Jammed Packings of Hard Particles

Aleksandar Donev

A Dissertation

Presented to the Faculty

of Princeton University

in Candidacy for the Degree

of Doctor of Philosophy

Recommended for Acceptance

by the Program in

Applied and Computational Mathematics

September, 2006

Abstract

This dissertation describes an investigation of jammed packings of frictionless hard particles, including the computer generation of (nearly) jammed packings, the development of mathematical criteria and algorithms to verify jamming, and computational and experimental studies of disordered and ordered hard-sphere and hard-ellipsoid packings.

In the first part of this dissertation a mathematical framework for understanding jamming in packings of hard particles is developed. Algorithms to model hard-particle systems, and in particular, a collision-driven molecular dynamics algorithm for the simulation of dense packings of hard spheres, ellipsoids and superellipsoids, are designed. This algorithm is used to implement a generalization of the Lubachevsky-Stillinger algorithm to generate disordered packings of hard spheres and hard ellipsoids. It is found that the density and average contact number of the random packings rises sharply, but continuously, as asphericity is introduced, leading to hypostatic packings much denser than well-known random sphere packings. A mathematical theory of jamming for packings of spherical and nonspherical particles, as well as algorithms to test whether a packing is (nearly) jammed are developed, verifying that our packings are jammed. A molecular-dynamics algorithm to calculate the (non-equilibrium) free energy of nearly jammed packings of hard particles is designed and implemented.

In the second part of this dissertation the properties of disordered and ordered packings of hard particles are studied. Investigated are correlations, including short-ranged order in the pair-correlation function, as well as long-ranged density fluctuations in the structure factor, for hard sphere packings in both three and higher dimensions. An unusual multitude of near contacts persistent with dimensionality, as well as a decorrelation for distances beyond contact as dimension increases, are found. Comparisons find good agreement between

computational and experimental results for packing of hard ellipsoids in finite containers. The densest known ordered packing of hard ellipsoids is discovered using molecular dynamics. Finally, the thermodynamics of dense systems of hard-particles is investigated. The phase-diagram of hard rectangles of aspect ratio two (dominos) is determined. Finally, it is demonstrated that there is no ideal glass transition for binary systems of hard disks.

Acknowledgements

I am indebted to a number of people who have made my stay at Princeton both productive and enjoyable. I thank my family and friends for their support and love, without which I could not have finished this dissertation. Thanks to my partner Christian for loving me back and for constantly making fun of my work.

I am sincerely grateful and deeply indebted to my advisor Dr. Sal Torquato, and to Dr. Frank Stillinger, for their continued intellectual and professional guidance, and their sincere care for my research and my career development. My work would not have been possible without close collaboration with Dr. Robert Connelly (Cornell University) and Dr. Paul Chaikin.

Many professors, postdocs, and students have helped me or collaborated with me; thanks to Dr. Robert Vanderbei, Weining Man, Josh Burton, Monica Skoge, and all the students in the Torquato research group. Thanks also to colleagues around the world that have shared their advice or codes with me, and to the open-source community for making computing a shared resource. My stay at Princeton University would not have been the same without the LGBT student center, the Program in Theatre in Dance, the people at PRISM, and of course, the students, faculty and staff in PACM.

Work presented in this Dissertation was supported in part by the Petroleum Research Fund under Grant No. 36967-AC9, and by the National Science Foundation under Grant Nos. DMR-0213706 and DMS-0312067.

Contents

Chapter 1

Introduction and Overview

This Dissertation presents the results of collaborative computational, theoretical, and experimental studies of dense packings of frictionless impenetrable objects, henceforth, *hard particles*, in a d-dimensional Euclidean space $\mathcal{R}^d$, where the particles occupy a certain (*covering* or *volume*) *fraction* (or *density*) ϕ of the total volume. We focus extensively on low dimensions, especially two ($d = 2$) and three ($d = 3$) dimensions, on smooth convex particle shapes, and packings where all particles have the same or very similar shape. Furthermore, our primary interest is understanding *jammed* packings, i.e., packings where the particles are packed tightly and cannot move freely due to the impenetrability constraints.

Although the packings we investigate can be modeled purely within classical geometry, our investigation focuses on hard-particle packings relevant to material science. Packings of hard particles interacting only with infinite repulsive pairwise forces on contact are applicable as models of complex many-body systems because repulsive interactions are the primary factor in determining their structure. Hard-particle packings are therefore widely used as simple models for granular materials [1, 2, 3], glasses [4], liquids [5], and other random media[6], to mention a few examples. Furthermore, hard-particle packings, and especially hard-sphere packings, have inspired mathematicians and been the source of numerous challenging (many still open) theoretical problems [7, 8, 9, 10]. Of particular interest are *jammed* hard-particle packings, where particles become trapped and the packing gains mechanical rigidity and resists further densification and/or flow. There are still many open theoretical problems in the field, such as the precise identification and quantitative description of the maximally random jammed (MRJ) state, the identification of packing structures with extremal properties (for example, the lowest or highest density jammed packings), understanding the thermodynamics of model glass formers, the quantification of disorder via order metrics, and others. Additionally, the results we obtain can have significant practical applications such as production of advanced ceramic materials and rocket fuels, developing constitutive equations for the mechanical behavior of granular media, understanding why eggs are not spherical, and others.

This dissertation is the product of a long-term research program that consists of five main areas, all of which are represented in subsequent chapters:

- Developing algorithms to generate ordered (crystalline) and disordered (glassy) packings of spherical and nonspherical particles.
- Understanding the phenomenon of jamming in hard-particle packings and developing algorithms to test whether a packing is jammed.
- Statistical characterization of the generated jammed packings, including density, correlation functions, and order metrics.
- Understanding the thermodynamics of nearly jammed packings of hard particles and their relevance to the glass transition.

Material presented in this Dissertation has previously appeared in the following peer-reviewed publications:

1. A. Donev, S. Torquato, F. H. Stillinger, and R. Connelly. A Linear Programming Algorithm to Test for Jamming in Hard-Sphere Packings. *J. Comp. Phys.*, 197(1):139–166, 2004.

