangements with one or more particles missing. Then, in the region where the d-charge distribution changes from mainly single d-charges -1 to mainly clusters of three d-charges ± 1 ($105 < N < 160$), the packing is not efficient. This results in a larger radius and energy for the GEM configuration. For $N > 160$ the packing becomes more symmetric again in the sense that it is then possible to have 12 equal clusters. This results in a better packing for the configurations

and thus the radius and energy become lower again. However, both radius and energy are now unable to asymptotically reach the situation for a flat surface.

That the total number of d-charges for $N < 200$ tends to 12 and 36, suggests that the system strives towards a symmetry as far as d-charges are concerned.

We now understand why the deviation from the general trend in the energy as a function of N shows a behaviour attributed to an extra phenomenon. It turns out that we have modelled all the contributions to the system energy as smooth functions of N , while we actually find that the growth of the d-charge clusters occurs largely in jumps. The fluctuations connected with the more or less discrete growth of the d-charge clusters spans whole ranges of N -values.

6.3.4 The symmetry of the configurations

In this section we discuss the symmetry properties found for the GEM configurations as listed in Table A in the column labelled 'S'.

For a sphere certain symmetry groups are relevant that are not applicable in flat 2D arrangements (see Sect. 2.6.4). The symmetry properties are determined as discussed in Sect. 2.6.4.

Results

In Table 6.4 we list for each symmetry group, except C_1 , which N -values form configurations possessing this symmetry. There are certain N -values for which a symmetry group can be expected on *a priori* grounds, for example those related to Plato's regular solids, the regular polyhedra with $N=4, 6, 8, 12$ and 20 . In our simulations we do find the expected symmetry for $N=4$, the tetrahedron, $N=6$, the octahedron, and $N=12$, the icosahedron. But for $N=8$ and 20 we do not find the cube and dodecahedron respectively. In stead we find a twisted cube with D_{4d} symmetry and for $N=20$ we find a D_{3h} symmetry. These latter configurations can be understood in terms of the number of nearest neighbour interactions for each particle. The cube and dodecahedron have square and pentagonal faces respectively, the found configurations have triangular faces.

Analysis

Even though—for increasing N —at $N=13$ for the first time a GEM configuration is found that does not have any rotational symmetry (it only has a mirror plane), and $N=23$ is the first N -value that does not have any symmetry, for $N < 100$ more than 10% of the configurations have a high symmetry, such as T , O , or I .

A large variation in symmetry as a function of N is observed; There are N -values that form high symmetry configurations while neighbouring N -values have

symmetry	N -values
C_s	13, 15, 16, 19, 26, 31, 33, 35, 43, 59, 68, 76, 79, 90, 113
C_2	28, 34, 57, 58, 70, 81, 82, 83, 89, 103, 104, 106, 115, 118, 120, 125, 133, 135, 152, 177, 185, 190, 195, 198
C_{2v}	14, 42, 46, 51, 61, 67, 77
C_3	39, 40, 49, 52, 129, 151
C_{3v}	7, 10
C_{4v}	5
C_{5v}	11, 71
D_2	36, 54, 56, 62, 86, 88, 96, 108, 142
D_3	21, 27, 30, 41, 45, 50, 63, 65, 66, 69, 75, 78, 84, 87, 101, 192
D_{3h}	3, 9, 20
D_4	80
D_{4d}	8, 18
D_{5h}	17, 72
D_6	92, 98, 110
D_{6d}	38
D_∞	2
T	60, 64, 100
T_d	4, 22
O	24, 48, 126, 134
O_h	6, 44
I_h	12, 32

Table 6.4: symmetry groups and N -values.

low or even no symmetry. An example of this is the series with $N=23$, 24 and 25 where $N=24$ has high symmetry while $N=23$ and 25 have no symmetry at all.

Except for the special case $N=2$, where the particles form a line and thus an infinite number of rotational symmetry operators and mirror planes exist, we never find a higher symmetry axis than 6-fold.

For the main axis appearing in the symmetry groups C and D, we find that 2-fold and 3-fold axis are most popular. A 4, 5 or 6-fold axis is found less often. Only a minority of the N -values can form configurations with T , O or I symmetry.

For N larger than 105 most lowest energy configurations found have no symmetry (group name C_1). This can be due to the circumstance that in these cases the GEM arrangement not yet has been determined. It is more likely, however, that due to the d-charge growth no symmetric arrangement can be formed.

Discussion

That often a 2- or 3-fold axis is found has clear consequences for the position of the main axis. The 2-fold axis mostly passes through the edges between particles and sometimes through particles, the 3-fold axis mainly passes through the centre of triangular faces (formed by three mutually neighbouring particles), but also through particles. The 4-fold rotational axis mainly passes through squares made

symmetry	value
C	1
D	2
T	3
O	4
I	5

Table 6.5: Numerical value for the symmetry classes. See text for explanation.

by 4 particles, and sometimes through particles. The 5-fold rotational axis most often is found to pass through particles that are d-charges -1 but it is possible that it goes through a pentagon formed by 5 particles (this is the case for $N=11$ and 72). The 6-fold axis is found to pass through particles with d-charge 0.

After comparing the deviation of the energy $E_g(N)$ from the model ($E_d(N)$ in Fig. 6.14) with the symmetry data in Table A, we find that there is a correlation. In Fig. 6.19 the average $E_d(N)$ against the symmetry class is plotted for $N < 100$ and for $2 \leq N \leq 200$. The symmetry classes are shown in Table 6.5. In that table C stands for all groups with name C , similarly the D groups. The symbol T stands for tetrahedral, O for octahedral and I for icosahedral symmetry (see Sect. 2.6.4). For $N > 105$ almost no configuration has global symmetry, if these are not used (the line labelled $N < 100$ in Fig. 6.19) we find a stronger correlation than for $2 \leq N \leq 200$.

The reason for not finding more than 6-fold rotational symmetry is that on a flat surface the particles strive towards this type of symmetry and the placement of the d-charges (12 in number) poses a limit on the available maximum symmetry to 6-fold.

Although Plato's solids are perfectly regular, and possess many symmetry operators, for $N=8$ and 20 they do not form a stable potential energy minimum. In these two cases Plato's configurations are not triangularly built-up [158].

Another well known structure is the truncated icosahedron or bucky ball with $N=60$. In our simulations we do not find this configuration. The used circularly symmetric potential does not favour this structure, just as with hard-particles, Coulomb and logarithmic potentials [48, 31].

For $N > 105$ extra d-charges are present. With these it becomes harder to achieve high symmetries. For example the clusters with d-charges -1, 1, -1 not only have to be positioned symmetrically, the orientation of the clusters also has to match. Between $N = 105$ and $N = 170$ we find that there are more than 12 d-charges -1 but not enough to form 12 clusters of d-charges -1, +1, -1. It not possible to form icosahedral symmetry since no multiple of 12 d-charges is present. Other symmetries may be possible but are simply not found in our studies.

Since the clusters of d-charges present for $N > 105$ have an odd number




Figure 6.19: Correlation between symmetry and deviation from the fitted function.

of particles which have an orientation it is not possible to form for example icosahedral symmetry since this needs to have a 5-fold rotational axis through the d-charges (see also Sect. 6.6).

Many of the local minima have high symmetry. Sometimes the GEM configuration has low or even no symmetry while local minima exist that do have symmetry (see also Sect. 6.2). In these cases the low symmetry arrangements have a better distribution of the curvature stress and thus lower energy.

Compared to the Coulomb and Logarithmic potential [56] it is remarkable that the optimal LJ configurations have a low symmetry over the whole N -range. In Fig. 6.20 we show as a function of N when the LJ or the Coulomb potential gives a higher symmetry for the optimum configuration. We see that the Coulomb potential on average gives more symmetrical configurations. Especially in the range of the transition from single d-charges to clusters of d-charges there is almost exclusively C_1 symmetry for the LJ potential, therefore in this range ($105 < N < 160$) the Coulomb gives higher or equal symmetry.

There are two main differences between the Coulomb and Logarithmic potential compared with the LJ potential; the LJ potential is short ranged and has an attractive tail. We first discuss if the range of the attractive part is of importance and then if the fact that there is an attractive part responsible for the low symmetry arrangements.




Figure 6.20: Comparison between the symmetry properties of equilibrium N -particle configurations with the LJ and Coulomb potentials. The label “Higher” indicates that the LJ configuration has higher symmetry than the Coulomb configuration.

It might be that long range potentials organise particles at a larger distance than short range potentials, thus causing higher symmetry. To test whether the spatial extension of the potential is important, we made a number of adaptations to the LJ potential. In these we tried to leave the hard repulsive core and the minimum in the potential essentially undisturbed while extending the tail. Results for these new potentials indicate no increase in symmetry (data not shown).

To test if the presence of the attractive part of the LJ potential is responsible for the low symmetry we perform simulation experiments with increased density, such that the most important contributions to the potential energy comes from the repulsive core in stead of the attractive tail. These simulation experiments are discussed in Sect. 6.5. The results indicate that the attractive part of the potential is responsible for the relatively low symmetry of the LJ GEM configurations.

6.3.5 Enantiometric configurations

In this section we study the “handedness” that arrangements can have. If an arrangement is enantiometric, a second and different arrangement exists with exactly the same properties such as energy and symmetry. These two arrangements can, however, not be mapped onto each other by applying rotation operators (proper isometries); they have opposite chirality or handedness related by an improper isometry (reflection-rotation).

To detect if an arrangement is enantiometric we use the algorithm described by Erber and Hockney [48] and Bergersen *et al.* [32] (see Sect. 2.6.4). We test all the configurations found for $2 \leq N \leq 200$ for this property.

Results

The first 9 enantiometric configurations are $N=23, 25, 29, 37, 47, 53, 55, 61, 73$, all odd N -values. The first even N -values for enantiometry are 74, 94, 102. For all $22 < N < 100$ there are only 17 enantiometric GEM configurations of which only 2 with even N . All our $N > 144$ GEM results are found to be enantiometric. By comparing the N -values for enantiometry with the symmetry groups listed in Table A we can check that the enantiometric configurations have no rotational or mirror symmetry.

It is not surprising that large N configurations are more often enantiometric than those at small N . At the larger N it is more difficult to form configurations that have high symmetry. Random states are in general enantiometric, when mirrored they can not be brought back to the original by rotations alone. Thus the observation that more enantiometric states are found at large N is—in our view—not necessarily an indication of a new level of complexity, as stated by Erber *et al.* [48], but rather an indication of (pseudo)randomness. The distribution of enantiometric configuration in our data is similar to that for the Coulomb potential [48]. The pseudo randomness at large N may partly reflect that our GEM results there are not yet the true global minima.

6.3.6 Dipole moment

In this section we study the “mass dipole moment” $C_m(N)$ of the GEM configurations of Table A.

The value of $C_m(N)$, the length of the vector from the centre of the sphere to the centre of mass, indicates how well balanced the particles are spread over the sphere.

The centre of mass is simply calculated from the sum of particle coordinates within a configuration

$$C_m(N) = \left| \sum_{i=1}^N \mathbf{r}_i \right|, \quad (6.13)$$

where \mathbf{r}_i is the position vector of particle i .

Results

In Fig. 6.21 we plot $C_m(N)$. We see that for some (ranges of) N -values C_m (nearly) vanishes, while others have $C_m \sim \mathcal{O}(1)$. No clear trend in $C_m(N)$ is visible, over the whole N -range we find $C_m(N)$ of $\mathcal{O}(1)$ or smaller.

The first N -value with a non-vanishing C_m -value is $N=5$. Even for the larger N -values ($N=110, 126, 134$) some configurations have a vanishingly small C_m . There are two configurations with remarkably high dipole moments. These are the configurations of $N=46$ and 191.




Figure 6.21: The size of the centre of mass for the LJ potential as a function of N .

symmetry group	average dipole moment
C	0.512
D	$2 * 10^{-6}$
T	$3 * 10^{-7}$
O	$8 * 10^{-8}$
I	$\approx 10^{-15}$

Table 6.6: The average size of the dipole moment as a function of symmetry group.

Since the average radius of the LJ solutions increases with N , C_m might also tend to increase. In Fig. 6.21, however, we see no such tendency. This indicates that the relative dipole moment, normalised by the radius of the configuration, $c_m(N) = C_m(N)/R_g(N)$ would diminish, certainly for large N .

If an arrangement at small N has a non-vanishing dipole moment at all, it is likely to have a large dipole moment. On the other hand, at large N statistical averaging results in a relatively low dipole moment.

There is a relation between symmetry and dipole moment. Highly symmetric configurations usually have a zero or very small dipole moment, low symmetry configurations are less ordered and have higher dipole moments. In Table 6.6 the average size of the dipole moment for the symmetry groups is shown. Configurations with C symmetry have, on average, a high dipole moment, arrangements with higher symmetry have significantly less dipole moments. The two configurations found with I symmetry have zero dipole moments within the precision of our calculations. Since small N are more likely to have high symmetries, they also have more configurations with vanishingly small dipole moments.

Global energy minima configurations depend on the interaction potential used.

For the logarithmic interaction it can be shown [32] that the dipole moment always vanishes. For the Coulomb potential it was found [48] that the dipole moment is smaller than 10^{-2} for $N \leq 112$. For the generalised Coulomb potential of order $n > 1$ the size of the dipole moment is $\mathcal{O}(1)$, while random configurations have dipole moments $\mathcal{O}(N^{-1/2})$. For the mentioned potentials a unit radius was used.

6.4 Global and local energy minima

In this section we report on all potential energy minima (global and local) that we obtained. These minima are too numerous to list in a table, therefore we present our data in figures. We discuss the distribution of the minima in the E - N plane and the number of minima found as a function of N .

The local minima presented in this section are obtained using the SD method alone. For each N 1000 random starts are made.

By studying all the found minima we gain insight in the state space as a function of N . If correlations in the minima as a function of N exist, these are of importance when we consider systems in equilibrium in a grand canonical ensemble where the number of particles is allowed to fluctuate.

To investigate whether correlations exist as a function of N , we have to compare configurations of neighbouring N -values. We use the following operational method: if a given configuration with $N-1$ particles is strongly correlated with a given configuration with N particles, then inserting a particle “cautiously” into the first configuration would lead—without a structural transition (Sect. 6.2)—to the configuration with N particles. The other way around: if from a configuration with N particles one particle is cautiously removed, this might produce the correlated configuration with $N-1$ particles without a large scale rearrangement.

The single particle is slowly inserted or removed by changing a parameter p , in small steps (0.01) between zero (particle completely absent) and unity (particle completely present), and applying SD after every step to relax the configuration. The energy and forces of the single particle are multiplied with p . As in Sect. 6.2, where the radius is changed in steps, we measure the Euclidean distance between successive configurations. Following this procedure it is found that, if no large Euclidean distances as a function of p are travelled, no transition takes place.

A new single particle is inserted at the position with the Largest Minimal Distance (LMD) to any of the already present particles. This position is approximated by using a discrete set of 30000 well distributed points (generalised spiral points described by Zhou [31]) over the sphere and calculating the LMD for each of these points. For the removal of a particle we simply take one of the N particles. In the absence of symmetry we have to repeat this for all N particles since they might all lead to a different configuration at $N-1$. If symmetry is present less particles need to be taken out.

The results in this section suggest another simulation where the radius *and* the number of particles can change. The simulation is started with a random configuration for a random number of particles (between 10 and 50) on a sphere with random radius (between 1 and 10). First we use SD without changing the radius. This often causes a cap configuration (see Sect. 6.2) to form. We use SD including optimisation of the radius to find a (locally) optimum configuration. Then we start the following loop. First, and independent of each other, a particle is cautiously inserted in the largest hole in the configuration and the worst particle (in terms of energy) is cautiously removed from the configuration. We then have three configurations with N_1-1 , N_1 and N_1+1 . The best of these configurations (in terms of energy) is used further to insert and remove a particle, and so on, giving at step i : N_i-1 , N_i and N_i+1 . Whenever the configuration with N_i particles turns out the best, the process stops. If this is repeated for many random instances we can find which configurations have a large capture basin in the variables N and R .

The number of minima $M(N)$ and their separation in energy is related to the complexity of the problem of locating the GEM arrangements. If many local minima are present and the difference in energy between them is small, finding of the GEM arrangement is difficult.

The distribution $M(N)$ is approximated by counting for all N the number of different minima found. We give data only for N up to 120. For larger N we are not confident that all or most minima are found. Enantiometrically related configurations (see Sect. 2.6.4 and 6.3.5) are counted only once. Also we restrict the range of the initial radius to be within the range described at the beginning of Sect. 6.3 because we do not want to include the many local minima with largely open configurations for $R \gg R_{\text{eq}}$ (see Sect. 6.2).

Results

In Fig. 6.22 we plot the energy minima obtained with the SD method. The lowest energy points are the presumed GEM configurations as a function of N . If the same minimum is found more than once, still only one point is plotted. In Fig. 6.22 small N -values have only relatively few minima. High N -values have so many minima that they are not visible separately on the energy scale used. We note that, for $N > 120$, our data generated with Simulated Annealing results in lower minima (see Fig. 6.13).

The frequency of occurrence shows that highly symmetric GEM configurations are usually the most probable. For low symmetry arrangements it is observed that local minima can be more probable. For $N > 120$ we observe that the probability of minima is a smooth distribution with a peak just above the lowest energy found. In the probability distribution of the minima as a function of energy we find gaps for several N -values. In Fig. 6.22 a correlation in gaps for different N is clearly visible. Starting from a lowest energy point at some N , to-




Figure 6.22: The energy minima as a function of N , found from 1000 random starts for each N with the SD method.

wards smaller N , a band of local minima appears. The gaps are situated between the bands and other minima. The bands in the E - N plane start in relatively deep GEM configurations (see Sect. 6.3.1), examples are $N=44$, 48, 72 and 78.

Analysis

The bands in Fig. 6.22 form an intriguing phenomenon. An obvious hypothesis is that the configurations form essentially one arrangement in the variables N and R , i.e. they share a common structure or “building plan”. From Fig. 6.22 it is clear that in particular the deeper GEM have bands towards smaller N . In order to test whether the configurations in a band really have a common arrangement, we perform a number of simulation studies where a single particle is cautiously inserted in—or removed from—the configurations, as described above.

In the simulations where one particle is removed at a time from a N -particle configuration, not allowing a structural transition to take place, many local minima—and some global minima—at smaller N are found. For successive removal of a single particle we find that more than one path can be taken for some N , depending on the exact location in state space. The SD method allows different routes for different configurations corresponding to the same E - N minimum.

Two examples of this process are shown in Figs. 6.23 and 6.24 where the Euclidean distance during the “cautious” insertion and removal is plotted. The first figure shows a test case of inserting a particle in an arrangement for $N=25$ and then removing the same particle again. We observe transitions and conclude that the process is essentially reversible. The example in Fig. 6.24 shows a situation with $N=43$ and 44. There no transition at all takes place and here again the process is essentially reversible. The large peak in the Euclidean distance near $p = 0$ is no structural transition. It turns out that the repulsive core still has a relatively large width at $p = 0.01$. At $p = 0$ the particles in the arrangement shift (slightly) towards the gap.

Starting with the GEM configuration at $N=48$, we remove—in all possible ways—one particle. For the arrangements with $N=47$ thus generated, not requiring large scale rearrangement, we repeat a similar simulation. We continue this process towards $N=46$, 45 and 44. In Fig. 6.25 we plot the series of minima generated in the above described simulations. A band of minima in the same range as in Fig. 6.22 is found. For clarity we also plot the energy of the GEM configurations for $44 \leq N \leq 48$. The GEM configurations for $N=46$ and 47 are also found in our particle removal simulations.

The N -values that “attract” local minima from smaller N have a structure that is stable even when only a part of it is formed. Just as in the case where the curvature is changed we can expect that there is only a limited range in N where such a structure is (meta)stable. One observation is that inserting a particle into a configuration with a completed building plan causes a structural transition. This is due to the fact that those configurations are usually tightly packed. Then




Figure 6.23: The process of “cautiously” inserting and removing a single particle for $N=25$ and 26, respectively. Structural transitions take place.




Figure 6.24: The same as Fig. 6.23 for $N=44$ and 43. Here no structural transition takes place.

fitting a new particle is difficult and a new structure to accommodate the new particle is formed, hence the asymmetry of the bands. On the small N side there may also be a limit to the stability. If a large fraction of the optimum arrangement is removed from the system, the edge becomes relatively large and the rest of the particles is likely to collapse into a configuration with smaller radius.

We now discuss results for simulations where it is tried to insert and remove particles to lower the energy (per particle). The distribution of N -values where the process stops (where insertion and removal of a particle fails to lower the energy) is plotted in Fig. 6.26. Quite remarkably, certain N -values appear to be strongly preferred. During these simulations it turns out that on average more particles are inserted ($\approx 85\%$) than removed. Especially often found N -values




Figure 6.25: Minima generated without transitions by successive removal of particles from the global minimum at $N=48$. The energies of the global minima for $N=44$ to 48 are also plotted.

are 12, 17, 27, 32, 36, 44 and 48.

In Fig. 6.27 we plot, with logarithmic y-axis, the number of different minima found as a function of N . The number of different minima clearly rises approximately exponentially. We notice that three regions in N are visible. For $N < 10$ only one single minimum is found. For $11 < N < 20$ the number of minima rises fast. Then, for $N > 20$, on average the number of minima rises approximately exponential over the remaining N -range shown. Some N -values have a remark-




Figure 6.26: The number of times the process ended in a minimum with N particles.

ably small number of different minima. Examples of these are $N=24$ and 48. An exponential function of N is fitted through the data for $N > 70$ where the number of different minima seems to have settled into a trend. The fitted exponential, shown in Fig. 6.27, has the form $2.14 \exp(0.0367N)$.

Discussion

The data indicate that for small N -values in general the true GEM arrangement is found very often. For larger N the average energy of the found minima shifts away from the energy of the GEM arrangement. Then the distribution of minima resulting from a large number of randomly initiated runs as a function of energy becomes a broad peak with tails on both sides. The shapes of these distributions indicate that the state space for large problem sizes consists of many, nearly degenerate, energy minima. The number of minima as a function of N becomes very large for N near 200. The landscape over the state space is covered with local minima which makes finding the GEM arrangement very hard.

The relatively deep GEM configurations turn out to be correlated with local minima at smaller N . At smaller N the configurations of the minima already “point” towards the arrangement of the deep GEM at larger N . This indicates that—if the number of particles is not kept constant during a simulation—the




Figure 6.27: The number of different minima $M(N)$ as a function of N .

bands in the E - N plane all drain towards the relatively deep GEM configurations. This tends towards a relatively large population of GEM configurations at specific N -values. In other words, in R - N state space some specific N -values have a much larger capture basin than nearby values.

We also studied the data found by Erber and Hockney [48] for the Coulomb potential but did not find a similar behaviour as described above. For the Coulomb case the radius is fixed and thus open (cap) configurations with $R > R_{\text{eq}}$ are not found. If part of a structure is missing the repulsive force makes sure that the other particles spread over the entire surface of the sphere.

The GEM configurations that induce local minima at smaller N are relatively deep in terms of energy and have high density as well as high symmetry. The symmetric basic arrangements according to a “building plan” have apparently a stability spread over the whole surface of the sphere, and thus are also stable if parts are missing. In fact the minimum of $N=25$ shown in Fig. 6.6 is part of the icosahedral arrangement of $N=32$.

We find that the number of minima rises exponentially with N . Erber and Hockney [48] have shown that the number of distinct minima for the Coulomb problem also rises exponentially: $0.382 \exp(0.0497N)$. Different from Erber and Hockney, we count chiral states as one state, and we only count minima within a limited range for the radius (near that of the radius expected for the GEM

configurations). Despite this, we still find many more minima than observed for the Coulomb problem.

The attractive part of the LJ provides stability to whole classes of configurations that could not exist in the case of purely repulsive forces (see Sect. 6.2). If we also include these in $M(N)$, we have many more minima than for the Coulomb potential.

6.5 Arrangements under compression

In Sect. 6.2 we studied the (meta)stable configurations at $R > R_{\text{eq}}$. In the previous sections we explored the global energy minima (GEM) at zero pressure. In this section we simulate systems at a higher density than at equilibrium when the radius, thus the 2D “volume”, is a free parameter. In Sect. 6.3.4 it was concluded that the low symmetry of the GEM arrangements for the LJ potential (compared with the Coulomb potential) is not a result of the short range of the attractive potential. Here we test if the presence of an attractive part causes the lower symmetry. We test this by determining the symmetry groups of systems as a function of density.

When the 2D volume of the LJ system is reduced with respect to the free equilibrium value, the system becomes compressed; a net pressure results, directed outward of the spherical surface. Under these conditions the repulsive core in the LJ potential becomes more important than the attractive tail part. On the basis of the Coulomb results [48] it can be expected that the symmetry of a high density LJ system is higher than at zero external pressure. It can also be expected that—apart from the 12 inherent d-charges—d-charges $\neq 0$ pairs will be suppressed with respect to zero pressure.

The results shown in this section are obtained with the SA followed by SD method (see Sect. 3.2). In these simulations we use the NVT ensemble. The density ρ is varied from 1.1 to 1.9 in steps of 0.1. At higher density we use a higher start temperature, the other parameters are taken equal to those described in Sect. 6.3. For each combination of N and ρ we perform 25 runs. Only the lowest energy found for each N are discussed.

Results

In Table 6.7 we show the results for $N=10$ to 49. The increase in density causes a change in symmetry properties. For some N we performed simulations up to $\rho = 5$. and found that the same symmetry is present as at $\rho = 1.9$ in the table below (data not shown).

For several N -values the symmetry changes. Mostly it is an increase in symmetry. For $N=21$ and 36 we observe that the symmetry can also drop. In both cases the 2-fold symmetry axes belonging to the D symmetry group are lost, at

$N=21$, however, mirror planes are formed. For $N=33$ we observe that within the shown density range the symmetry can drop and then rise again. The symmetry rises spectacularly for $N=14,16,28,29,40$ and 42 . If at zero pressure already a high symmetry is present the symmetry properties do not change with increase of density. If at zero pressure low symmetry is present it usually increases with density. Compared to the zero pressure GEM arrangements we indeed find that the number of d-charges different from -1 and 0 decreases with increasing density. The only N value that has d-charges $+1$ at $\rho = 1.9$ is 47 while at zero pressure $N=23, 29, 31, 33, 40, 42, 43, 46$ and 47 have d-charges $+1$ (see Table A).

For $N=200$ at zero pressure, 12 clusters of disclinations are present (see Fig. 6.18). When in (independent) simulations the density is increased in steps, some of these clusters disappear to leave single d-charges (see Fig. 6.28).




Figure 6.28: The Voronoi polygons of the configuration with $N=200$ at $\rho=2.60$.

Analysis

As was already shown in Sect. 6.2, the minima have a stability range in the variable R . Here we find that at high density this is also the case. For many