2. A. Donev, S. Torquato, F. H. Stillinger, and R. Connelly. Jamming in Hard Sphere and Disk Packings. *J. App. Phys.*, 95(3):989, 2004.
3. S. Torquato, A. Donev, and F. H. Stillinger. Breakdown of Elasticity Theory for Jammed Hard-Particle Packings: Conical Nonlinear Constitutive Theory. *Int. J. Solids Structures*, 40(25):7143 – 7153, 2003.
4. A. Donev, I. Cisse, D. Sachs, E. A. Variano, F. H. Stillinger, R. Connelly, S. Torquato, and P. M. Chaikin. Improving the Density of Jammed Disordered Packings using Ellipsoids. *Science*, 303:990–993, 2004.
5. A. Donev, P. M. Chaikin, F. H. Stillinger, and S. Torquato. Unusually Dense Crystal Packings of Ellipsoids. *Phys. Rev. Lett.*, 92:255506, 2004.
6. A. Donev, S. Torquato, F. H. Stillinger, and R. Connelly. Comment on "Jamming at zero temperature and zero applied stress: The epitome of disorder". *Phys. Rev. E*, 70:043301, 2004.
7. A. Donev, S. Torquato, and F. H. Stillinger. Neighbor List Collision-Driven Molecular Dynamics Simulation for Nonspherical Particles: I. Algorithmic Details II. Applications to Ellipses and Ellipsoids. *J. Comp. Phys.*, 202(2):737–764, 765–793, 2005.
8. A. Donev, S. Torquato, and F. H. Stillinger. Pair Correlation Function Characteristics of Nearly Jammed Disordered and Ordered Hard-Sphere Packings. *Phys. Rev. E*, 71:011105, 2005.
9. W. Man, A. Donev, F. H. Stillinger, M. Sullivan, William B. Russel, D. Heeger, S. Inati, S. Torquato, and P. M. Chaikin. Experiments on Random Packing of Ellipsoids. *Phys. Rev. Lett.*, 94:198001, 2005.
10. A. Donev, S. Torquato, and F. H. Stillinger. Unexpected Density Fluctuations in Jammed Disordered Sphere Packings. *Phys. Rev. Lett.*, 95(9):090604, 2005.
11. A. Donev, J. Burton, F. H. Stillinger, and S. Torquato. Tetratic Order in the Phase Behavior of a Hard-Rectangle System. *Phys. Rev. B*, 73:054109, 2006.
12. A. Donev, F. H. Stillinger, and S. Torquato. Do Binary Hard Disks Exhibit an Ideal Glass Transition? *Phys. Rev. Lett.*, 96:225502, 2006.
13. M. Skoge, A. Donev, F. H. Stillinger, and S. Torquato. Packing Hyperspheres in High-Dimensional Euclidean Spaces. Submitted to *Phys. Rev. E*, 2006.
14. A. Donev, F. H. Stillinger, and S. Torquato. Calculating the Free Energy of Nearly Jammed Hard-Particle Packings Using Molecular Dynamics. Submitted to *J. Comp. Phys*, 2006.
15. A. Donev, R. Connelly, F. H. Stillinger, and S. Torquato. Hypostatic Jammed Packings of Hard Ellipsoids. In preparation, 2006.
16. A. Donev, F. H. Stillinger, and S. Torquato. Configurational Entropy of Binary Hard-Disk Glasses. In preparation, 2006.

In this chapter we present some background material for subsequent chapters, aiming at developing an intuitive (physical) understanding of the essential underlying concepts and methodology before presenting a more formal approach. The reader is referred to Ref. [6] for background on the thermodynamics and statistical characterization of particle media. We also review the main findings of subsequent papers, pointing to the publications where the results were first published, available in electronic form at `http://cherrypit.princeton.edu/donev` (the webpage also contains additional resources related to this work, such as source codes, animations, etc.). We review some of the notational conventions we use in subsequent chapters at the end of this chapter.

1.1 Generating Packings using Molecular Dynamics

In this work we focus on computationally-generated (nearly) jammed packings of hard spherical and non-spherical particles. In this section we give a high-level overview of the molecular dynamics algorithm used to generate hard-particle packings and study their thermodynamic and kinetic properties.

The hard-particle interaction potential is singular, being zero if the two particles do not overlap, and infinite otherwise. Classical molecular-dynamics methods are not suitable for modeling such singular potentials, ad we have developed a collision-driven (event-driven) *molecular dynamics* (MD) algorithm for frictionless

perfectly-elastic hard particles [11]. The particles move freely inbetween elastic collisions with other particles, when they simply bounce away immediately, exchanging collisional momentum (impulse) over an infinitesimally short time period. The goal of the MD algorithm is to efficiently predict and process the sequence of elastic collisions between the particles given their initial positions and velocities.

We have used our collision-driven MD algorithm to implement a generalization of the *Lubachevsky-Stillinger* (LS) sphere packing algorithm [12, 13]. The generalization handles nonspherical particles (in particular, our implementation handles spherical, ellipsoidal, or superellipsoidal particles) and provides numerous improvements to the numerical efficiency of the algorithm and also increases the generality of the algorithm; however, the fundamental ideas are the same as in original Lubachevsky-Stillinger algorithm. The method is a hard-particle molecular dynamics algorithm for producing dense disordered as well as ordered packings. Small particles are randomly distributed and randomly oriented in a box with periodic boundary conditions and without any overlap. For this purpose one can use *random sequential addition* (RSA) [14], where particles are added to the simulation region one by one, ensuring that they do not overlap any of the previously-added particles. After the initial particle positions are determined, the particles are given (linear and angular) velocities and their motion followed as they collide elastically and also expand uniformly (i.e., preserving their shape) at a certain *growth* or *expansion rate* γ, using the MD algorithm.

Asymptotically, as the density increases, a *jammed packing* with a diverging collision rate and (locally) maximal *jamming density* ϕ_J is reached. Intuitively, a jammed (compactly packed, mechanically stable) packing is one where the particles are locked in their positions despite thermal agitation (shaking) and potentially boundary deformations (external loading). Depending on the boundary conditions and whether collective particle rearrangements take place, one can define different jamming categories, organized hierarchically into *local*, *collective* and *strict jamming* in Ref. [15]. The algorithm never reaches the true jammed state, and the particles have some *free volume* to rattle which shrinks as the reduced *pressure* $p = PV/NkT$ diverges. Given sufficient computational resources, our careful numerical implementation is able to reach the jamming point almost within full numerical precision, corresponding to pressures $p > 10^{12}$.

As the jamming point is approached, each particle i collides repeatedly only with a small set of *neighbors* $\mathcal{N}(i)$. In the jamming limit, these neighboring particles touch to form the *contact network* of the packing, exerting compressive forces on each other but not being able to move despite thermal agitation (shaking). These forces form a *force network*, often called *force chains* in the physics literature [16] due to the distinctive highly heterogeneous character of the force spatial distribution, concentrated along chains. The net time-averaged force exerted on each particle is zero, since the particle does not move. By averaging the total exchanged momentum between each pair of recently collided particles, one obtains a measure of the contact force between pairs of particles (with some arbitrary proportionality constant). An example of a jammed packing and its contact network is illustrated in two dimensions for binary hard disks in Fig. 1.1.

Note that the neighbors $\mathcal{N}(i)$ included in the contact network at the jamming point are only those neighbors which touch the particle i in the jamming limit. Often, a geometric definition of *first neighbors* is used instead, which considers all particles within a certain cutoff distance to be neighbors. Such a definition is somewhat arbitrary and is not unique because of its dependence on the cutoff distance. A better definition of geometric neighbors is to use the *Voronoi tessellation* of the packing. For a point pattern (or a monodisperse sphere packing), the Voronoi cell of a given point i is defined as the set of points that are closer (in some distance metric) to point i than to any other point. If the Euclidean distance metric is used, the cells are polyhedra and they tessellate, i.e., prefectly tile (partition), Euclidean space into a polyhedral cell complex known as the Voronoi tessellation. Particles that share faces of their respective Voronoi cells are called *Delaunay neighbours* and the set of Delaunay neighbors triangulates space into simplices (triangles in two dimensions, tetrahedra in three dimensions, etc.). Note that the contact network of the packing is a subset of the Delaunay triangulation, since every pair of particles that touches is also a pair of Delaunay neighbors. However, not all Delaunay neighbors are necessarily touching. Note that for polydisperse sphere packings a slight modification of the ordinary Euclidean distance metric can be used to define a Voronoi tessellation that is a polyhedral cell complex [17]. For nonspherical particles, however, it is not trivial to generalize the idea of a Voronoi tessellation [18]; obvious generalizations lead to Voronoi cells that either do not tile space or they are not connected.

The LS packing protocol we have described has several important tuning parameters that affect the resulting packings. In particular, the initial configuration, the distribution of velocities, as well as the particle growth rate, determine what kind of jammed packing is produced. We will discuss these issues

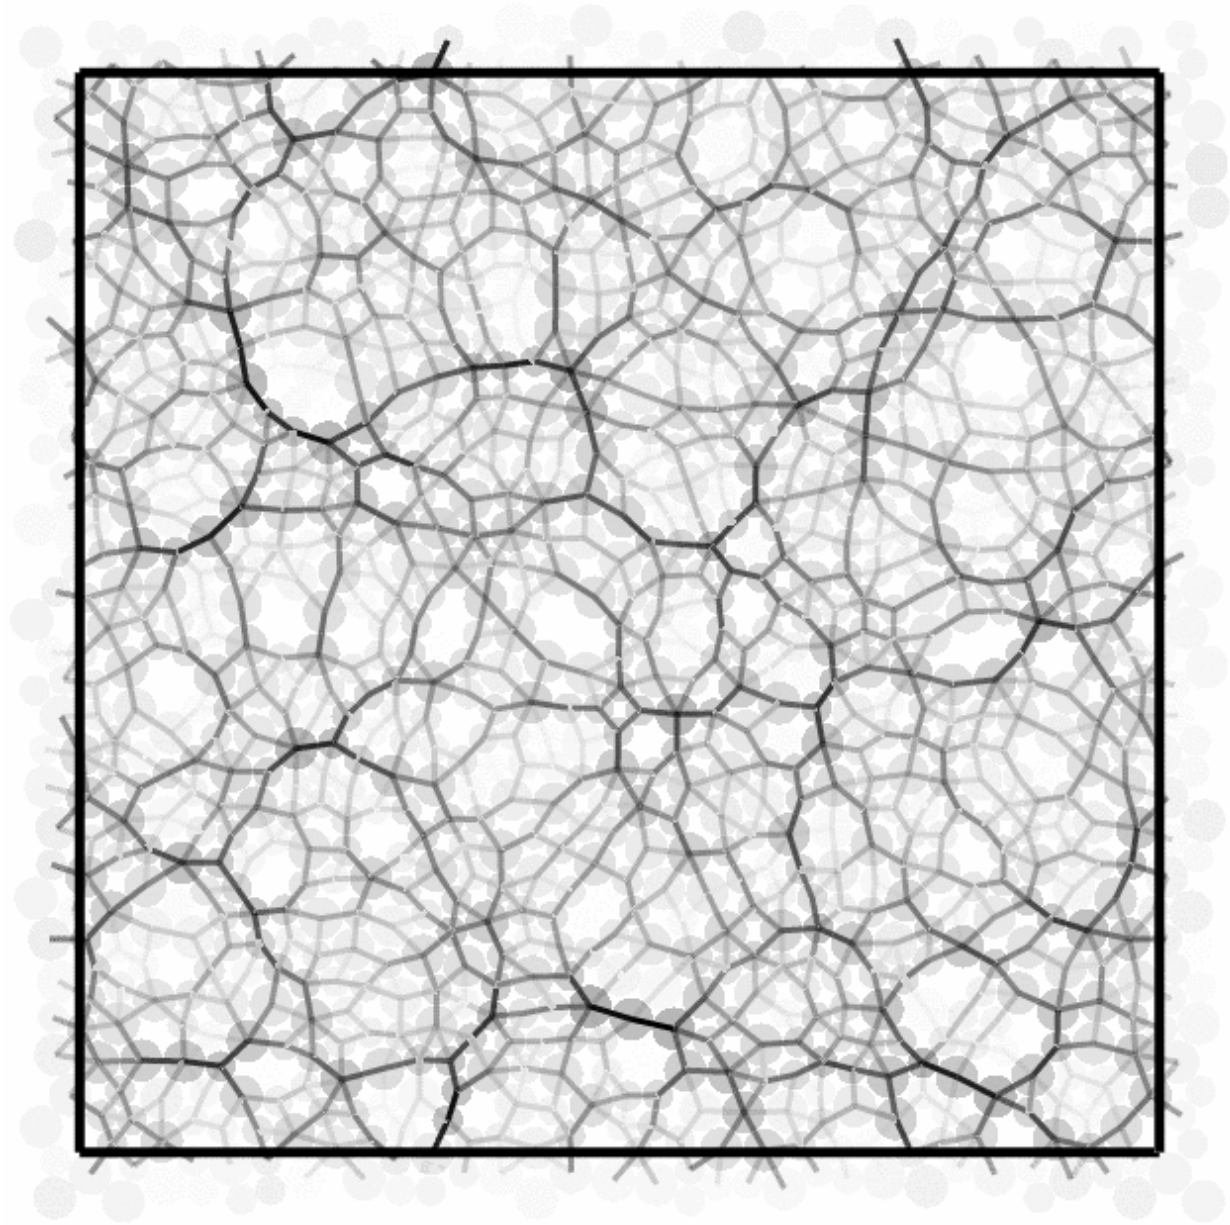

Figure 1.1: A network of interparticle forces (darker lines indicate stronger forces) in a bidisperse jammed packing of hard disks, where the larger disks are 1.4 times larger than the small disks. The contact network was obtained by averaging the total exchanged momentum between colliding particles over a long period of time during the final stages of the Lubachevsky-Stillinger packing algorithm. Darker particles collide more frequently then lighter ones.

shortly, however, first we review some alternative algorithms to generate dense packings of hard particles, that have appeared in the physics and engineering literature.

1.1.1 Other Packing Algorithms

Many of the tools we develop in this work are relevant to the implementation and analysis of other algorithms for generating dense packings of hard particles as well. In particular, the algorithms we develop in Chapters 2 and 3 can be used to implement any kind of molecular dynamics (MD) or Monte Carlo (MC) algorithm for either hard or soft nonspherical particles. The analysis of hard particle packings we perform focuses on the final (static) jammed packing, and is therefore applicable to any packing regardless of the protocol used to produce it. It is therefore useful to review some other packing algorithms (for spherical particles, unless indicated otherwise) that have appeared in various papers. We give a characteristic reference first followed by a short summary of the algorithm:

[19] Many groups have used a *sedimentation* packing algorithm, in which spheres are dropped in a gravitational field onto an initial disordered substrate, and each sphere rolls on top of other spheres downhill until it forms three contacts with spheres below it, which puts in a stable configuration. These kinds of algorithms are thought to produce a "random-loose packing", in the sense that these packings are less dense (approximately 0.6) and are thought to be representative of granular materials when shaking is not present (i.e., gravity dominates the interparticle forces). The produced packings have strong vertical anisotropy, and it has been attempted to correct for this by using *seed growth* algorithms, in which one adds particles one-by-one to an initial small seed cluster [20]. Ref. [19] combines sedimentation with rearrangement (shaking), similarly to Ref. [21], in order to achieve states transitioning from random "loose" to "close-packed" structures.

[22] The algorithm constructs the Voronoi tessellation of a dilute disk packing, and then moves each disk to the center of the largest inscribed circle in the respective Voronoi polyhedron. The common disk diameter is increased as much as non-overlap allows, and the process is repeated many times. They believe their algorithm produces random loose packings (RLP).