N	$P = 0(\rho)$	$\rho = 1.1$	1.2	1.3	1.4	1.5	1.6	1.7	1.8	1.9
10	C_{3v} (0.987308)	C_{3v}	C_{3v}	C_{3v}	C_{3v}	C_{3v}	C_{3v}	C_{3v}	C_{3v}	C_{3v}
11	C_{5v} (0.988761)	C_{5v}	C_{5v}	C_{5v}	C_{5v}	C_{5v}	C_{5v}	C_{5v}	C_{5v}	C_{5v}
12	I_h (1.07529)	I_h	I_h	I_h	I_h	I_h	I_h	I_h	I_h	I_h
13	C_s (0.98722)	C_{3v}	C_{3v}	C_{3v}	C_{3v}	C_{3v}	C_{3v}	C_{3v}	C_{3v}	C_{3v}
14	C_{2v} (1.0037)	D_2	D_2	D_2	D_2	D_2	D_2	D_{6d}	D_{6d}	D_{6d}
15	C_s (1.0243)	D_3	D_3	D_3	D_3	D_3	D_3	D_3	D_3	D_3
16	C_s (1.03086)	C_s	D_{4d}	D_{4d}	D_{4d}	D_{4d}	D_{4d}	D_{4d}	D_{4d}	D_{4d}
17	D_{5h} (1.05067)	D_{5h}	D_{5h}	D_{5h}	D_{5h}	D_{5h}	D_{5h}	D_{5h}	D_{5h}	D_{5h}
18	D_{4d} (1.04919)	D_{4d}	D_{4d}	D_{4d}	D_{4d}	D_{4d}	D_{4d}	D_{4d}	D_{4d}	D_{4d}
19	C_s (1.02793)	C_s	C_2	C_2	C_2	C_2	C_2	C_2	C_2	C_2
20	D_{3h} (1.06737)	D_{3h}	D_{3h}	D_{3h}	D_{3h}	D_{3h}	D_{3h}	D_{3h}	D_{3h}	D_{3h}
21	D_3 (1.05791)	D_3	C_{2v}	C_{2v}	C_{2v}	C_{2v}	C_{2v}	C_{2v}	C_{2v}	C_{2v}
22	T_d (1.07032)	T_d	T_d	T_d	T_d	T_d	T_d	T_d	T_d	T_d
23	C_1 (1.04467)	C_1	C_1	C_1	C_1	C_1	C_1	C_2	C_2	C_2
24	O (1.08631)	O	O	O	O	O	O	O	O	O
25	C_1 (1.05963)	C_1	C_2	C_3	C_3	C_3	C_3	C_3	C_3	C_3
26	C_s (1.06587)	C_s	C_s	C_s	C_s	C_s	C_s	C_2	C_2	C_2
27	D_3 (1.08808)	D_3	D_3	D_3	D_3	D_3	D_3	D_3	D_3	D_3
28	C_2 (1.07299)	C_{3v}	T	T	T	T	T	T	T	T
29	C_1 (1.06693)	C_1	C_1	C_s	C_2	D_{3h}	D_{3h}	D_{3h}	D_{3h}	D_{3h}
30	D_3 (1.08558)	D_3	D_3	D_3	D_3	D_3	D_3	D_3	D_3	D_3
31	C_s (1.08444)	C_s	C_s	C_s	C_s	C_{2v}	C_{2v}	C_{2v}	C_{3v}	C_{3v}
32	I_h (1.10538)	I_h	I_h	I_h	I_h	I_h	I_h	I_h	I_h	I_h
33	C_s (1.07154)	C_s	C_s	C_s	C_s	C_s	C_s	C_1	C_1	C_2
34	C_2 (1.08175)	C_2	C_2	C_2	D_2	D_2	D_2	D_2	D_2	D_2
35	C_s (1.08314)	C_s	C_s	C_s	C_s	C_s	C_s	C_s	C_s	C_s
36	D_2 (1.08908)	D_2	D_2	D_2	C_2	C_2	C_2	C_2	C_2	C_2
37	C_1 (1.08717)	C_2	C_2	C_2	C_2	C_2	C_2	C_2	C_2	C_2
38	D_{6d} (1.09759)	D_{6d}	D_{6d}	D_{6d}	D_{6d}	D_{6d}	D_{6d}	D_{6d}	D_{6d}	D_{6d}
39	C_3 (1.08771)	C_3	C_3	C_3	D_{3h}	D_{3h}	D_{3h}	D_{3h}	D_{3h}	D_{3h}
40	C_3 (1.09083)	C_3	C_3	C_3	T_d	T_d	T_d	T_d	T_d	T_d
41	D_3 (1.09097)	D_3	D_3	D_{3h}	D_{3h}	D_{3h}	D_{3h}	D_{3h}	D_{3h}	D_{3h}
42	C_{2v} (1.08797)	C_{2v}	D_{5h}	D_{5h}	D_{5h}	D_{5h}	D_{5h}	D_{5h}	D_{5h}	D_{5h}
43	C_s (1.09764)	C_s	C_{2v}	C_{3v}	C_{3v}	C_{3v}	C_{3v}	C_{3v}	C_{3v}	C_{3v}
44	O_h (1.1131)	O_h	O_h	O_h	O_h	O_h	O_h	O_h	O_h	O_h
45	D_3 (1.10392)	D_3	D_3	D_3	D_3	D_3	D_3	D_3	D_3	D_3
46	C_s (1.0808)	C_{2v}	C_{2v}	C_{2v}	C_{2v}	C_{2v}	C_{2v}	C_{2v}	C_{2v}	C_{2v}
47	C_1 (1.098)	C_1	C_1	C_1	C_1	C_1	C_s	C_s	C_s	C_s
48	O (1.11302)	O	O	O	O	O	O	O	O	O
49	C_3 (1.09747)	C_3	C_3	C_3	C_3	C_3	C_3	C_3	C_3	C_3

Table 6.7: The symmetry groups for various N as a function of density. The first column lists the symmetry group at zero pressure, the density of those configurations is indicated between brackets.

condition	N -values
Sym(LJ)>Sym(C)	11, 13, 25, 29, 30, 33, 43
Sym(C)>Sym(LJ)	10, 16, 19, 23, 27, 35, 36, 37, 46

Table 6.8: The N -values where either LJ or Coulomb gives higher symmetry at $\rho = 1.9$ in the range $10 \leq N < 50$.

N -values when the density is increased, configurations are formed with higher symmetry than at zero pressure.

For $N=10$ the contribution to the total energy from the attractive part of the potential energy function becomes smaller than that of the repulsive part around the density 1.14, for $N=30$ this is $\rho = 1.26$. It is possible that this is the reason that around these densities most changes in symmetry occur. As soon as the contribution of the attractive part of the potential has become negligible we can approximate the interaction to be a generalised Coulomb potential of order $n = 12$. From then on an increase in density is not important anymore since in the calculation of the energy the radius contribution can be separated from the configuration contribution. In Sect. 6.2 it was already observed that for $R < R_{\text{eq}}$ few transitions occur.

Compared to the data for the Coulomb potential [31] we now find, at $\rho = 1.9$, that the symmetry of our configurations is on average the same. In Table 6.8 we show which N -values have higher or lower symmetry than Coulomb. The N -values that give equal symmetry are not listed. If we compare the LJ configurations at $\rho = 1.9$ with those of the Coulomb potential we find that in 10 cases the Coulomb interaction gives higher symmetry, in 7 cases LJ gives higher symmetry and in 24 cases they have an equal symmetry group. At zero pressure these figures are 23, 3 and 14 respectively.

At high density it is costly to form d-charges since these result in a less efficient packing. Therefore we see that the d-charge distribution in almost all cases tend to have 12 d-charges -1 in association with only d-charges 0.

In the simulations with $N=200$ it is observed that the clusters of d-charges disappear to leave single d-charges -1. Although the symmetry does not rise (it remains to be C_1) it becomes easier for high symmetry to emerge; the degrees of freedom associated with the orientation of the clusters of d-charges are now removed.

We thus find that the presence of the attractive term in the LJ potential is the reason for the low symmetry arrangements compared to those of the purely repulsive Coulomb and logarithmic potentials.

6.6 Icosahedral arrangements

In this section we discuss simulations of the LJ-system on a spherical surface with enforced symmetry. Based on the observations in Sect. 6.3.4 we know that arrangements with icosahedral symmetry often have a relatively low energy compared to other symmetries (see Fig. 6.19). Therefore we choose to fix icosahedral symmetry in the simulations discussed below.

The simulations with large N take so many CPU-cycles and have so many minima that it is not possible in practice to leave all degrees of freedom intact. In order to expand our search for low energy arrangements to $N > 200$ we can use the requirement of symmetry to simulate large systems since the number of degrees of freedom are then diminished greatly.

The necessary d-charges for spherical configurations are in these simulations also distributed with icosahedral symmetry. If the final arrangement with the enforced symmetry is not stable, the arrangement resulting if the symmetry requirement is lifted, is expected to have (clusters of) d-charges distributed icosahedrally. In Sect. 6.3.3 we find that for the global minima the average distribution of d-charges is also nearly icosahedrally. So even if the found icosahedrally symmetric arrangement is not stable we can find a reasonably good minimum.

For the simulations presented here we use a special implementation of SA and SD in addition to the standard SD method. The special implementations are discussed in Sect. 2.8.2. They are used to minimise the energy for a system with built-in icosahedral symmetry. The standard SD method is used afterwards to check whether the icosahedral arrangement obtained is stable when the symmetry requirement is lifted.

As with the standard SA we have to use a NPT ensemble with zero pressure. The parameters used in the simulations are appropriate for simulations with much smaller N since the number of degrees of freedom is reduced. For the number of Metropolis steps per particle, where one particle now stands for a number of particles connected via the icosahedral symmetry, we use the number $12A + 30B + 20C + 60D + 60E$. The different types of particles (A to E) are discussed in Sect. 2.8.2. For the cool rate we use 0.95.

The parameters that determine what kind of icosahedral structure is generated are A, B, C, D and E . The values we use for these parameters are chosen as follows. For the parameters A, B and C we use all combinations of the values 0 and 1. The value of D is varied from 0 to 4. The value of E is varied from 0 to 9 but instead of a full range it is always taken such that there is a reasonable balance between the number of particles on the edges and the number of particles in the faces (Eq. 2.17 in Sect. 2.8.2).

For each combination of parameters A, B, C, D and E we perform 50 independent minimisations. We only show results of the lowest energies observed for each of the combinations.

Results

In Table B in the appendix the results are shown of the minimisation simulations. In the first column we give the total number of particles N present for the combination of A, \dots, E given in the next 5 columns. Then the energy and radius of the arrangement with enforced icosahedral symmetry are shown and in the last two columns we list the energy and radius of the arrangement that is obtained when applying the standard SD method.

Analysis

As before, with the standard SA and SD we can not be sure that the true global minimum arrangement with icosahedral symmetry has been found. Even if we do find the global minimum it is not guaranteed that this is the global minimum if the symmetry requirement is lifted.

In the table we see that many simulations yield configurations with icosahedral symmetry that are stable even if the symmetry requirement is lifted. If we plot the difference in energy between the configuration with and without the symmetry requirement on a logarithmic scale we get Fig. 6.29. Some of the found minima have near zero difference and do not show up in the figure. We see that




Figure 6.29: Difference in energy between configurations with enforced and lifted symmetry.

for the smaller N we find configurations that are truly stable. For $N > 600$ the generated configurations are not stable if the symmetry requirement is lifted. Then we often see a drop in energy of $\mathcal{O}(1)$. There also seems to be a number

of situations where the difference in energy is around 10^{-10} . In these cases the configuration hardly changes and can thus be considered to be remained identical.

Discussion

The number of particles that we were able to use for these simulations is much larger than for the arrangements without symmetry requirements.

Often high symmetry arrangements can be formed but they are not the most optimal arrangements. If we consider the difference between the fit on the energy as a function of N in Sect. 6.3.1, see Fig. 6.30, then we see that almost all configurations have a higher energy.




Figure 6.30: Difference in energy between the model for the energy as a function of N and the configurations from this section where the symmetry requirement has been lifted. Stable arrangements are indicated.

In Sect. 6.3.3 we show that clusters of d-charges form for $N > 105$. The number of d-charges per cluster is odd and consists of one more d-charge -1 than d-charges +1. Then it is difficult to find configurations with several 5-fold rotational axes. It is, for example, possible to form those symmetry axes if the clusters consist of 11 d-charges. In that case one d-charge -1 is positioned in the centre of the cluster and 5 dislocations are placed symmetrically around it. According to our simple model of the growth of the clusters as a function of N (Eq. 6.12) this would happen around $N \approx 2200$.

Chapter 7

Discussion and conclusions on the physics results

In this chapter the results presented in Chapters 5 and 6 are discussed in a common frame.

The thermodynamic studies on the $N=100$ system reported in Chapt. 5, do not provide a T_c value for the transition from liquid to solid. In the $E(T)$ and $P(T)$ curves no indication of the transition is visible. The derivative of $E(T)$, the specific heat C_v , however, does show a peak around $T = 0.8$ and a change in behaviour around $T = 0.35$. From the data on the structure factor S_k we deduce that, according to theoretical predictions and data from Udink, a transition should occur in the range $0.1 < T < 0.4$. The radial distribution function $g(r)$ shows a smooth decrease in peak height for $T > 0.3$.

Thus, for the $N=100$ system, we can conclude that between $T \approx 0.8$ and $T \approx 0.3$ freezing takes place. The changes in system behaviour are more abrupt than expected from a broadening of the freezing temperature due to the small system size. It seems that freezing starts with a small fraction of the system solidifying at $T=0.8$, followed by gradual growth of the crystalline part until it extends over the whole system. This indicates that possibly a hierarchical ordering of the system is present where the d-charges and their surroundings behave differently from the rest of the system.

The effective interactions between d-charges show unlike-sign d-charges to attract and like-sign d-charges to repel each other. The attraction leads to clustering, the repulsion to a tendency for icosahedral symmetry of the d-charge distribution in frozen arrangements. Between $T \approx 0.8$ and $T \approx 0.2$ the number of d-charge clusters has a plateau while the density of d-charges changes strongly. This indicates that in this region the large d-charges clusters disappear. First experiments on the mobility (or at low temperatures volatility) indicate that in the freezing temperature range the d-charges start to be less and less volatile going outward to the first and second nearest neighbour rings around a d-charge (data not shown). This could indicate that the d-charges and their immediate sur-

roundings are frozen earlier than the bulk of the system. Further investigations, however, are needed for verification.

The energy and pressure for small N systems is higher than for a large flat system. We find that the packing efficiency, and the related ordering, gives the main contribution to this difference. The influence of optimum packing order, which is dependent on N (see Chapt. 6) already starts to play an important role during early freezing.

Besides the system size N , the symmetry also plays a role in the location of the transition temperature. For $N=32$ (icosahedral symmetry) it is observed that at temperatures around the transition, both orientational and positional order are longer ranged than for $N=33$ (C_s symmetry). For $N=48$ (octahedral symmetry) compared to $N=33$ this is also observed.

Although the spherical $N=4000$ system behaves more like a large flat system than for the smaller spherical systems used in Chapt. 5, we note that during freezing the energy does not drop as fast as for Udink's flat system. In Chapt. 6 we report evidence that this difference is probably due to the (N -dependent) relatively large number of d-charges present in large spherical LJ systems, even at zero temperature.

In Chapt. 6 the energy landscape over the state space at zero temperature is explored as a function of N and R . For large R compared to R_{eq} , the "open" arrangements where the particles cover only part of the sphere, we find that many different metastable arrangements exist, each with a range of stability in curvature. If the curvature is changed to the point where an arrangement is no longer metastable, a structural transition occurs. Then the particles rearrange to another arrangement with a different 2D topological (Voronoi) distribution. During the transition d-charges can be incorporated and local stress redistributed. In general, as the radius changes, an energy barrier is present between the two arrangements in their deepest energy state.

If the radius is around R_{eq} the system closes over the sphere. The number of different minima with $R \approx R_{\text{eq}}$ rises exponentially with N . We first discuss the properties of the Global Energy Minima (GEM) and then global *and* local minima together.

The general trend in the energy of the GEM arrangements, $E_g(N)$, is modelled by a smooth function, fitted to the data while ignoring packing effects. We then find that certain N -values have configurations with significantly lower energy than indicated by the general trend. These configurations on average have relatively high density, high symmetry, a small dipole moment, are not enantiometric and have a large capture basin.

The general trend in $R(N)$ indicates that small N systems have low densities. Towards larger N the density levels off at a density higher than that of an optimum flat regular lattice. The symmetry of the GEM arrangements as a function of N indicates that the LJ potential forms less symmetric configurations than is the case for the Coulomb and logarithmic potentials. The dipole moment is of

$O(1)$ which can be expected for a system where the ordering comes from a 12^{th} power repulsive core; no trend in N is apparent. Many of the GEM arrangements are enantiometric, which is a property easily obtained in random configurations.