[21, 23] In the Jodrey-Tory (JT) algorithm the spheres here have an outer d^{out} and an inner d^{in} diameter, the inner being the physically allowed one of the shortest pair distance, and the outer representing a repulsive core used to densify the packing. The outer diameter is slowly reduced during the simulation until it meets the inner, at which point a valid homogeneous isotropic dense packing is produced.
At each step, outer overlaps are eliminated by displacing the spheres. In the original algorithm, a complicated geometrical scheme was proposed on how to displace the spheres, but later a *force-biased* modification was developed and implemented efficiently. It is described well in Ref. [24]. In this algorithm, a central repulsive force of magnitude

$$f_{ij} = d_i^{\text{out}} d_j^{\text{out}} \left[\frac{4r_{ij}^2}{\left(d_i^{\text{out}} + d_j^{\text{out}}\right)^2} - 1 \right]$$

is imagined to act between the spheres i and j (this form is suitable for polydisperse mixtures as well), if their outer shells overlap. A simple gradient descent step is taken to move the spheres to a new position,

$$\mathbf{r}_i \leftarrow \mathbf{r}_i + \frac{\alpha}{d_i^{\text{in}}} \sum_j \mathbf{f}_{ij}$$

and the process is repeated anew. The force-biased JT algorithm is generalized to binary spheres and combined with a vibration phase in which random displacements are given to the spheres in Ref. [25] and is also used in Refs. [24, 26]. The algorithm has been generalized to non-spherical particles as well [27], and is essentially a coordinate-descent energy minimization algorithm where the interaction potential is continuously updated by changing the sphere inner and outer diameters.

[27] This paper describes a generalization of the force-biased JT algorithm for producing packings of spherocylinders. For spherocylinders particle overlap is easy to determine by looking at the shortest distance

between the lines of symmetry of the spherocylinders and comparing it to the diameter D. Repulsive forces proportional to the overlap are constructed for contacting particles and then particles are translated and rotated in the direction of the resultant force in an attempt to eliminate overlap. The system is compressed uniformly by a small amount, which induces new overlap and the whole process is repeated many times until overlaps cannot be eliminated.

[28] The Zinchenko algorithm maintains an isostatic contact network and grows the sphere at a uniform rate, maintaining this list of contacts. Maintaining a disordered contact network suppresses crystallization. The initial contact network is formed by placing sticky spheres randomly and growing the spheres in a sticky way, so that contacts are frozen as soon as they are made. When an extra contact is formed during the growth, a self-stress is determined which resists the further growth of the particles (this is unique up to scaling). The negative reaction forces correspond to contacts that can be broken in order to restore isostaticity. When all the contact forces are negative, then no contact can be broken and a final stat is reached. The algorithm produces disordered jammed packings with a very uniform density of about $0.637 \approx \frac{2}{\pi}$, and is difficult to implement efficiently.

[29] The local minima of an inverse-power repulsive potential is used to find the inherent structures of a hard-sphere liquid. The order of the limits is that for a fixed number of particles N the power in the interaction potential is increased to infinity first, and then N is taken to the thermodynamic limit. The spectrum of the Hessian matrix is analysed as a function of the temperature and melting is related to the existence of negative eigenvalues.

[30] A search for the densest packings of N disks on a sphere is performed using repulsive potential energy minimization. The repulsive energy is inverse power (as in Ref. [29])

$$\varepsilon = \sum_{ij} \left(\frac{c}{r_{ij}} \right)^p$$

for very large exponents p. In the limit $p \to \infty$ the shortest distance dominates the energy, so the minimum-energy configuration in fact maximizes the minimum distance between N points on a sphere (i.e., the densest packing). Random starting points are used in a local search for the global optimum. A combined gradient descent with Newton's method is used for the minimization, and the exponent is increased by doubling from 80 to $1,310,720$ so that high numerical accuracy is achieved. An important aspect of this work which is applicable to any similar algorithm is a final "exactification" stage. In this last stage, the interparticle gaps on the formed contact network are set to 0 to high numerical accuracy.

[31] Infinitely rapid quenching to zero temperature of Hertzian soft spheres is used to produce jammed binary hard-sphere packings. The quench is really a simple descent energy minimization, in this case using conjugate gradient methods. The minimal energy is zero if there is no overlap, otherwise it is positive. If the energy minimum is overlap-free, the sphere diameter is increased and the process repeated, until a jammed state is reached in which there are very small overlaps between the particles forming the contact network of the packing.

[32] Molecular dynamics (MD) with both strongly repulsive central and tangential friction contact forces is used to produce packings of frictional hard spheres in equilibrium under a gravitational field. Since tangential friction forces are history-dependent, MD is more suitable here then energy minimization. It is observed that frictionless packings are generically isostatic, but frictional packings transition from coordination number 6 to about 4.5 (i.e. hyperstatic) smoothly as friction coefficient is increased.

[33] The authors report a full-fledged discrete element method simulation of $10,000$ deformable spheres interacting with Hertz-Mindlin contact forces. They obtain a packing density of about 0.634(4) for frictionless packings. The parameters chosen are an estimate of real-life parameters for certain granular materials. They find that the coordination number goes to 4 as the applied compressive stress vanishes. They also analyse the force chains in the packing.

1.1.1.1 Why Lubachevsky-Stillinger?

Given the wealth of algorithms to produce dense packings of hard particles, why choose the LS algorithm? Before answering this question, it is worth grouping the algorithms described above into categories:

Serial Deposition These algorithms [19, 20, 34, 35] add particles to a system one by one, establishing contacts between the new particle and as many existing particles as possible, producing inhomogeneous anisotropic packings.

Overlap Elimination In these algorithms [24, 25, 23, 21, 36, 27] one tries to eliminate geometric overlaps between the particles by displacing the particles one by one, in the spirit of single-move Monte Carlo algorithms, until a point where overlaps can no longer be removed (if increasing the density) or the point when overlaps are first removed (if lowering the density). The particle displacements can either be random or force-biased.

Energy Minimization In these algorithms [29, 30, 31] one replaces the hard spheres with deformable (soft) spheres and then looks for potential energy minima (inherent structures). This can either be done for an interaction potential that is very stiff and mimics the hard-sphere potential, or it can be used in order to eliminate overlaps by using collective particle motions instead of using single-particle motions.

Molecular Dynamics In these algorithms [12, 13, 37, 33, 32] molecular dynamics on either hard or soft particles is used to density the system, either by increasing the density or by using dissipative dynamics in a gravitational field.

Contact Network Based The only packing algorithm in this category that we are familiar is the Zinchenko algorithm [28], however, this class of algorithms deserves further investigation. The idea is, as in serial deposition algorithms, to build the packing by maintaining a contact network, however, unlike in serial deposition, this contact network involves all particles from the start, and is updated globally to maintain the contacts in question.

We are primarily interested in generating large homogeneous and isotropic systems that are jammed. As we will demonstrate in this work, jamming is a collective phenomenon and the jamming condition is a very strict condition. Therefore packing algorithms that are not designed to produce jammed packings, such as serial deposition algorithms, will typically not produce jammed packings. Single-particle based algorithms, such as overlap elimination by displacing single particles, will easily get stuck in configurations that are not jammed, since finding unjamming motions requires collective displacements. This is demonstrated by the fact that the packings we produce using the LS algorithm are significantly denser that have previously been obtained for ellipsoids using RSA, sedimentation, or shaking (MC-like) packing protocols [14, 38, 34, 35]. The Zinchenko algorithm itself is very complicated and relies on assumptions that we will show to be violated for large disordered packings of ellipsoids.

This elimination leaves us with energy-minimization or molecular dynamics algorithms. We will demonstrate in this work that energy minimization algorithms are indeed appropriate for generating jammed packings, in that under appropriate conditions they guarantee jamming for the final packings. However, these algorithms are typically difficult to implement efficiently and also they are rather unphysical. Molecular dynamics on the other hand is not only a good way to generate packings, since it too produces jammed packings, but it is also very physical and can be used to study thermodynamics and kinetics in hard-particle systems. Additionally, we will see that MD can be implemented very efficiently for hard particles, where the only particle interactions are via particle collisions. The high precision of such collision-driven MD enables us to reach previously unavailable high densities and to produce tightly jammed ellipsoid packings with several thousand particles. In the future, algorithms based on energy minimization should also be implemented and compared to the LS algorithm, especially fo4r nonspherical particles.

1.2 Disordered Jammed Packings of Hard Spheres

The concept of jamming will be defined precisely in subsequent chapters and analyzed in extensive detail. Another concept that is central to this work, but is not yet well-understood, is that of a *disordered*, *random*

or *amorphous* packing (we will use all of these terms interchangeably). The problem of random packing is among the oldest in science and relates to the ancient (economically important) problem of how much grain a barrel can hold. If the particles are not carefully arranged to achieve the best packing density, but are rather "randomly" placed in a container and shaken, we expect to obtain a density lower than that of the best ordered or crystal packing (having a density of $\phi \approx 0.74$ for spheres [39], as achieved by Barlow packings). There are many experimental and computational algorithms which produce a relatively robust packing fraction $\varphi \approx 0.64$ for randomly packed monodisperse spheres [40, 41]. This number, widely designated as the *random close packing* (RCP) density, is not universal but generally depends on the packing protocol, as discussed shortly. Illustrations of the difference between disordered and ordered packings of spheres are shown in Fig. 1.2.

Disordered packings correspond to the *liquid* state, thermodynamically stable at low densities, while the ordered packings correspond to the *crystal* state, thermodynamically-stable at higher densities. We will not review the thermodynamics of hard-sphere systems in detail, and refer the reader to Ref. [6] for more background and additional references [c.f. Section 3.3 in Ref [6]]. It is important to point out that hard-particle systems are athermal, and the thermodynamic properties are solely a function of the density (volume fraction) ϕ [43]. One must consider a hard-particle system at a positive temperature ($kT = 1$), since the free-energy of the hard-sphere system consists entirely of the entropic term, $F = -TS$. At positive temperature, the time-averaged thermal motion of the particles leads to a well-defined free energy and its derivatives with respect to strain, i.e., stress and elastic moduli (bulk and shear modulus for isotropic states), exist just as for soft-particle systems (where however the limit $T \to 0$ is thermodynamically well-defined).

In three dimensions, extensive computational investigations have found clear evidence for a first-order phase transition from an isotropic liquid (*fluid*) phase[1] to a crystal (*solid*) phase corresponding to a Barlow packing[2], with a coexistence region spanning from $\phi \approx 0.49$ (the *freezing* point) to $\phi \approx 0.55$ (the *melting* point). In two dimensions, it is not yet known with confidence whether there is also a first-order phase transition with a freezing point at $\phi \approx 0.69$ and a very narrow coexistence region, or whether the transition from the disordered liquid to the ordered triangular (hexagonal) crystal is a continuous transition. An important theory for such a continuous transition was proposed by Kosterlitz and Thouless and augmented by Halperin, Nelson and Young (KTHNY), and it predicts that a two-dimensional crystal (possessing quasi-long-range translational order and long-range orientational order) undergoes a continuous melting transition via the unbinding of dislocations into a hexatic phase (short-range translational order and quasi-long-range orientational order), which in turn continuously melts into the disordered liquid (short-range translational and orientational order) via the unbinding of disclinations [47, 48, 49]. Strong finite-size effects in two dimensions make it difficult to numerically ascertain the validity of the KTHNY scenario. It is clear however that the liquid-solid transition for monodisperse hard disks is nearly continuous, and that the liquid crystallizes very easily, unlike in three dimensions, where there is a nucleation barrier that inhibits the crystallization process. On the other hand, crystallization is strongly inhibited for sufficiently polydisperse systems of disks [50, 51].

Next we briefly illustrate how the behavior LS algorithm is related to the thermodynamic properties of the system, for hard spheres in three dimensions. In Fig. 1.3 we show the *equation of state* (EOS) of a hard sphere system as the density is increased slowly, through the growth of particles at an expansion rate γ, starting from a liquid. The pressure of the system is measured in the MD algorithm by averaging over a time period that is as small as possible but sufficiently large to average over many collisions. Instead of plotting the reduced pressure p directly, we use the well-known fact that near jamming the reduced pressure is asymptotically given by the free-volume equation of state [52],

$$p = \frac{PV}{Nk_BT} = \frac{1}{\delta} = \frac{d}{1 - \phi/\phi_J}, \tag{1.1}$$

which can be inverted to give an estimate $\tilde{\phi}_J$ of the jamming density,

$$\tilde{\phi}_J = \frac{\phi}{1 - d/p}. \tag{1.2}$$

[1] It is technically more appropriate to call this the fluid (gas-liquid) phase, since there is no cohesion for hard particles and thus no gas-liquid phase transition, however, in analogy with systems of soft particles we refer to the dense fluid as a liquid.

[2] It is not presently agreed-upon whether the true equilibrium solid phase is the FCC or the HCP crystal, although it is commonly believed that it is not a random stacking variant of the FCC lattice [44, 45, 46].

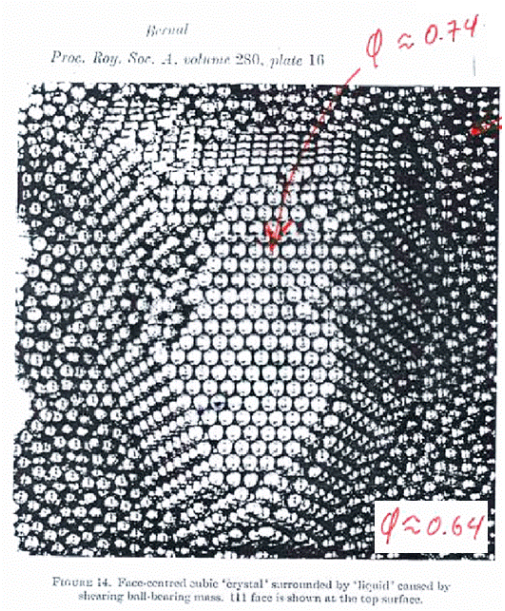

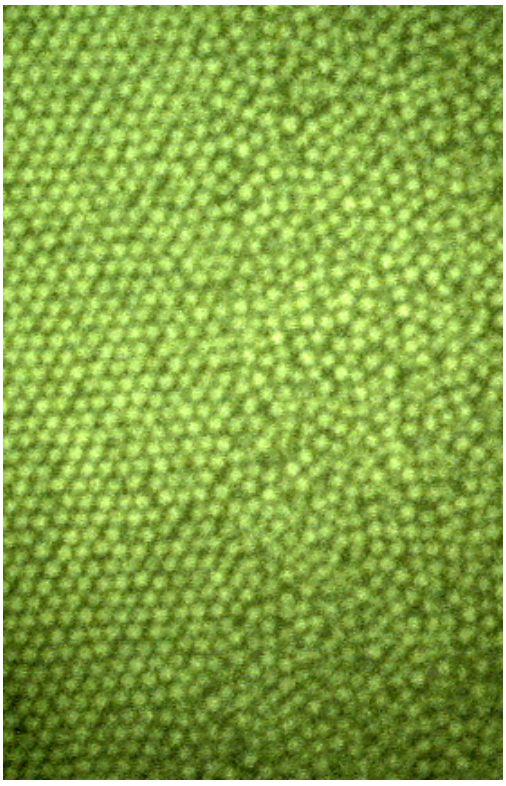

Figure 1.2: (*Left*) A carefully arranged ordered packing of ball bearings surrounded by a disordered pile of bearings, taken from Ref. [42]. (*Right*) [Courtesy of Paul Chaikin] A colloidal system of spherical particles in an external electric field, causing a density gradient from a less dense disordered liquid structure (right), to a more dense ordered crystal structure (left).