The d-charges have an interesting and unexpected behaviour as a function of N . We find that as N rises, so does the number of present d-charges. These, however, do not rise linearly but appear to grow in steps. It is expected that this effect is also present for much larger N . For $N < 105$, 12 and sometimes a few more, d-charges are found. Visual inspection shows that the d-charges are distributed far from each other, ideally forming an icosahedron. Then, for N up to $N=160$, The number of d-charges rises to 36, adding 12 dislocations. For larger N , in our data to 200, we find that 12 clusters of 3 d-charges each and with net d-charge -1 are present. The steep growth of the number of d-charges can be correlated with the data on the deviation from the general trend in $E_g(N)$ and the deviation from a general trend in $R(N)$. In a simple model treating the “lattice” as a medium and only taking account of d-charges, we find that transitions to larger clusters of d-charges occur at larger N . There is also a correlation with symmetry. Below $N=105$ highly symmetric arrangements are possible, above $N=160$ a symmetry in the number of d-charges per cluster and in the positions of the clusters is found, not in the configurations themselves, indicating a hierarchical build-up of the arrangements. In the region where the number of d-charges rises, hardly any symmetry can be found.

From the data on all the minima found (local *and* global) we find a new phenomenon not observed in the purely repulsive Coulomb and logarithmic potentials. Some N -values, especially those found to have GEM arrangements that are significantly better than predicted from the general trend, induce local minima at smaller N . These N -values have arrangements that are very stable. Even if part of these arrangements is missing they can still form global and local minima at smaller N . These form bands in the E - N plane. There is a limit to the range of stability in N for these arrangements. Since a range of stability is found in R and for some arrangements also in N , we conclude that these arrangements are related by a common “building plan”. Adding particles to partly finished arrangements often lowers the mean energy while completing the building plan. Hence, a range of N -values can “drain” towards a preferred completed building plan.

The complexity of the search for the GEM arrangements rises exponentially with N since the number of local minima rises exponentially with N while their energies are nearly degenerate. This indicates that without special implementations the number of simulations necessary becomes prohibitive for large N . In order to find solutions for larger N -values, degrees of freedom can be removed. The two arrangements we find with icosahedral symmetry have significantly lower energies than indicated by the general trend. From the data on the growth of d-charges, however, we can deduce that this type of symmetry—in the exact sense—is not expected to show up very much for larger N . However, building

this type of symmetry into the simulations can be used to generate start arrangements for further searches through state space.

If the radius is smaller than R_{eq} the system is compressed. The energy is high due to the repulsive hard core. The symmetry of the deepest energy arrangements in general increases with density. The number of present d-charges tends to the minimum number necessary (i.e. 12).

Under appropriate conditions the above described mechanisms (the energetically favoured GEM arrangements for specific N -values, the draining of open arrangements towards these GEM arrangements and the steep growth of d-charge clusters) lead under appropriate conditions to a non-uniform size distribution. The energy preferences for specific N -values possessing preferred building plans, remain significant at least up to the temperature where freezing sets in.

Due to the use of an interaction potential with an attractive term (here LJ), allowing stable zero pressure arrangements with radii dependent on N , as well as on R (open arrangements), adds substantially to the understanding of natural spherical systems. Energetically favoured GEM arrangements that might lead to a non-uniform size distribution are also found for the Coulomb and logarithmic potentials, where often the same preferred N -values are found. For these arrangements high symmetry is the most striking property. For the LJ potential we find two mechanisms (draining of open arrangements towards energetically favoured N -values and steep growth of d-charge clusters) that can lead to a non-uniform size distribution, different from what was known experimentally for purely repulsive potentials.

Compared to the Coulomb and logarithmic potential we find that the LJ GEM arrangements on average have low symmetry. In the LJ system the inherent d-charges can, for $N > 100$, distort the lattice creating more d-charges, where highly symmetric arrangements can not be formed anymore. Increasing the energy present in the lattice of particles without d-charge by enforcing high density, causes the d-charge clusters to disappear, which allows for high symmetry again.

The discrete size distribution for the spherical natural systems discussed in the introduction of this thesis, are observed at finite (room) temperature. For our model system we find that the energy difference between arrangements remains significant upto the beginning of the freezing transition.

The natural systems are formed in a self-assembly process. Although we do not precisely mimic such a process we find that in our model system preferred building-plans exist that correspond to minimum energy arrangements in the variables N and R . With this information we infer that in a self-assembly process taking place in the combined state space of all possible N and R systems, the arrangement can follow the building plan energy minimum towards larger N , finally arriving at the energetically most preferred N and R .

In conclusion, in our computer experiments we have obtained preliminary evidence for different mechanisms that individually or together can lead to pre-

ferred sizes of spherical 2D LJ systems. The amount of computing power needed to study in a similar way realistic N values and interactions for biomembrane vesicles and viruses is enormous and beyond current computing capabilities.

Summary

Seeking an explanation for the discrete size distribution of some naturally occurring spherical shells formed by self-assembly—such as biomembrane vesicles and viruses—provides a first physics motivation for the research presented in this thesis. Attempting to capture in a simple model aspects of these natural systems that are relevant for their intrinsic stability, we consider N particles with Lennard-Jones (LJ) interaction confined to the surface of a sphere with radius R . The formation of the energetically optimal arrangements is mimicked by crystallisation on a sphere.

Many crystallisation studies have been performed for flat 2D systems. However the thermodynamics of similar spherical systems can be different due to the presence of disclinations (“d-charges”), needed for curvature, and distorting the triangular lattice that is optimal for flat systems. The inherent d-charges can act as seeds for a hierarchical order. From this point of view, understanding the thermodynamic behaviour and properties of spherical LJ systems—and making a comparison with flat systems—provides a second physics motivation for the work presented in this thesis.

After the introductory chapter where the research motivations, the concepts involved, and some historical notes are given, the relevant theoretical background and a concise literature overview is presented in Chapt. 2.

The present work amounts to experimentation by computer simulation. Simulations of the adopted model are performed and properties of the model system are measured and analysed. In the crystallisation studies we mainly use Simulated Annealing (SA) with a Monte Carlo simulation at each temperature (T), and a stochastically initiated Steepest Descent (SD) energy minimisation method, both at fixed N -values. These are discussed in detail in Chapt. 3. From methodological studies we conclude that SA in comparison with stochastically applied SD is more efficient in finding the most stable solutions: the global energy minima (GEM). We observe an optimum in the quality versus calculation time, dependent on details of the energy landscape over the system state space (varying with N).

Since the SA method is computationally highly intensive and hundred thousands independent simulations are needed, applying some form of parallelism is mandatory. In Chapt. 4 we discuss four different parallelisation strategies. The

vectorisation of the simulation code is relatively easy and on a vectorsupercomputer high speedup can be attained. For the two decomposition methods, systolic and data parallel, it turns out that the first is not capable of scaling—with constant quality—to a large number of processors. The latter method turns out to be very efficient if the number of processors is chosen in accordance with the N value. The combination of the two decomposition methods, the hybrid implementation, is efficiently scalable up to a large number of processors. From a time complexity analysis we conclude that on modern parallel—possibly heterogeneous—computer systems especially the data decomposition approach should be adapted. Because of our access to large amounts of networked computational resources (Condor) we mainly use job parallelism, with decomposition of the parameter space.

Temperature dependent simulations of the crystallisation process, are reported in Chapt. 5. From the behaviour of the spherical LJ-system at constant density (constant R) as a function of T , we deduce that the freezing transition for small system sizes occurs over a broad T range below the transition for a large flat system. For $N=100$, where the T dependent studies are concentrated, we find that freezing takes place in the T range between $T=0.8$ and $T=0.2$. The radial distribution functions for the various d-charges indicate that like-sign d-charges effectively repel each other while unlike-sign d-charges attract each other. The attraction leads to clustering and the repulsion to a tendency for icosahedral symmetry of the distribution of d-charge clusters with net d-charge -1. During freezing already systematic thermodynamic differences show-up as a function of N . These appear to be correlated with the symmetry and packing efficiency, consequently with long range positional and bond-orientational order, for the zero temperature GEM states.

Resulting crystalline arrangements at zero temperature and zero external pressure are presented in Chapt. 6. There we find that minima in the energy landscape over the combined state spaces for different N are present, with a range of metastability in both curvature radius R and N . At low curvature (large R) many different minima are observed corresponding to “open” arrangements, not completely covering the sphere. Where arrangements are “closed”, in general we find local minima very near the global minimum. The number of local minima for closed arrangements rises exponentially with N . The GEM results clearly show energetic preference for specific N -values compared to a smooth general trend in energy, fitted without packing constraints taken into account. For larger system sizes preference for specific ranges in N is observed. The preferred arrangements have on average a relatively high density and symmetry. The distortion in the triangular lattice sites surrounding the inherent d-charges tends to increase with system size (“d-charge polarisation”). We observe that the growth of d-charge cluster size in the studied N range occurs—because of the tendency to symmetry—steeply from 12 to 36. In the transition range for N the arrangements have relatively low symmetry, low density and they are less optimal than the general trend in the energy as a function of N .

In an extensive SD study of global and local minima, energetically preferred “building plans” become manifest. These look like metastable closed arrangements, usually with a high symmetry, where a number of particles can be missing (making them open), while still retaining metastability. In a simple self-assembly inspired simulation these—in N —growing systems tend to follow a chain of intermediate incomplete arrangements of the same building plan, leading preferentially (relatively frequently) to the nearby energetically most favoured N value with the building plan completed.

It is found that GEM LJ-arrangements have on average a lower symmetry than those for a purely repulsive interaction. Upon radial compression, which increases the influence of the LJ repulsive core, indeed the LJ GEM states tend to increase their level of symmetry significantly and d-charge growth tends to be suppressed. Results of a series of simulations with built-in icosahedral symmetry at large N are also presented. The thesis ends with an overall discussion and conclusions.

The obtained results add to our understanding of spherical 2D many particle systems in general, and Lennard-Jones systems in particular. They give hints towards physical mechanisms that could cause discreteness. More elaborate studies are needed to show if the followed approach in the future can be developed towards modelling natural spherical shells more realistically.

Samenvatting

Het zoeken naar een verklaring voor de discrete grootteverdeling van sommige in de natuur voorkomende bolschilvormige systemen gevormd door zelf-assemblage—zoals biomembraan blaasjes (vesicles) en virussen—vormt een eerste fysieke motivatie voor het onderzoek gepresenteerd in dit proefschrift. In onze poging aspecten relevant voor de intrinsieke stabiliteit van deze natuurlijke systemen te vatten in een simpel model, beschouwen wij N deeltjes met Lennard-Jones (LJ) interactie op het oppervlak van een bol met straal R . Het ontstaan van een energetisch optimale ordening wordt gemodelleerd met behulp van kristallisatie op een bol.

Veel kristallisatie studies zijn uitgevoerd in vlakke 2D systemen. Echter de thermodynamica van vergelijkbare systemen op bolvormige oppervlakken kan verschillen door de aanwezigheid van disclinaties (“d-ladingen”), een noodzakelijk gevolg van de kromming, welke het voor vlakke systemen optimale hexagonale rooster verstoren. De inherente d-ladingen kunnen de rol hebben van beginpunten in een hiërarchische ordening. Vanuit dit oogpunt vormt het begrijpen van het thermodynamische gedrag en de eigenschappen van bolvormige 2D LJ systemen—en het vergelijken met vlakke systemen—een tweede fysieke motivatie voor het werk beschreven in dit proefschrift.

Na het introductie-hoofdstuk met de motivaties voor het onderzoek, de gebruikte begrippen en enige historische opmerkingen, wordt de relevante theoretische achtergrond en een beknopt literatuuroverzicht weergegeven in Hoofdstuk 2.

Het huidige werk behelst experimenten met behulp van computersimulatie. Simulaties van het gebruikte modelsysteem worden uitgevoerd en de eigenschappen van het model worden gemeten en geanalyseerd. Voor de kristallisatiestudies gebruiken wij hoofdzakelijk het Simulated Annealing (SA) algoritme met een Monte Carlo simulatie bij elke temperatuur (T), en een stochastisch geïnitieerde Steepest Descent (SD) energie minimalisatie algoritme, beiden met een vaste N -waarde. Deze algoritmen worden in detail besproken in Hoofdstuk 3. Uit methodologische studies concluderen wij dat SA in vergelijking met SD efficiënter is in het vinden van de meest stabiele oplossingen: de globale energie minima (GEM). We nemen waar dat er een optimum is in de kwaliteit ten opzichte van de rekentijd, afhankelijk van de eigenschappen van het energielandschap over de toestandsruimte van het systeem (welke varieert met N).

Omdat de SA methode zeer rekenintensief is en er honderdduizenden onafhankelijke simulaties nodig zijn, is het toepassen van een of andere vorm van parallelisme noodzakelijk. In Hoofdstuk 4 beschouwen wij vier verschillende paralleliserings-strategieën. Vectorisering van de simulatiecode is relatief makkelijk en met een vectorsupercomputer kan een hoge tijdwinst behaald worden. Voor de twee decompositie methoden, systolisch en data-parallel, blijkt dat de eerste niet met behoud aan kwaliteit van de resultaten schaalbaar is naar een groot aantal processoren. De tweede methode blijkt zeer efficiënt te zijn als het aantal processoren wordt afgestemd op het aantal deeltjes. De combinatie van de twee decompositie methoden, de hybride implementatie, is efficiënt schaalbaar naar een groot aantal processoren. Vanuit een tijd-complexiteits-analyse concluderen wij dat op moderne parallelle—mogelijk heterogene—computer systemen vooral de data-decompositie aangepast moet worden. Omdat wij toegang hebben tot een grote hoeveelheid rekenkracht in de vorm van door netwerken verbonden werkstations via het Condor systeem, gebruiken wij voornamelijk “jobparallelisme”, met decompositie van de parameter ruimte.

Temperatuur-afhankelijke simulaties van het kristallisatieproces worden beschreven in Hoofdstuk 5. Uit het gedrag van het bolvormige 2D LJ-systeem bij een constante dichtheid (constante R) als een functie van T , leiden wij af dat bevrozing van een klein systeem optreedt over een breed temperatuurgebied beneden het vriespunt van een groot vlak systeem. Voor $N=100$, waar wij ons in de temperatuur-afhankelijke studies op concentreren, vinden wij dat het bevrozen plaats vindt in het temperatuurgebied van $T=0.8$ tot $T=0.2$. De radiële distributiefuncties voor de verschillende d-ladingen geven aan dat d-ladingen met gelijk teken effectief elkaar afstoten en met ongelijk teken elkaar aantrekken. De aantrekking leidt tot concentratie (clustering) van d-ladingen en de afstoting tot een tendens naar icosaeëdrische (regelmatig twintigvlak) symmetrie in de ordening van clusters met netto d-lading -1 . Reeds gedurende het bevrozen vertonen zich systematische thermodynamische verschillen bij variatie van N . Deze blijken gecorreleerd te zijn met de symmetrie en pakkingefficiëntie, dus met lange drachts positie- en bindings-oriëntatie-ordening, van de GEM toestanden bij de nultemperatuur.