Since the pressure increases very rapidly near jamming, it is more convenient to plot the estimated jamming density $\tilde{\phi}_J(\phi)$ instead of the pressure $p(\phi)$, and we do this in Fig. 1.3 for several different expansion rates. Of course, the particles can also shrink ($\gamma < 0$) starting from a dense configuration such as the FCC crystal, in which case the EOS along the crystal branch can be obtained for densities ranging from the crystal jamming density to the freezing density $\phi \approx 0.50$.

In the limit $\gamma \to 0$ we obtain the thermodynamic equilibrium behavior of the hard-sphere system, and for small enough γ the system is in quasi-equilibrium, in the sense that the rate of density change is slow enough to allow for full relaxation of the system. For finite γ the system is not in true equilibrium, and in fact, as the relaxation time of the liquid grows due to the increase in density and increased proximity to jamming, the system may become trapped in a glassy state. This is exactly what is observed in Fig. 1.3. For small γ there is a first-order transition from the liquid to the crystal branch around the melting point $\phi \approx 0.55$, however, for larger γ there is a kinetic glass transition around $\phi \approx 0.6$ leading to non-equilibrium glassy states that eventually produce random jammed packings with jamming density $\phi_J \approx 0.64$. If the expansion rate is intermediate *partial crystallization* occurs leading to the formation of small nucleated crystallites inside a random packing matrix, or the formed polycrystal can be distorted and have multiple grain boundaries between crystallites with different orientations. In summary, the packings produced by the LS algorithm will have a density in the range $0.64 - 0.74$, depending on the exact parameters used in the packing-generation protocol. However, a wide range of expansion rates γ take liquid configurations into jammed packings with density in the range $\phi_J \approx 0.635 - 0.655$, which are disordered packings showing no apparent signs of local or global ordering.

1.2.1 The Maximally Random Jammed (MRJ) State

Using the variety of algorithms we have described, we can produce a variety of packings, either by using different algorithms, or by using one algorithm and changing some tunable parameter (such as γ in the LS algorithm). Focusing just on jammed packings, a challenging problem is their enumeration and classification. Since one cannot enumerate all possible packings even for a small number of particles, it is desirable to devise a small set of parameters that can characterize packings well. Two important scalar properties of packings are the *density (packing fraction)* ϕ and *order metric* ψ. For any two states X and Y, $\psi_X > \psi_Y$ implies that state X is to be considered as more ordered than state Y. Candidates for such an order metric include various translational and orientational order parameters [6], but the search for better order metrics is still very active [56]. One promising avenue is exploring the connections between entropy (information content) and packing disorder [57].

We can use a hypothetical scalar order-metric ψ to measure the amount of order in a packing, such that $\psi = 1$ corresponds to fully ordered (for example, the perfect FCC crystal), and $\psi = 0$ corresponds to perfectly disordered (Poisson distribution of sphere centers) packings. Figure 1.4 from Ref. [58] shows a conjectured region of feasible hard-sphere packings in the $\phi - \psi$ plane. It is clear that only a small subset of this feasible region will be occupied by jammed packings (for a *given* jamming category), as schematically indicated in Fig. 1.4. Several limit points in this region are particularly interesting:

- Point A corresponds to the lowest-density jammed packing, and its location strongly depends on the jamming category used. It can be shown that there are zero-density locally jammed disk packings (see references and discussion in Ref. [59]), and the same is true for sphere packings. However, for collectively and strictly jammed packings, it is not known what are the lowest possible densities.

- Point B corresponds to the most dense jammed packing. It has of course already been identified to be a triangular packing for disks and the FCC/HCP variant lattice for spheres. But much less is known about polydisperse packings [7, 59], or packings of nonspherical particles.

- MRJ point represents the maximally random jammed (MRJ) state [58], which has recently supplanted the ill-defined "random close packed" (RCP) state. The RCP state was widely believed to have a packing fraction $\varphi \approx 0.63 - 0.64$ in three dimensions. The MRJ state is the most disordered jammed packing in a given jamming category (locally, collectively or strictly jammed). The MRJ state is well-defined for a given jamming category and choice of order metric. This work focuses primarily on the generation and analysis of packings representative of the MRJ state for collective and strict jamming, for both spherical and non-spherical particles.

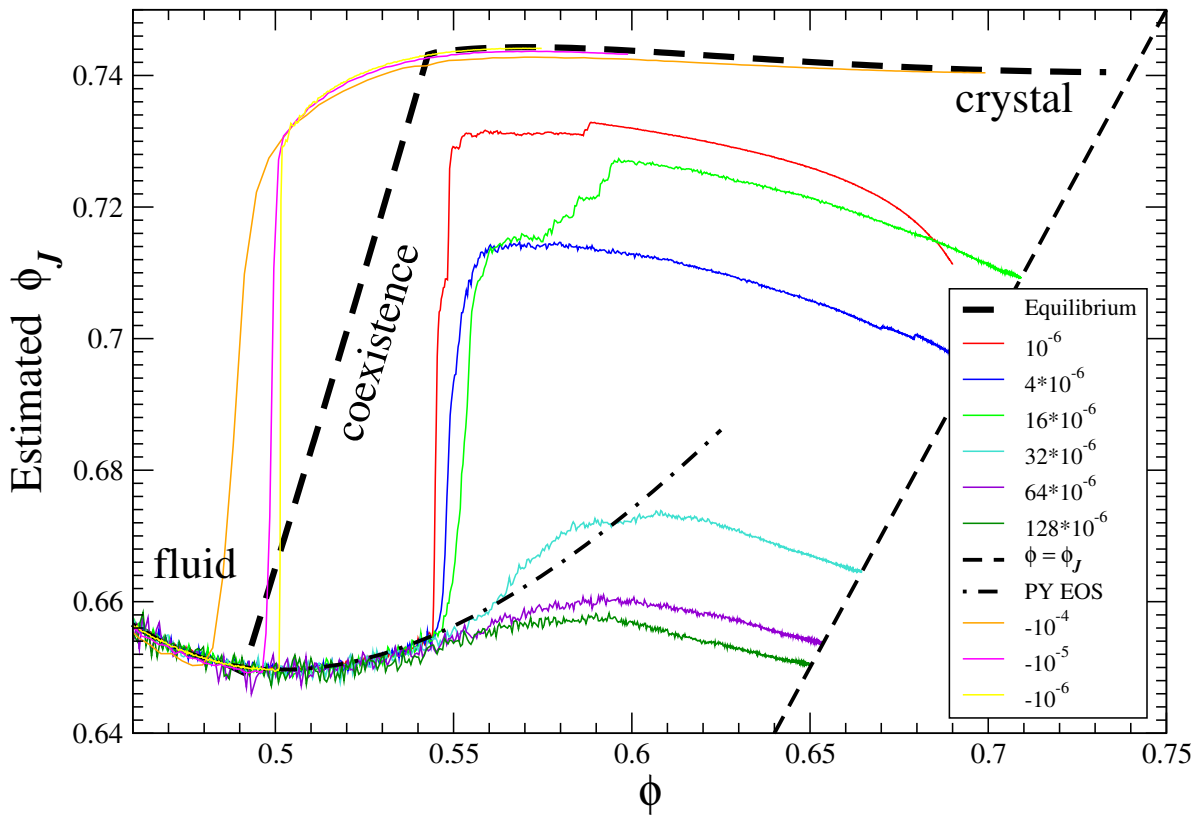

Figure 1.3: The estimated jamming packing fraction $\tilde{\phi}_J$ as a function of packing fraction ϕ for $d = 3$. Shown are systems of 4096 spheres with various expansion rates, showing the crystallization that occurs for sufficiently slow particle expansion and the kinetic glass transition for fast particle expansion, leading to disordered packings. Also shown are results for systems of 10976 spheres placed in an FCC lattice with negative expansion rates $\gamma = -10^{-4}$, -10^{-5}, and -10^{-6} (last three curves). For comparison, we plot approximations to the equilibrium EOS for the fluid phase, the coexistence region, and the crystal phase [53, 54], as well as the Percus-Yevick (PY) [55] EOS for the fluid phase.