De resulterende kristallijne verdelingen bij temperatuur en externe druk gelijk aan nul, worden beschreven in Hoofdstuk 6. Daar vinden wij dat er minima bestaan in het energielandschap over de gecombineerde toestandsruimten voor verschillende N met een stabiliteitsbereik in zowel R als N . Bij geringe kromming (grote R) bestaan er veel verschillende minima die corresponderen met “open” verdelingen waarbij het boloppervlak slechts ten dele wordt bezet. Voor “gesloten” verdelingen zijn er locale minima dicht bij de GEM toestand. Het aantal van deze locale minima neemt exponentieel toe met N . De GEM resultaten geven duidelijk aan dat er een energetische voorkeur bestaat voor specifieke N -waarden in vergelijking met het gladde algemene energieverloop, waarbij de beperkingen in iedere pakking niet zijn meegenomen. Voor grotere systemen

blijkt dat er een energetische voorkeur bestaat voor bepaalde gebieden in N . De voorkeursverdelingen hebben gemiddeld een relatief hoge dichtheid en symmetrie. De verstoring in het hexagonale rooster rondom de inherente d-ladingen heeft de neiging toe te nemen met de systeemgrootte (“d-ladings polarisatie”). Wij vinden dat de groei van d-ladings-cluster grootte in het bestudeerde N -gebied—vanwege de neiging naar symmetrie—stijl verloopt van 12 naar 36. In het overgangsgebied in N zijn de verdelingen van relatief lage symmetrie en dichtheid, en energetisch minder gunstig dan het algemene energieverloop als functie van N .

In een uitvoerige SD-studie van globale en lokale minima, doen zich energetisch gunstige “bouwplannen” gelden. Deze vertonen de structuur van stabiele gesloten verdelingen, meestal met een hoge symmetrie, waarin een aantal deeltjes ontbreken (daarom zijn ze open), maar toch metastabiel zijn. In een vereenvoudigde—op het zelf-assemblage proces geïnspireerde—simulatie, volgen de in N groeiende systemen hetzelfde bouwplan, bij voorkeur (relatief veelvuldig) leidend naar de energetisch meest gunstige N -waarde behorend bij het voltooide bouwplan.

Wij vinden dat GEM LJ-verdelingen gemiddeld een lagere symmetrie hebben dan verdelingen voor zuiver afstotende interacties. Bij hogere dichtheid, waar de afstotende LJ kern belangrijker wordt, treedt inderdaad een significante stijging van de symmetrie van de GEM toestanden op, en de toename van d-ladingen wordt onderdrukt. De resultaten van een serie simulaties waarin de icosaeëdrische symmetrie is ingebouwd worden besproken. Het proefschrift eindigt met een discussie en conclusies over de fysische resultaten.

De bereikte resultaten verhogen ons begrip van bolvormige 2D veeldeeltjes systemen in het algemeen, en van Lennard-Jones systemen in het bijzonder. Ze geven mogelijke fysische mechanismen aan die kunnen leiden tot discretisatie in grootte. Uitgebreidere studies zijn nodig om aan te tonen of de gevolgde aanpak in de toekomst kan worden uitgebouwd tot een realistischer modellering van natuurlijke bolschilvormige systemen.

Nawoord

Het werk beschreven in dit proefschrift heb ik met veel plezier uitgevoerd. Voor het stimuleren en begeleiden van mijn onderzoek ben ik veel dank verschuldigd aan mijn promotor Peter Sloot en aan mijn copromotor René van Dantzig. Zij vulden elkaar goed aan en ik heb op verschillende gebieden veel van hun kunnen leren. Hun bijdrage heeft een belangrijke invloed op het werk en het proefschrift gehad.

Van Peter heb ik veel hulp gehad op het gebied van High Performance Computing. Ook hield hij de grote lijnen in het oog. Ik ben blij dat hij mijn humor altijd goed heeft aangevoeld. Dankzij René heb ik veel gebruik kunnen maken van High Throughput Computing, tevens stond hij mij met veel oog voor detail, met raad en daad terzijde. Via herhaaldelijke correcties van mijn teksten door met name René heeft dit proefschrift de huidige kwaliteit en leesbaarheid verkregen.

De discussies met Daan Frenkel en de bezoeken aan de werkbijeenkomsten van zijn groep op het AMOLF hebben mijn inzichten in thermodynamische computermodellen en simulatietechnieken verdiept.

Voordat ik aan dit project begon, had Arne ter Laak al een begin gemaakt met dit onderzoek. Dit gaf mij een vliegende start.

Diederik Burer heeft gedurende zijn stage zowel een sequentiële als parallelle Moleculaire Dynamica implementatie van de simulatie code gemaakt. In dit proefschrift heb ik dankbaar gebruik gemaakt van zijn interessante resultaten. Zijn open persoonlijkheid en humor zullen mij nog lang bijblijven.

Then, I gratefully acknowledge the contribution to my work from Miron Livny and his group at the University of Wisconsin in Madison. The enormous amount of computing power needed for this project could be acquired thanks to access to the large Condor pools in Madison and the possibility to use large amounts of computing power from other institutes. In particular from the University of Bologna, but also from NIKHEF Amsterdam, the University of Amsterdam, Delft University of Technology, and during shorter periods from CERN, Warsaw University and JINR in Dubna. These cooperations allowed to expand the work beyond original expectations.

I wish to thank J. Edmundson for his symmetry test algorithms that provided a basis for developing part of my analysis.

Veel van het plezier in mijn werk komt van de sfeer in onze groep en de hulp-

vaardigheid die ik genoten heb. Veel dank aan mijn collega's: Alfons, Andy, Arjan, Arjen, Benno, Berry, David, Dick, Diederik, Drona, Frank, Gera, Giel, Hugh, Jaap, Jan, Joep, Marcel, Martin, Peter, Rob en Walter. Speciaal gedenkwaardig zijn de wandeltochten door de Ardennen. Dankzij hopman Benno is er niemand ten prooi gevallen aan "killer snails" of roofmussen. Aan deze tochten, het zeilen en diverse conferentie bezoeken zal ik nog regelmatig met veel genoegen terugdenken.

Randvoorwaarden voor prettig werken is een goed systeembeheer en Secretariaat. Hiervoor veel dank aan Frans, Gert, Hugo, Jan, Laura, Monique, Richard en Virginie. Helaas is Ina niet meer onder ons, haar kunde en vriendelijkheid zal ik niet vergeten.

Alle nog niet genoemde collega's en medewerkers van WINS die mijn werk mogelijk hebben gemaakt of anderszins hebben ondersteund: bedankt!

Naast het werk heeft een aantal mensen gezorgd dat ik regelmatig stoom kon afblazen. De "Bijlmer crew", met als belangrijkste exponenten Annemiek, Guus, Marko en Ronald, wordt bedankt voor de met alcohol benevelde avonden. Ik bedank Arjan voor de vele nachten vol met hilarische en ietwat onrealistische ideeën. Jacco wil ik bedanken voor de vele keren dat ik langs kon komen voor discussies, wijn en darten.

Ik bedank mijn ouders, en met name mijn moeder, dat ze mij gestimuleerd hebben om toch vooral door te leren. Ik denk dat ik wel mag stellen dat ik dit advies ter harte genomen heb. Jolanda, Sigrid, Henk en Richard wil ik bedanken voor hun interesse in mijn bezigheden en hun steun.

Tenslotte wil ik Bibian bedanken voor haar volkomen vertrouwen in mij. Haar morele steun en liefde heb ik zeer nodig gehad de afgelopen jaren. En nog.

Appendix A

Energy minima for $2 \leq N \leq 200$

Table with for each N the lowest energy state. Of each state the energy (E), radius (R), coordination number distribution and symmetry (S) group are given.

N	E	R	d-charges					S
			-2	-1	0	1	2	
2	-0.5000000000000000	0.5000000000000000	0	0	0	0	0	D_∞
3	-1.0000000000000000	0.597883634403098	0	0	0	0	0	D_{3h}
4	-1.5000000000000000	0.612372437659237	4	0	0	0	0	T_d
5	-1.696124031007751	0.701013442578522	2	3	0	0	0	C_{4v}
6	-2.118677042801556	0.703946850539554	0	6	0	0	0	O_h
7	-2.004472838219701	0.784996743186216	1	3	3	0	0	C_{3v}
8	-2.263564708273913	0.813761566788938	0	4	4	0	0	D_{4d}
9	-2.354168318846044	0.853466867104140	0	3	6	0	0	D_{3h}
10	-2.391701447066171	0.897777352533909	0	3	6	1	0	C_{3v}
11	-2.546447320800543	0.940905005832123	0	2	8	1	0	C_{5v}
12	-2.799795573726781	0.942373155293769	0	0	12	0	0	I_h
13	-2.448969977079452	1.023669635576872	0	1	10	2	0	C_s
14	-2.533348412368524	1.053553039688811	0	0	12	2	0	C_{2v}
15	-2.604254157790561	1.079511939730079	0	0	12	3	0	C_s
16	-2.629397082118763	1.111359350373031	0	0	12	4	0	C_s
17	-2.729322104279705	1.134715561188902	0	0	12	5	0	D_{5h}
18	-2.720681589556329	1.168431468280902	0	2	8	8	0	D_{4d}
19	-2.692582368657945	1.212803379999540	0	0	12	7	0	C_s
20	-2.812968910809319	1.221102153780282	0	0	12	8	0	D_{3h}
21	-2.761570772615670	1.256842828104036	0	0	12	9	0	D_3
22	-2.814127147712486	1.278935841126762	0	0	12	10	0	T_d
23	-2.798822415652630	1.323636299622494	0	1	11	10	1	C_1
24	-2.923589586974076	1.325942483975014	0	0	12	12	0	O
25	-2.790744681094829	1.370215612837342	0	0	12	13	0	C_1
26	-2.838736266635434	1.393253846649313	0	0	12	14	0	C_s
27	-2.915404947139873	1.405226921912913	0	0	12	15	0	D_3
28	-2.834795918128242	1.441037402745361	0	0	12	16	0	C_2
29	-2.836943974529299	1.470703678854806	0	0	13	15	1	C_1
30	-2.907879985813801	1.482942826360981	0	0	12	18	0	D_3
31	-2.914845518605604	1.508252602796314	0	0	13	17	1	C_s

32	-2.980094374796669	1.517799565207946	0	0	12	20	0	I_h
33	-2.857364512086475	1.565483089899774	0	0	14	17	2	C_s
34	-2.879973993426453	1.581507917153482	0	0	12	22	0	C_2
35	-2.897476940814867	1.603568457619434	0	0	12	23	0	C_s
36	-2.926074652855166	1.621869926830591	0	0	12	24	0	D_2
37	-2.907674813755474	1.645689466050403	0	0	12	25	0	C_1
38	-2.972138629420019	1.659845132455830	0	0	12	26	0	D_{6d}
39	-2.913516422413766	1.689160940136996	0	0	12	27	0	C_3
40	-2.935326506884592	1.708230437866873	0	0	15	22	3	C_3
41	-2.929234832467207	1.729343588048532	0	0	12	29	0	D_3
42	-2.957197169053093	1.752712826655803	0	0	14	26	2	C_{2v}
43	-2.986878602027695	1.765631923177463	0	0	13	29	1	C_s
44	-3.039204045421928	1.773591356738300	0	0	12	32	0	O_h
45	-2.989550533031941	1.801077961986298	0	0	12	33	0	D_3
46	-2.967979946789361	1.831567835990777	0	0	14	30	2	C_{2v}
47	-2.991666695749649	1.845624666183277	0	0	13	33	1	C_1
48	-3.043710551686008	1.852527998433596	0	0	12	36	0	O
49	-2.961259366808124	1.884936044669402	0	0	12	37	0	C_3
50	-2.984583094280973	1.899868057548674	0	0	12	38	0	D_3
51	-2.991978366226236	1.916517435225234	0	0	12	39	0	C_{2v}
52	-2.988779822951595	1.937974048925100	0	0	15	34	3	C_3
53	-2.974816531472269	1.961895379519709	0	0	16	33	4	C_1
54	-3.008467736524301	1.971909776092477	0	0	16	34	4	D_2
55	-2.973425131692696	1.997949144046892	0	0	16	35	4	C_1
56	-3.005994640031623	2.007994951024024	0	0	16	36	4	D_2
57	-2.994392848924415	2.027327589626184	0	0	16	37	4	C_2
58	-2.997051140312090	2.044677513303251	0	0	12	46	0	C_2
59	-3.010013049113121	2.060451925426821	0	0	15	41	3	C_s
60	-3.032786493941144	2.072914876983660	0	0	12	48	0	T
61	-3.005979486009581	2.099984988857275	0	0	18	37	6	C_{2v}
62	-3.021158756037753	2.109459052546311	0	0	16	42	4	D_2
63	-3.022031670943843	2.125802560540739	0	0	12	51	0	D_3
64	-3.020706178561096	2.143018548376299	0	0	12	52	0	T
65	-3.023038735530462	2.159849647418891	0	0	12	53	0	D_3
66	-3.034797798662134	2.173766775635421	0	0	12	54	0	D_3
67	-3.033218086181598	2.190095552387426	0	0	16	47	4	C_{2v}
68	-3.015064415231199	2.212957714506952	0	0	16	48	4	C_s
69	-3.034610208057660	2.221952321872294	0	0	12	57	0	D_3
70	-3.042587333369762	2.236128944654069	0	0	12	58	0	C_2
71	-3.048006785146135	2.253609149146347	0	0	13	57	1	C_{5v}
72	-3.057702042394375	2.264321813953698	0	0	12	60	0	D_{5h}
73	-3.019407304509814	2.292383040075030	0	0	14	57	2	C_1
74	-3.030862374610875	2.303497514639481	0	0	14	58	2	C_1
75	-3.044593521225616	2.315406949106456	0	0	18	51	6	D_3
76	-3.043315920486785	2.332264011995436	0	0	17	54	5	C_s
77	-3.052751366790504	2.342682708578410	0	0	16	57	4	C_{2v}
78	-3.068799176876690	2.355651081549880	0	0	18	54	6	D_3
79	-3.047756931955313	2.373368483194140	0	0	13	65	1	C_s
80	-3.055012900625701	2.387441437502674	0	0	12	68	0	D_4