1.2.2 Polydisperse Packings

The LS packing algorithm as well as the concept of MRJ can easily be generalized to packings of nonspherical particles as well as packings of particles of different sizes. We will extensively discuss monodisperse packings of nonspherical particles in this work, and also occasionally make use of *polydisperse* packings as well. Here we briefly comment on the impact polydispersity has on disordered jammed packings of hard spheres, without trying to give a comprehensive review of the literature.

The impact of polydispersity on packing densities and structure has been studied in the literature primarily for hard spheres. Incorporating polydispersity in the LS algorithm for hard spheres amounts to allowing each sphere i to have a different diameter D_i and to grow at its own rate γ_i [60]. In order to preserve the probability distribution of relative particle sizes $P_D(D/\bar{D})$, where $\bar{D}$ is the average sphere diameter, the LS algorithm is started with the sphere diameters drawn from the specified distribution P_D, and the particle expansion rates are made proportional to the particle diameters, $\gamma_i = dD_i/dt \sim D_i$. The simplest case occurs when there are only two different sphere diameters of size ratio $\kappa = D_1/D_2 > 1$, i.e., for *bidisperse* packings. The statistical properties of the MRJ state, and in particular the jamming density, depends on both the *size dispersity* κ, and the *composition* of the packing, i.e., the fraction of particles that are large, $0 < x < 1$. Several of the packing algorithms discussed in Section 1.1.1 have been used to study random packings of hard disks and spheres [60, 25, 61, 62, 17, 63, 64, 65, 66]. Amorphous jammed sphere packings with a continuous

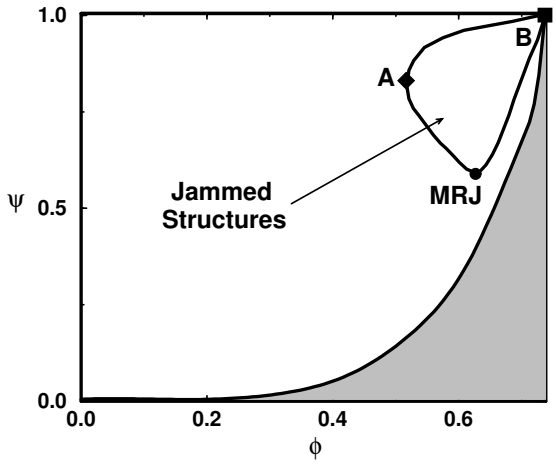

Figure 1.4: A highly schematic plot of the subspace in the density-disorder ($\phi - \psi$) plane where strictly jammed three-dimensional packings exist. Point A corresponds to the lowest-density jammed packing, and it is intuitive to expect that a certain ordering will be needed to produce low-density jammed packings. Point B corresponds to the most dense jammed packing. Point MRJ represents the maximally random jammed state. This is the most disordered jammed packing. We conjecture that the Lubachevsky-Stillinger packing algorithm [12, 13] typically produces packings along the right (maximally dense) branch, and we do not know of an algorithm that produces packings along the left (minimally dense) branch.

distribution of diameters have also been simulated [63, 64, 65]. We note that the generation of packings with large size dispersity is computationally very challenging (see the discussion in Chapter 3).

Crystal (ordered) binary packings of hard disks and spheres have also been discussed [61, 67, 68], although little is known rigorously. It is believed that for small size dispersity the crystals are phase-separated, with the two different particle sizes demixing and forming separate monodisperse crystal structures. In analogy to the generation of dense packings for monodisperse spheres, using the LS algorithm it is possible to create amorphous packings of bidisperse spheres at higher packing fractions than the limit of the fluid phase in the equilibrium system. However, extensive computational and experimental research has shown that compared to the monodisperse case, hard-sphere systems with sufficient size polydispersity (for example, $\kappa > 1.2$) tend to remain amorphous over a more broad range of packing fractions, and crystallization (i.e., crystal nucleation) is strongly suppressed [69, 60]. This is especially pronounced in two dimensions, where monodisperse hard disk systems crystallize readily and do not form amorphous jammed packings. In three dimensions crystal nucleation is an activated process and amorphous hard sphere packings are easily generated, however, crystal packings are thermodynamically strongly favored [70]. Bidisperse systems are thus used as model glass formers [51], and we will study the thermodynamics of binary hard-disk systems in Chapter 13. For polydisperse systems in which there is a continuous distribution of particle radii, the phase diagram and the existence and structure of thermodynamic crystal and glassy phases is still actively debated.

1.3 Summary of Research

In this section we summarize the research that has appeared in publications relating to this Dissertation. The results of this research are presented in subsequent chapters, with some added technical details.

1.3.1 Generating Jammed Hard-Particle Packings

We already discussed briefly the modified Lubachevsky-Stillinger algorithm that we use to generate the packings studied in this work. In Chapter 2 we give technical details on overlap potentials for nonspherical particles, used extensively in subsequent chapters. In Chapter 3 we discuss our collision-driven MD algorithm for nonspherical hard-particles, and in particular, hard ellipsoids.

1.3.1.1 Collision-Driven Molecular Dynamics Algorithm

Our collision-driven MD algorithm for hard-particle systems was described in detail in a two-part series of papers [11] and is presented in Chapter 3. In the first part, we presented in considerable detail a collision-driven molecular dynamics algorithm for a system of nonspherical particles, within a parallelepiped simulation domain, under both periodic or hard-wall boundary conditions. The algorithm extends previous event-driven molecular dynamics algorithms for spheres, and is most efficient when applied to systems of particles with relatively small aspect ratios and with small variations in size. We presented a novel partial-update near-neighbor list (NNL) algorithm that is superior to previous algorithms at high densities, without compromising the correctness of the algorithm. This efficiency of the algorithm was further increased for systems of very aspherical particles by using bounding sphere complexes (BSC). These techniques will be useful in any particle-based simulation, including Monte Carlo and time-driven molecular dynamics. Additionally, the algorithm allows for a non-vanishing rate of deformation of the boundary, which can be used to model macroscopic strain and also alleviate boundary effects for small systems.