81	-3.040443111613370	2.405496835695788	0	0	16	61	4	C_2
82	-3.036629161025914	2.421820588971989	0	0	16	62	4	C_2
83	-3.039736810516653	2.434795696871155	0	0	14	67	2	C_2
84	-3.051926932935712	2.448182481980792	0	0	12	72	0	D_3
85	-3.041467105409450	2.464409004969828	0	0	15	67	3	C_1
86	-3.044445037863451	2.478604188594150	0	0	12	74	0	D_2
87	-3.048615801155628	2.493077153950914	0	0	18	63	6	D_3
88	-3.059034645236533	2.504241465416455	0	0	12	76	0	D_2
89	-3.054668510384810	2.519344738182221	0	0	12	77	0	C_2
90	-3.049877137471776	2.540613112183650	0	0	16	70	4	C_s
91	-3.052638760612254	2.551677109090950	0	0	15	73	3	C_1
92	-3.068391387273855	2.559448724569366	0	0	12	80	0	D_6
93	-3.061580134667455	2.578325483646422	0	0	16	73	4	C_1
94	-3.070146514964345	2.588610373791938	0	0	14	78	2	C_1
95	-3.072084399988970	2.600111381805944	0	0	13	81	1	C_1
96	-3.081425625803969	2.609596703568227	0	0	12	84	0	D_2
97	-3.076135067207275	2.626721288868675	0	0	15	79	3	C_1
98	-3.090663894975623	2.635565739020858	0	0	12	86	0	D_6
99	-3.068095087770885	2.653858334618032	0	0	13	85	1	C_1
100	-3.075249310690257	2.663546522522910	0	0	12	88	0	T
101	-3.071708478005182	2.678676628924295	0	0	12	89	0	D_3
102	-3.076205852417322	2.692536029989609	0	0	13	88	1	C_1
103	-3.079858373389464	2.702954624202101	0	0	12	91	0	C_2
104	-3.087489673137782	2.714674249635795	0	0	12	92	0	C_2
105	-3.069764576307374	2.736845130822934	0	0	19	79	7	C_1
106	-3.072160353246025	2.747420325960794	0	0	16	86	4	C_2
107	-3.076797005712412	2.757575253277746	0	0	16	87	4	C_1
108	-3.085351413191244	2.766675673934666	0	0	16	88	4	D_2
109	-3.070297424622890	2.791876508194072	0	0	20	81	8	C_1
110	-3.073320515205445	2.794488290080189	0	0	12	98	0	D_6
111	-3.074449327677529	2.814741297900506	0	0	21	82	7	C_1
112	-3.073604112657678	2.824143307160662	0	0	16	92	4	C_1
113	-3.077528428107247	2.836259164448559	0	0	18	90	4	C_s
114	-3.077496922372325	2.854835788105537	0	0	20	86	8	C_1
115	-3.083128845321891	2.862330006539489	0	0	18	91	6	C_2
116	-3.080362396947650	2.875555469800447	0	0	17	94	5	C_1
117	-3.081073442697559	2.888864081502031	0	0	19	91	7	C_1
118	-3.084263678897037	2.900961541460660	0	0	18	94	6	C_2
119	-3.083438853744268	2.911113111283086	0	0	17	97	5	C_1
120	-3.088398471288236	2.924439498447144	0	0	22	88	10	C_2
121	-3.088093362420383	2.937669505381374	0	0	21	91	9	C_1
122	-3.088159482846123	2.947235958134594	0	0	21	92	9	C_1
123	-3.089366461830898	2.957466536428576	0	0	18	99	6	C_1
124	-3.092400847342048	2.969134041155647	0	0	19	99	5	C_1
125	-3.095068828327779	2.978067018734893	0	0	18	101	6	C_2
126	-3.100850044117253	2.986325970656095	0	0	24	96	0	O
127	-3.095151141742201	3.004704971800503	0	0	17	105	5	C_1
128	-3.096888075106277	3.014665785390008	0	0	16	108	4	C_1
129	-3.098316057023640	3.024828419306609	0	0	15	111	3	C_3

130	-3.103471238049764	3.040399569127699	0	0	19	104	7	C_1
131	-3.097707717375346	3.052003103659474	0	0	22	100	8	C_1
132	-3.102630726578489	3.061159789567351	0	0	18	108	6	C_1
133	-3.102951229779080	3.071393348487630	0	0	16	113	4	C_2
134	-3.103386861016297	3.078759347558850	0	0	24	104	0	O
135	-3.101055226567179	3.092163686671983	0	0	16	115	4	C_2
136	-3.105348876078927	3.110217048189944	0	0	21	106	9	C_1
137	-3.109978263589873	3.120984330660666	0	0	22	105	10	C_1
138	-3.110505431525404	3.129560806432384	0	0	20	110	8	C_1
139	-3.111035514524276	3.140794740156268	0	0	20	111	8	C_1
140	-3.111971786880358	3.150410402410642	0	0	19	114	7	C_1
141	-3.110939438934202	3.161439523981784	0	0	20	113	8	C_1
142	-3.112483255720369	3.170039357404192	0	0	20	114	8	D_2
143	-3.106175750329642	3.190420128454142	0	0	23	109	11	C_1
144	-3.108221399689585	3.198578091349305	0	0	23	111	9	C_1
145	-3.114053511847666	3.208952135748989	0	0	22	113	10	C_1
146	-3.115975584949435	3.221888423244658	0	0	24	110	12	C_1
147	-3.113922407297399	3.230960932715713	0	0	22	115	10	C_1
148	-3.118508719035905	3.240234954169562	0	0	22	116	10	C_1
149	-3.118803517629518	3.252644916206779	0	0	23	115	11	C_1
150	-3.128182173670630	3.260137032630535	0	0	29	106	14	C_1
151	-3.128472773014002	3.273122889986749	0	0	24	115	12	C_3
152	-3.125524342603588	3.282077715854691	0	0	22	120	10	C_2
153	-3.130896531897768	3.291565201540283	0	0	22	121	10	C_1
154	-3.127005007010226	3.305359595686113	0	0	24	118	12	C_1
155	-3.132862706538736	3.312386924021478	0	0	24	120	10	C_1
156	-3.127844478080208	3.322675173258684	0	0	21	126	9	C_1
157	-3.131124904973331	3.334680130528265	0	0	23	123	11	C_1
158	-3.133146802879490	3.343811287785840	0	0	24	123	10	C_1
159	-3.133144334099244	3.352410719794076	0	0	24	125	8	C_1
160	-3.129909678635975	3.365588982581071	0	0	23	126	11	C_1
161	-3.129721444033362	3.375317381165808	0	0	23	127	11	C_1
162	-3.129459902293785	3.385975387683052	0	0	23	128	11	C_1
163	-3.132873338761440	3.395961708610330	0	0	22	131	10	C_1
164	-3.134884748249570	3.405299806397327	0	0	22	132	10	C_1
165	-3.131566474182685	3.416170671780072	0	0	22	133	10	C_1
166	-3.136590295472705	3.426329586307788	0	0	23	132	11	C_1
167	-3.135451640262044	3.437332963291338	0	0	24	131	12	C_1
168	-3.136678867724460	3.447589317158606	0	0	24	132	12	C_1
169	-3.135170939270683	3.456976495180792	0	0	24	133	12	C_1
170	-3.134942196241432	3.467642804146652	0	0	23	136	11	C_1
171	-3.140983709291654	3.476977682769265	0	0	24	135	12	C_1
172	-3.139302825850949	3.486693285069663	0	0	24	136	12	C_1
173	-3.140563338610686	3.494628748214911	0	0	21	143	9	C_1
174	-3.142182978369526	3.506380213208114	0	0	24	138	12	C_1
175	-3.141080577886517	3.515473274474434	0	0	23	141	11	C_1
176	-3.144540704778469	3.524971756942792	0	0	23	142	11	C_1
177	-3.152597719358375	3.533970574289147	0	0	24	141	12	C_2
178	-3.147652937449270	3.544449916384940	0	0	24	142	12	C_1

179	-3.146002042746655	3.554691664005756	0	0	24	143	12	C_1
180	-3.146214212187226	3.564175946389665	0	0	24	144	12	C_1
181	-3.150821602930910	3.572562184575907	0	0	23	147	11	C_1
182	-3.150601079740954	3.581030930736014	0	0	22	150	10	C_1
183	-3.152818607079521	3.592592388291539	0	0	24	147	12	C_1
184	-3.151194638019939	3.602484516634498	0	0	24	148	12	C_1
185	-3.156344034402267	3.610518189009657	0	0	24	149	12	C_2
186	-3.153910186180270	3.620859105803043	0	0	24	150	12	C_1
187	-3.153600127082671	3.630863189592233	0	0	24	151	12	C_1
188	-3.158556195410998	3.638931964728188	0	0	24	152	12	C_1
189	-3.155454762449801	3.649219412940159	0	0	24	153	12	C_1
190	-3.160387892650515	3.657905693919613	0	0	24	154	12	C_2
191	-3.158125596187284	3.670128971890075	0	0	25	154	11	C_1
192	-3.174933738472386	3.673165160750332	0	0	24	156	12	D_3
193	-3.161625768983771	3.684622302462571	0	0	23	159	11	C_1
194	-3.161354619492706	3.695001055737960	0	0	24	158	12	C_1
195	-3.161715211878604	3.704139336030599	0	0	24	159	12	C_2
196	-3.164648708988006	3.713157374307327	0	0	24	160	12	C_1
197	-3.166429333188981	3.721755937513140	0	0	24	161	12	C_1
198	-3.166965274636082	3.731319964188270	0	0	24	162	12	C_2
199	-3.161323530591782	3.741213153036003	0	0	23	165	11	C_1
200	-3.163419118298085	3.750331427267608	0	0	24	164	12	C_1

Appendix B

Minima with built-in icosahedral symmetry

The optimal states with icosahedral symmetry, E_{ann} are the found energies after annealing with enforced symmetry, E_{icosd} is the energy after SD with enforced sym. and E_{sd} is the energy after applying SD without enforced symmetry.

N	A	B	C	D	E	E_{sdico}	R_{sdico}	E_{sd}	R_{sd}
60	0	0	0	1	0	-1.8456548820	2.4364728156	-1.8456548820	2.4364727433
72	1	0	0	1	0	-2.0710491688	2.6415535352	-2.0710491688	2.6415534770
80	0	0	1	0	1	-2.9988253510	2.4223303895	-2.9988253510	2.4223303215
80	0	0	1	1	0	-2.9988253510	2.4223303886	-2.9988253510	2.4223302837
90	0	1	0	0	1	-2.3184568247	2.8219687277	-2.3184568247	2.8219686925
90	0	1	0	1	0	-1.5691326302	3.3229274431	-1.5691326302	3.3229273638
92	1	0	1	0	1	-2.9518841896	2.5762077434	-2.9518841896	2.5762076936
92	1	0	1	1	0	-2.9518841896	2.5762077426	-2.9518841896	2.5762076936
102	1	1	0	0	1	-2.5041552455	2.9102852165	-2.5041552455	2.9102851652
102	1	1	0	1	0	-1.7415706679	3.5392964796	-1.7415706679	3.5392962807
110	0	1	1	0	1	-3.0911374960	2.8106796448	-3.0911374960	2.8106795951
110	0	1	1	1	0	-1.7915784347	3.2439749062	-1.7915784347	3.2439747977
120	0	0	0	1	1	-2.7491084512	3.0175272300	-3.1090231045	2.9285242146
120	0	0	0	2	0	-1.4303226221	4.2178272319	-1.4303226228	4.2178250956
122	1	1	1	0	1	-3.0656933775	2.9392795722	-3.0656933775	2.9392795334
122	1	1	1	1	0	-1.7999757684	3.4787591578	-1.7999757684	3.4787590498
132	1	0	0	1	1	-2.4539122039	3.3382059712	-2.4539122207	3.3382190079
132	1	0	0	2	0	-1.5709038405	4.4387899103	-1.5709038415	4.4387947890
140	0	0	1	1	1	-1.9549766163	3.6578568775	-3.0373360917	3.2047508029
140	0	0	1	2	0	-1.3516979226	4.1836958295	-1.3516979226	4.1836954374
150	0	1	0	1	1	-3.1040869608	3.2585889217	-3.1040869622	3.2585876903
152	1	0	1	1	1	-2.0786706986	3.7506005315	-2.0786708759	3.7506633624
152	1	0	1	2	0	-1.4545865608	4.4134582596	-1.4545865638	4.4134554254
162	1	1	0	1	1	-3.0112886213	3.4000151820	-3.0112886213	3.4000151274
170	0	1	1	1	1	-2.2257332501	3.7979118194	-3.1126885428	3.4686566921
180	0	0	0	2	1	-2.3843159117	4.1269998770	-2.3843159804	4.1269139066
182	1	1	1	1	1	-2.3360764759	3.8724461367	-3.1032270521	3.6015700857
192	1	0	0	2	1	-2.3272246973	4.3312953371	-2.3272247001	4.3312776753

192 APPENDIX B. MINIMA WITH BUILT-IN ICOSAHEDRAL SYMMETRY

200	0	0	1	2	1	-2.4992927976	4.0447032876	-3.1398369456	3.7464078485
210	0	1	0	2	1	-2.1612201379	5.0273496145	-2.1612201767	5.0272088658
212	1	0	1	2	1	-2.3583120774	4.2488441537	-3.1199538036	3.8650429606
230	0	1	1	2	1	-2.0271027552	5.0047389539	-2.0271027553	5.0047379238
242	1	1	1	2	1	-1.9882161260	5.2179958065	-2.9498069681	4.3461943952
270	0	1	0	2	2	-2.4563911892	4.9874106927	-2.4564668933	4.9880516412
282	1	1	0	2	2	-2.3929833259	5.1770202079	-3.1054480907	4.5057799973
290	0	1	1	2	2	-2.5367827397	4.9425684789	-2.5392788808	4.9408546681
302	1	1	1	2	2	-2.4387620106	5.1051876260	-3.0884102537	4.6663294766
320	0	0	1	3	2	-2.1738736433	5.8717406073	-2.5191713034	5.4983281376
330	0	1	0	2	3	-2.9953894800	4.9539073018	-2.9953894803	4.9539121139
332	1	0	1	3	2	-2.1252572457	6.0654856438	-2.9891913981	5.0013971879
342	1	1	0	2	3	-2.8822159161	5.0917604261	-2.8822159173	5.0917791025
350	0	1	1	2	3	-3.1338279034	4.9355651521	-3.1338279073	4.9355937897
360	0	0	0	3	3	-2.6012145314	5.8306949664	-2.6012145324	5.8306719569
362	1	1	1	2	3	-3.0716828961	5.0566679404	-3.0716828979	5.0566880788
372	1	0	0	3	3	-2.5134241939	6.0140140611	-2.5134247703	6.0133584548
380	0	0	1	3	3	-2.4512317131	5.7952556126	-3.2168273300	5.1375689729
392	1	0	1	3	3	-2.4299991017	5.9828436010	-2.4299991020	5.9828302736
420	0	0	0	3	4	-2.8786710289	5.8357569011	-2.8786710517	5.8356595499
432	1	0	0	3	4	-2.7909085471	5.9770184260	-2.7909085723	5.9771291397
440	0	0	1	3	4	-2.8628061888	5.7967871000	-2.8628062224	5.7969018820
450	0	1	0	3	4	-2.5837171500	6.7119661209	-2.5837475388	6.7129003452
452	1	0	1	3	4	-2.7612298609	5.9397106838	-2.7612298614	5.9397244454
462	1	1	0	3	4	-2.5093481453	6.8827680127	-2.9114419390	6.1315326362
470	0	1	1	3	4	-2.4976213677	6.6924383354	-2.6282767915	6.7083378239
482	1	1	1	3	4	-2.3486445394	6.8049565525	-3.1071026455	5.9042925382
510	0	1	0	3	5	-2.8070836643	6.7239300766	-2.8070836877	6.7238085928
522	1	1	0	3	5	-2.7348454134	6.8670230711	-2.7348454410	6.8671637096
530	0	1	1	3	5	-2.7326595549	6.7097289203	-2.7326595810	6.7095996723
542	1	1	1	3	5	-2.5844490574	6.7483371568	-3.1171676607	6.2541478636
570	0	1	0	3	6	-2.8768429994	6.5957928266	-2.8789015422	6.5950430314
582	1	1	0	3	6	-2.8654081591	6.7939748479	-2.8654081749	6.7938769026
590	0	1	1	3	6	-2.3932406660	8.5056905129	-2.8693149370	7.0923628756
600	0	0	0	4	6	-2.7593382370	7.6149231902	-2.7593382538	7.6148020264
602	1	1	1	3	6	-2.6694129871	7.0430724440	-3.0652822353	6.6382431750
612	1	0	0	4	6	-2.3986703188	7.7206556798	-3.0565056855	6.7743653051
620	0	0	1	4	6	-2.4620933057	7.4411163264	-3.1534719375	6.6258162164
632	1	0	1	4	6	-2.3065044467	7.6255517573	-3.0793552119	6.7818316117
660	0	0	0	4	7	-2.7215977957	7.5388851315	-2.7254437948	7.5593763112
672	1	0	0	4	7	-2.8024848825	7.6866326181	-2.8919181480	7.6487147917
680	0	0	1	4	7	-2.1538156610	9.0033481074	-2.8841141485	7.5387748937
692	1	0	1	4	7	-2.6618668338	7.6021787543	-3.1025483189	7.0607898620
720	0	0	0	4	8	-2.9356531332	7.5232392780	-2.9368806757	7.5280504027
732	1	0	0	4	8	-2.3227790945	9.5384656109	-2.6654245217	8.9165954740
740	0	0	1	4	8	-2.9558262914	7.4260906010	-2.9573636109	7.4246924670
752	1	0	1	4	8	-2.8676004886	7.5431687389	-3.0971181026	7.3638440573
770	0	1	1	4	8	-2.4536386456	8.4766440587	-2.9656060597	7.6629044410
782	1	1	1	4	8	-2.0693137067	9.8578328927	-3.0634344400	7.5990307520

Bibliography

- [1] W.S. Bont. Geometric progression in the size of membrane vesicles, nuclei and cells. *European Journal of Cell Biology*, 39:485–499, 1985.
- [2] A. Klug. Architectural design of spherical viruses. *Nature*, 303:378–379, 1983.
- [3] M.J. Wenninger. *Spherical models*. Cambridge University Press, 1979.
- [4] J.J. Thomson. On the structure of the atom. *Phil. Mag. S. 6*, 7(39):237–265, 1904.
- [5] P.M.L. Tammes. On the origin of number and arrangement of the places of exit on the surface of pollen-grains. *Recueil des Travaux Botaniques Neerlandais*, 27:1–84, 1930.
- [6] R. Buckminster Fuller and E.J. Applewhite. *Synergetics*. Macmillan Publishing Company, 1982.
- [7] L. Föpple. Stabile anordnungen von electronen im atom. *J. Reine Angew. Math.*, 141:251–302, 1912.
- [8] H.S.M. Coxeter. The problem of packing a number of equal nonoverlapping circles on a sphere. *Trans. N.Y. Acad. Sci. Ser. II*, 24:320–331, 1962.
- [9] L. Fejes Toth. *Regular Figures*, volume 48 of *International series of monographs on pure and applied mathematics*. Pergamon press, 1964.
- [10] R.M. Robinson. Finite sets of points on a sphere with each nearest to five others. *Math. Annalen*, 179:296–318, 1969.
- [11] R. van Dantzig. Private communications.
- [12] P.M.A. Sloot, J.A. Kaandorp, and A. Schoneveld. Dynamic Complex Systems (DCS): a new approach to parallel computing in computational physics. Technical Report TR-CS-95-08, University of Amsterdam, 1995.
- [13] J.H. Reif and S.R. Tate. The complexity of N-body simulation. In A. Lingas, R. Karlsoon, and S. Carlsson, editors, *Automata, Languages and Programming Proceedings*, pages 163–175, 1993.
- [14] L.T. Wille and J. Vennik. Computational complexity of the ground-state determination of atomic clusters. *J. Phys. A: Math. Gen.*, 18:L419–L422, 1985.
- [15] J.P.K. Doye. *The structure, thermodynamics and dynamics of atomic clusters*. PhD thesis, University of Cambridge, Department of Chemistry, Cambridge, 1996. (and references therein).
- [16] T. V. Ramakrishnan and M. Yussouff. First-principles order-parameter theory of freezing. *Phys. Rev. B*, 19(5):2775–, 1979.
- [17] W.J. Moore. *Physical Chemistry*. Prentice-hall, New Jersey, 1972.

- [18] D. Frenkel and J.P. McTague. Computer simulations of freezing and supercooled liquids. *Ann. Rev. Phys. Chem.*, 31:491–521, 1980.
- [19] K. J. Strandburg. Two-dimensional melting. *Reviews of Modern Physics*, 60(1):161–207, 1988.
- [20] N.W. Ashcroft and N.D. Mermin. *Solid State Physics*. Saunders College publishing, 1976.
- [21] N.D. Mermin. Crystalline order in two dimensions. *Phys. Rev.*, 176(1):250–254, 1968.
- [22] F.F. Abraham. Melting in two dimensions is first order : an isothermal-isobaric Monte Carlo study. *Phys. Rev. Lett.*, 44(7):463–466, 1980.
- [23] Y. Imry and L. Gunther. Fluctuations and physical properties of the two-dimensional crystal lattice. *Phys. Rev. B*, 3(11):3939–3945, 1971.
- [24] M.A. Glaser and N.A. Clark. Melting and liquid structure in two dimensions. *Adv. in Chem. Phys.*, LXXXIII:543–709, 1993.
- [25] M.P. Allen and D.J. Tildesley. *Computer Simulation of Liquids*. Oxford science publication, Clarendon press, 1987.
- [26] R. Collins. *Phase Transitions and Critical Phenomena*, volume 2. Academic Press, New York, 1972.
- [27] I.K. Crain. The Monte Carlo generation of random polygons. *Computers and Geosciences*, 4:131–141, 1978.
- [28] H.W. Kroto, J.R. Heath, S.C. O’Brien, R.F. Curl, and R.E. Smalley. C60 : Buckminsterfullerene. *Nature*, 318:162–163, 1985.
- [29] J. Han and R. Jaffe. Energetics and geometries of carbon nanoconic tips. to appear in *J. Chem. Phys.*
- [30] J. Han. Energetics and structures of fullerene crop circles. *Chem. Phys. Lett.*, 1997.
- [31] Yanmu Zhou. *Arrangements of points on the sphere*. PhD thesis, Graduate School, University of South Florida, Tampa, Florida, 1995.
- [32] B. Bergersen, D. Boal, and P. Palfy-Muhoray. Equilibrium configurations of particles on a sphere: the case of logarithmic interactions. *J. Phys. A: Math. Gen.*, 27:2579–2586, 1994.
- [33] F.W. de Wette, R.E. Allen, D.S. Hughes, and A. Rahman. Crystallisation with a Lennard-Jones potential: a computer experiment. *Phys. Lett.*, 29A(9):548–549, 1969.
- [34] D. Henderson, editor. *Fundamentals of inhomogeneous fluids*. Marcel Dekker inc., New York, 1992.
- [35] H.T. Croft, K.J. Falconer, and R.K. Guy. *Unsolved problems in geometry*. Springer-Verlag, New York, 1991.
- [36] J.P. Hansen, D. Levesque, and J.J. Weis. Self-diffusion in the two-dimensional, classical electron gas. *Phys. Rev. Lett.*, 43(14):979–982, 1979.
- [37] L. Euler. Elementa doctrinae solidorum. *Novi Commentarii Academiae Scientiarum Petropolitanae*, 4:109–140, 1758.
- [38] I. Lakatos. *Proofs and refutations*. Cambridge University Press, 1976.
- [39] S. P. Giarritta, M. Ferrario, and P. V. Giaquinta. Statistical geometry of hard particles on a sphere : analysis of defects at high density. *Physica A*, 201:649–665, 1993.

- [40] R. Eppenga and D. Frenkel. Monte Carlo study of the isotropic and nematic phases of infinitely thin hard platelets. *Mol. Phys.*, 52(6):1303–1334, 1984.
- [41] D. Frenkel. Monte Carlo simulations. In S.C. Parker C.R.A. Catlow and M.P. Allen, editors, *Computer Modelling of Fluids, Polymers and Solids*, page 113. Kluwer, Dordrecht, 1990.
- [42] V.I. Harismiadis, J. Vorholz, and A.Z. Panagiotopoulos. Efficient pressure estimation in molecular simulations without evaluating the virial. *J. Chem. Phys.*, 105:8469–8470, 1996.
- [43] D. Frenkel. *Free-energy computation and first-order phase transitions*. Soc. Italiana di Fisica - Bologna - Italia, 1986.
- [44] J. Møller. *Lecture Notes in Statistics, Lectures on Random Voronoi Tessellations*, volume 87. Springer-Verlag, 1994.
- [45] S.P. Giarritta, M. Ferrario, and P.V. Giaquinta. Statistical geometry of hard particles on a sphere. *Physica A*, 187:456–474, 1992.
- [46] M. Hamermesh. *Group theory and its application to physical problems*. Dover Publications Inc. New York, 1989.
- [47] A. Schoenflies. *Krystallsysteme und Krystalstruktur*, volume XII. Springer, Repring 1984.
- [48] T. Erber and G.M. Hockney. Complex systems : Equilibrium configurations of N equal charges on a sphere ($2 \leq N \leq 112$). Technical Report FERMILAB-PUB-95-075-T, Fermi National Accelerator Laboratory, 1995.
- [49] C. Udink. *Melting in two-dimensional systems; a computer simulation study*. PhD thesis, University of Amsterdam, 1987.
- [50] J.M. Kosterlitz and D.J. Thouless. Long range order and metastability in two-dimensional solids and superfluids. *J. Phys. C*, 5:L124–L126, 1972.
- [51] J.M. Kosterlitz and D.J. Thouless. Ordering, metastability and phase transitions in two-dimensional systems. *J. Phys. C*, 6:1181–1203, 1972.
- [52] D. R. Nelson and B. I. Halperin. Dislocation-mediated melting in two-dimensions. *Phys. Rev. B*, 19(5):2457–2487, 1979.
- [53] A. P. Young. Melting and the vector coulomb gas in two dimensions. *Phys. Rev. B*, 19(4):1855–1866, 1979.
- [54] D.S. Fisher, B.I. Halperin, and R. Morf. Defects in the two-dimensional electron solid and imperfections for melting. *Phys. Rev. B*, 20(11):4692–4712, 1979.
- [55] S.T. Chui. Grain-boundary theory of melting in two dimensions. *Phys. Rev. B*, 28(1):178–194, 1983.
- [56] E.A. Rakhmanov, E.B. Saff, and Y.M. Zhou. Electrons on the sphere. In R.M. Ali, St. Ruscheweyh, and E.B. Saff, editors, *Computational Methods and Function Theory*, pages 111–127. World Scientific, 1995.
- [57] N.W. Ashcroft and J. Lekner. Structure and resistivity of liquid metals. *Phys. Rev.*, 145(1):83–90, 1966.
- [58] J. D. Weeks, D. Chandler, and H. C. Andersen. Role of repulsive forces in determining the equilibrium structure of simple liquids. *J. Chem. Phys.*, 54(12):5237–5247, 1971.

- [59] J.A. Barker and D. Henderson. Perturbation theory and equation of state for fluids. II. A successful theory of liquids. *J. Chem. Phys.*, 47(11):4714–4721, 1967.
- [60] L. Verlet and J. Weis. Perturbation theory for the thermodynamic properties of simple liquids. *Mol. Phys.*, 24(5):1013–1024, 1972.
- [61] L. Verlet. Computer “experiments” on classical fluids. II. Equilibrium correlation functions. *Phys. Rev.*, 165(1):201–212, 1968.
- [62] J.P. Hansen and L. Verlet. Phase transitions of the Lennard-Jones system. *Phys. Rev.*, 184(1):151–161, 1969.
- [63] T.V. Ramakrishnan and M. Yussouff. First-principles order-parameter theory of freezing. *Phys. Rev. B*, 19(5):2775–2794, 1979.
- [64] T.V. Ramakrishnan. Density-wave theory of first-order freezing in two dimensions. *Phys. Rev. Lett.*, 48(8):541–545, 1982.
- [65] J. Vollmer, W. Breymann, and R. Schilling. Number of metastable states of a chain with competing and anharmonic Φ^4 -like interactions. *Phys. Rev. B*, 47(18):11767–11773, 1993.
- [66] M.R. Hoare and J. McInnes. Statistical mechanics and morphology of very small atomic clusters. *Discuss. Faraday Soc.*, 61:12–24, 1976.
- [67] R.A. La Violette and F.H. Stillinger. Enumeration of random packings for atomic substances. *Phys. Rev. B*, 35(11):5446–5452, 1987.
- [68] F.H. Stillinger and T.A. Weber. Dynamics of structural transitions. *Phys. Rev. A*, 28(4):2408–2416, 1983.
- [69] T.A. Weber and F.H. Stillinger. The effect of density on the inherent structure in liquids. *J. Chem. Phys.*, 80(6):2742–2746, 1984.
- [70] J. R. Edmundson. The arrangement of point charges with tetrahedral and octahedral symmetry on the surface of a sphere with minimum coulombic potential energy. *Acta Cryst. A*, 49:648–654, 1993.
- [71] K. Chen, T. Kaplan, and M. Mostoller. Melting in two-dimensional Lennard-Jones systems: observation of a metastable hexatic phase. *Phys. Rev. Lett.*, 74(20):4019–4022, 1995.
- [72] J. Tobochnik and G.V. Chester. Monte Carlo study of melting in two dimensions. *Phys. Rev. B*, 25(11):6778–6798, 1982.
- [73] D. Frenkel and J.P. McTague. Evidence for an orientationally ordered two-dimensional fluid phase from molecular-dynamics calculations. *Phys. Rev. Lett.*, 42(24):1632–1635, 1979.
- [74] K.J. Strandburg, J.A. Zollweg, and G.V. Chester. Bond-angular order in two-dimensional Lennard-Jones and hard-disk systems. *Phys. Rev. B*, 30(5):2755–2759, 1984.
- [75] S. Toxvaerd. Computer simulation of melting in a two-dimensional Lennard-Jones system. *Phys. Rev. A*, 24(5):2735–2742, 1981.
- [76] A.F. Bakker, C. Bruin, and H.J. Hilhorst. Orientational order at the two-dimensional melting transition. *Phys. Rev. Lett.*, 52(6):449–452, 1984.
- [77] F.F. Abraham. The phases of two-dimensional matter, their transitions, and solid-state stability: a perspective via computer simulation of simple atomic systems. *Physics Reports*, 80(5):339–374, 1981.

- [78] P. Bladon and D. Frenkel. Dislocation unbinding in dense two-dimensional crystals. *Phys. Rev. Lett.*, 74(13):2519–2522, 1995.
- [79] P. Ballone, G. Pastore, M. Rovere, and M.P. Tosi. Liquid structure and freezing of the two-dimensional classical electron fluid. *J. Phys. C*, 18:4011–4019, 1985.
- [80] J.Q. Broughton, G.H. Gilmer, and J.D. Weeks. Molecular Dynamics study of melting in two dimensions. Inverse-twelfth-power interaction. *Phys. Rev. B*, 25(7):4651–4669, 1982.
- [81] H.S.M. Coxeter, P. Du Val, and H.T. Flather. *The fifty-nine icosahedra*. Springer, 1982.
- [82] M.J. Wenninger. *Dual models*. Cambridge, 1983.
- [83] M.J. Wenninger. *Polyhedron models*. Cambridge, 1971.
- [84] John Leech. Equilibrium of sets of particles on a sphere. *The mathematical gazette*, XLI:81–90, 1957.
- [85] M.G. Calkin, D. Kiang, and D.A. Tindall. Minimum-energy charge configurations. *Am. J. Phys.*, 55(2):157–158, 1987.
- [86] R. H. Frickel and B. V. Bronk. Symmetries of configurations of charges on a sphere. *Can. J. Chem.*, 66:2161–2165, 1988.
- [87] L. Glasser and A. G. Every. Energies and spacings of points charges on a sphere. *J. Phys. A: Math. Gen.*, 25:2473–2482, 1992.
- [88] J. R. Edmundson. The distribution of point charges on the surface of a sphere. *Acta Cryst. A*, 48:60–69, 1992.
- [89] J.B. Weinrach, K.L. Carter, D.W. Bennett, and H.K. McDowell. Point charge approximations to a spherical charge distribution. *Journal of chemical education*, 67(12):995–999, 1990.
- [90] T. W. Melnyk, O. Knop, and W. R. Smith. Extremal arrangements of points and unit charges on a sphere : equilibrium configurations revisited. *Can. J. Chem.*, 55:1745–1761, 1977.
- [91] B.W. Van de Waal. Equilibrium distributions of repulsive particles in a sphere. *Am. J. Phys.*, 56(7):583–584, 1988.
- [92] A.A. Berezin. Spontaneous symmetry breaking in classical systems. *Am. J. Phys.*, 53(11):1036–1037, 1985.
- [93] J. Tobochnik and P. M. Chapin. Monte Carlo simulation of hard spheres near random closest packing using spherical boundary conditions. *J. Chem. Phys.*, 88(9):5824–5830, 1988.
- [94] M. Goldberg. An improved packing of 33 equal circles on a sphere. *Elem. Math.*, 22:110–112, 1967.
- [95] W. Schreiner and K.W. Kratky. Computer simulation of hard-disc packings with spherical boundary conditions. *J. Chem. Soc., Faraday Trans. 2*, 78:379–389, 1982.
- [96] T. Tarnai and Zs. Gáspár. Multi-symmetric close packings of equal spheres on the spherical surface. *Acta Cryst. A*, 43:612–616, 1987.
- [97] E.A. Rakhmanov, E.B. Saff, and Y.M. Zhou. Minimal discrete energy on the sphere. *Mathematical Research Letters*, 1:647–662, 1994.
- [98] J.M. Voogd, P.M.A. Sloot, and R. v. Dantzig. Comparison of vector and parallel implementations of the Simulated Annealing algorithm. *FGCS*, 11:467–475, 1995.

- [99] P. M. A. Sloot, J. M. Voogd, D de Kanter, and L. O. Hertzberger. Simulated annealing: comparison of vector and parallel implementations. Technical Report CS-93-06, University of Amsterdam, Amsterdam, The Netherlands, 1993.
- [100] J. M. Voogd and P. M. A. Sloot. Simulated annealing on H.P.C. systems: Applied to crystallization with spherical boundary conditions. In A. Verbraeck and E. J. H. Kerckhoffs, editors, *European Simulation Symposium 1993*, pages 371–376, Delft, The Netherlands, 1993. Society for Computer Simulation International.
- [101] J. M. Voogd, R. van Dantzig, and P. M. A. Sloot. Two-dimensional crystallisation on spherical surfaces. In R. Gruber and M. Tomassini, editors, *Proceedings of the 6th joint EPS-APS International Conference on Physics Computing*, pages 463–466, Lugano, Switzerland, 1994. European Physical Society, American Physical Society.
- [102] J. M. Voogd, P. M. A. Sloot, and R. van Dantzig. Simulated annealing for N-body systems. In W. Gentzsch and U. Harms, editors, *lecture notes in computer science 796, high-performance computing and networking*, pages 293–298, Munich, Germany, 1994. Springer Verlag.
- [103] P.M.A. Sloot, A. Schoneveld, J.F. de Ronde, and J.A. Kaandorp. Large-scale simulations of complex systems, part I: conceptual framework. Technical Report Working Paper 97-07-070, Santa Fe Institute, 1997.
- [104] S. Kirkpatrick, C. D. Gelatt jr., and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.
- [105] D.E. Goldberg. *Genetic Algorithms in search, optimization and machine learning*. Addison-Wesley, 1989.
- [106] J. Hertz, A. Krogh, and R. Palmer. *Introduction to the theory of neural computing*. Addison-Wesley, 1994.
- [107] S. Chen, Z. Wang, X. Shan, and G. Doolen. Lattice Boltzmann computational fluid dynamics in three dimensions. *Journal of Statistical Physics*, 68(3/4):379–400, 1992.
- [108] P. J. M. van Laarhoven and E. H. L. Aarts. *Simulated Annealing : Theory and Applications*. Kluwer Academic Publishers, 1987. (and references therein).
- [109] R.O. Jones. Aluminum and gallium clusters - a comparative study using Simulated Annealing. *Z. Phys. D*, 26:23–27, 1993.
- [110] L.T. Wille. Searching potential energy surfaces by Simulated Annealing. *Nature*, 324:46–48, 1986.
- [111] L.T. Wille and J. Vennik. Electrostatic energy minimisation by Simulated Annealing. *J. Phys. A: Math. Gen.*, 18:L1113–L1117, 1985.
- [112] G. Seifert and R. O. Jones. Structure of phosphorus clusters by Simulated Annealing. *Z. Phys. D*, 26:349–351, 1993.
- [113] R.O. Jones. Simulated annealing study of neutral and charged clusters : Al(n) and Ga(n). *J. Chem. Phys.*, 99(2):1194–1205, 1993.
- [114] R. Biswas and D.R. Hamann. Simulated annealing of silicon atom clusters in Langevin Molecular Dynamics. *Phys. Rev. B*, 34(2):895–901, 1986.
- [115] J. B. Cole. The statistical mechanics of image recovery and pattern recognition. *Am. J. Phys.*, 59(9):839–842, 1991.
- [116] G.C. Fox and S.W. Otto. Concurrent computation and the theory of complex systems. In M.T. Heath, editor, *Hypercube Multiprocessors*, 1986.

- [117] J.F. de Ronde. *Mapping in high performance computing. A case study on finite element simulation*. PhD thesis, University of Amsterdam, Amsterdam, The Netherlands, 1997.
- [118] A. Schoneveld, J.F. de Ronde, and P.M.A. Sloot. On the complexity of task allocation. *Complexity*, 3(2):52–60, 1997.
- [119] H.W. Carter. Computer-aided design of integrated circuits. *Computer*, 19(4), 1986.
- [120] G.B. Sorkin. Simulated Annealing on fractals: Theoretical analysis and relevance for combinatorial optimisation. In W.J. Dally, editor, *Advanced Research in VLSI*, Proceedings of the 6th MIT Conference, pages 331–351, 1990.
- [121] E.D. Weinberger. Correlated and uncorrelated fitness landscapes and how to tell the difference. *Biological Cybernetics*, 63:325, 1990.
- [122] P.F. Stadler and R. Happel. Correlation structure of the landscape of the graph-bipartitioning problem. *J. Phys. A: Math. Gen.*, 25:3103–3110, 1992.
- [123] T. Erber and G.M. Hockney. Equilibrium configurations on N equal charges on a sphere. Technical Report FERMILAB-PUB-91/222-T, Fermi National Accelerator Laboratory, 1991.
- [124] N. Metropolis, A. W. Rosenbruth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *J. Chem. Phys.*, 21(6):1087–1092, 1953.
- [125] A.R. Leach. *Molecular modelling, principles and applications*. Addison Wesley Longman, 1996.
- [126] H. Ding, N. Karasawa, and W. A. Goddard III. Optimal spline cutoffs for Coulomb and van der Waals interactions. *Chem. Phys. Lett.*, 193(1,2,3):197–201, 1992.
- [127] P.M.A. Sloot. Modelling and simulation. Technical Report CS-94-12, University of Amsterdam, Amsterdam, The Netherlands, 1994.
- [128] L. Ingber. Very fast simulated re-annealing. *Math. Comput. Modelling*, 12(8):967–973, 1989.
- [129] D.A. Burer. Parallel Molecular Dynamics: Applied to crystallisation on a sphere. Master's thesis, University of Amsterdam, Amsterdam, The Netherlands, 1997.
- [130] P.F. Stadler. Correlation in landscapes of combinatorial optimization problems. *Europhys. Lett.*, 20(6):479–482, 1992.
- [131] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical recipes in C*. Cambridge University Press, 1992. Second edition.
- [132] K. Hwang. *Advanced computer architectures, parallelism, scalability and programability*. Series in computer engineering. Mc Graw-Hill, 1993.
- [133] P.M.A. Sloot. Modelling for parallel simulation: Possibilities and pitfalls. *Invited lecture Eurosim '95, Simulation Congress, Conference Proceedings, (Amsterdam, Netherlands)*, pages 29–44, 1995.
- [134] A.W. van Halderen, B.J. Overeinder, P.M.A. Sloot, R. van Dantzig, D.H.J. Epema, and M. Livny. Hierarchical resource management in the polder metacomputing initiative. to appear in *Parallel Computing*.
- [135] D. Abramson, I. Foster, J. Giddy, A. Lewis, R. Susic, R. Sutherst, and N. White. The nimrod computational workbench: A case study in desktop metacomputing. In *Australian Computer Science Conference (ACSC 97)*, Sydney Macquarie University, 1997.

- [136] L. Ingber. Simulated annealing: Practice versus theory. *J. Math. Comput. Modelling*, 18(11):29–58, 1993.
- [137] R. Azencott. *Simulated Annealing: Parallelization Techniques*. Wiley & Sons, 1992.
- [138] Y. Kim and M. Kim. A step-wise-overlapped parallel annealing algorithm on a message-passing multiprocessor system. *Concurrency: practice and experience*, 2(2):123–148, 1990.
- [139] A. Sohn. Parallel N-ary speculative computing of simulated annealing. *IEEE trans. on Par. and Distr. Systems*, 6(10):997–1005, 1995.
- [140] D. M. Jones and J. M. Goodfellow. Parallelization strategies for molecular simulation using the Monte Carlo algorithm. *J. of Comp. Chem.*, 14(2):127–137, 1993.
- [141] K. Esselink, L.D.J.C. Loyens, and B. Smit. Parallel Monte Carlo simulations. *Phys. Rev. E*, 51(2):1560–1568, 1995.
- [142] D. R. Greening. Parallel Simulated Annealing techniques. *Physica D*, 42:293–306, 1990.
- [143] A. Sohn. Parallel speculative computation of Simulated Annealing. In *1994 International conference on parallel processing*, 1994.
- [144] E. Aarts and J. Korst. *Simulated Annealing and Boltzmann machines*. Wiley-Interscience Series in discrete mathematics and optimization, 1989.
- [145] M. Livny, J. Basney, R. Raman, and T. Tannenbaum. Mechanisms for High Throughput Computing. *SPEEDUP Journal*, 11(1), 1997.
- [146] D.H.J. Epema, M. Livny, R. van Dantzig, X. Evers, and J. Prune. A worldwide flock of condors : Load sharing among workstation clusters. *FGCS*, 12:54–65, 1996.
- [147] E. H. L. Aarts, F. M. J. de Bont, E. H. E. Habers, and P. J. M. van Laarhoven. Parallel implementations of the statistical cooling algorithm. *North Holland Integration, the VLSI journal*, 4:209–238, 1986.
- [148] A. ter Laak, L. O. Hertzberger, and P. M. A. Sloot. Nonconvex continuous optimization experiments on a transputer system. In A. R. Allen, editor, *Transputer Systems - Ongoing Research*, page 251, Amsterdam, 1992. IOS Press.
- [149] A.G. Hoekstra, P.M.A. Sloot, F. van der Linden, M. van Muiswinkel, J.J.J. Vesseur, and L.O. Hertzberger. Native and generic parallel programming environments on a transputer and a PowerPC platform. *Concurrency: practise and experience*, 8(1):19–46, 1996.
- [150] A.G. Hoekstra. Private communications.
- [151] P. Hut, J.M. Arnold, J. Makino, S.L.W. McMillan, and T.L. Sterling. GRAPE-6: A petaflops prototype. In *The 1997 Petaflops Algorithms Workshop (PAL'97)*, 1997.
- [152] Y. Komeiji, M. Uebayasi, R. Takata, A. Shimizu, K. Itsukashi, and M. Taiji. Fast and accurate Molecular Dynamics simulation of a protein using a special-purpose computer. *J. Comp. Chem.*, 18:1546–1563, 1997.
- [153] J. K. Salmon. *Parallel hierarchical N-Body Methods*. PhD thesis, California institute of technology, Pasadena USA, 1991.
- [154] Y. Imry. Finite-size rounding of a first-order phase transition. *Phys. Rev. B*, 21(5):2042–2043, 1980.
- [155] M.E. Fisher and A.N. Berker. Scaling for first-order phase transitions in thermodynamic and finite systems. *Phys. Rev. B*, 26:2507–2513, 1982.

- [156] P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality. *Phys. Rev. A*, 38(1):364–374, 1988.
- [157] R. Lipowsky. The conformation of membranes. *Nature*, 349:475–481, 1991.
- [158] L.L. Whyte. Unique arrangements of points on a sphere. *Am. Math. Month.*, 59:606–611, 1952.