In the second part of this series of papers we specialized the algorithm to systems of ellipses and ellipsoids. The theoretical machinery needed to treat such particles, including the overlap potentials, was developed in full detail. We described an algorithm for predicting the time of collision for two moving ellipses or ellipsoids. We presented performance results for our implementation of the algorithm, demonstrating that, for dense systems of very aspherical ellipsoids, the novel techniques of using neighbor lists and bounding sphere complexes offer as much as two orders of magnitude improvement in efficiency over direct adaptations of traditional event-driven molecular dynamics algorithms. The practical utility of the algorithm was demonstrated by presenting several interesting physical applications, including the generation of jammed packings inside spherical containers, the study of contact force chains in jammed packings, and melting the densest-known equilibrium crystals of prolate spheroids.

1.3.2 Jamming in Hard-Particle Packings

An essential component of understanding dense hard-particle systems is rigorously understanding the terminal singular states representing *ideal jammed packings*. Such packings have been called *rigid packings* in the mathematics literature, and in fact rigidity theory provides the necessary tools for developing a mathematical framework for jamming in hard-particle packings. This framework provides the basis for better understanding of *nearly jammed* packings, including their thermodynamics and kinetics, in either thermal settings such as glasses, or athermal systems such as granular media. In Chapter 4 we discuss jamming in hard-sphere packings, which will give us the necessary background for extending the theory to packings of nonspherical particles in Chapter 5.

1.3.2.1 Hard Sphere Packings

Torquato and Stillinger [15] proposed a classification scheme of jammed packings into hierarchical categories of *locally*, *collectively* and *strictly jammed* configurations. In Ref. [71], we presented a rigorous and practical algorithm to assess whether an ideal hard-sphere packing in two or three dimensions is jammed according to the aforementioned categories, as described in detail in Chapter 4. The algorithm is based on linear programming (LP) and is applicable to regular as well as random packings of finite size with hard-wall and periodic boundary conditions. If the packing is not jammed, the algorithm yields representative multi-particle unjamming motions. Furthermore, we extended the jamming categories and the testing algorithm to packings with significant interparticle gaps. We described in detail two variants of the proposed randomized linear programming approach to test for jamming in hard-sphere packings. The first algorithm treats ideal packings in which particles form perfect contacts. Another algorithm treats the case of jamming in packings

with significant interparticle gaps. This extended algorithm allows one to explore more fully the nature of the feasible particle displacements. We implemented the algorithms and applied them to ordered as well as random packings of disks and spheres with periodic boundary conditions and gave some representative results for large disordered disk and sphere packings.

In Ref. [72] we applied the LP algorithms from Ref. [71] to test jamming categories of ordered lattices as well as random packings of disks and spheres under periodic boundary conditions. The random packings were produced computationally with a variety of packing generation algorithms, all of which should, in principle, produce at least collectively jammed packings. Our results highlighted the importance of jamming categories in characterizing particle packings. One important and interesting conclusion was that the amorphous monodisperse sphere packings with density $\varphi \approx 0.64$ were for practical purposes strictly jammed in three dimensions, but in two dimensions the monodisperse disk packings at previously reported "random close packed" densities of $\varphi \approx 0.83$ were not even collectively jammed. On the other hand, amorphous bidisperse disk packings with density of $\varphi \approx 0.84$ were virtually strictly jammed. This clearly demonstrated that one cannot judge "stability" in packings based solely on local criteria.

1.3.2.2 Hypostatic Packings of Hard Ellipsoids

As discussed in Chapter 5, testing for jamming in packings of nonspherical particles such as ellipsoids is significantly more difficult than for spheres. The reason is not just the increased complexity of keeping track of particle orientations, but rather, a fundamental aspect of contacts between nonspherical particles. Namely, impenetrability constraints are guaranteed to be strictly concave for spheres, which allows for rigorous linearization sufficiently close to the jamming point. However, this guarantee is not present for nonspherical particles and in fact the nonlinear character of the constraints needs to be taken into account. The rigidity class for jammed packings of ellipsoids belongs to the class of *prestress stability* as defined by Connelly *et al.* [73], and testing for it is in general an exponentially hard problem. In Ref. [74] we develop first- and second-order conditions for jamming, and demonstrate that ellipsoid packings can be jammed even though they are hypostatic. We apply an algorithm using these conditions to computer-generated hypostatic ellipsoid and ellipse packings and demonstrate that our algorithm does produce jammed packings, even close to the sphere point.

1.3.3 Statistical Properties of Random Hard-Sphere Packings

After having generated a nearly jammed packing and verified its jamming category rigorously, the next step is the analysis of the statistical and macroscopic properties of the jammed state itself. These include the interparticle force chains which resist particles' motions, the mechanical properties of the packings, the local and global geometrical structure of the systems, the correlations between particles, and the nature of the ordering present. We have performed some important analysis of MRJ packings of hard spheres. In Chapter 9 we focus on short-range particle-particle (pair) correlations by examining the features of the pair correlation function $g_2(r)$ near contact for three-dimensional nearly jammed packings. In Chapter 10 we extend these studies to long-range correlations by focusing on the structure factor $S(k)$ near the origin. We study pair correlations in sphere packings in higher dimensions (specifically, $d = 4, 5$ and 6) in Chapter 11.

1.3.3.1 Correlation Functions

In Ref. [75] we studied the approach to jamming in hard-sphere packings, and, in particular, the pair correlation function $g_2(r)$ around contact, both theoretically and computationally. We present these studies in Chapter 9. Our computational data unambiguously separated the narrowing delta-function contribution to g_2 due to emerging interparticle contacts from the background contribution due to near contacts. We also showed that disordered hard-sphere packings are strictly isostatic, i.e., the number of exact contacts in the jamming limit is exactly equal to the number of degrees of freedom, once rattlers are removed. For such isostatic packings, we derived a theoretical connection between the probability distribution of interparticle forces $P_f(f)$ and the contact contribution to g_2, and verified this relation computationally. We observed a clear maximum in P_f and a nonzero probability of zero force, as well as an unusual power-law divergence in the near-contact contribution to g_2, persistent even in the jamming limit. Additionally, we presented high-quality numerical data on the two discontinuities in the split-second peak of g_2, and use a shared-neighbor

analysis of the graph representing the contact network to study the local particle clusters responsible for the peculiar features. We also investigated partially crystallized packings along the transition from maximally disordered to fully ordered packings.

1.3.3.2 Density Fluctuations

Continuing on previous theoretical investigations of local density fluctuations in atomic systems by Torquato and Stillinger in 2003, we computationally studied jammed disordered hard-sphere packings of as many as one-million particles in Ref. [76] and we present the results in Chapter 10. We reported on the decay of the pair correlation function, demonstrating that it is consistent with exponentially damped oscillatory but without ruling out the possibility of a weak (quasi)long-ranged tail. We showed that local density fluctuations are suppressed from volume to surface ones, i.e., the generated packings are hyperuniform, with structure factor which vanishes at the origin to within 10^{-3}. The numerical results suggested a strange non-analytic linear dependence of the structure factor around the origin, while the direct pair correlation function showed a significant long-range tail outside the core, unlike the stable liquid. Finally, we demonstrated that the generated packings are saturated and all voids in the mechanically rigid backbone of the packing are filled with rattlers. Our results illuminated many open questions about density fluctuations in glassy atomic systems.

1.3.3.3 Higher Dimensions

In Ref. [77] and Chapter 11 we present the first study of disordered jammed hard-sphere packings in four and five dimensions using event-driven molecular dynamics. We give the first estimates for the densities of the maximally random jammed states in four and five dimensions and calculate the pair correlation function $g_2(r)$ and structure factor $S(k)$ for these states. We find that both $g_2(r)$ and $S(k)$ are significantly damped in four and five dimensions compared to three dimensions, consistent with a decorrelation principle recently proposed Torquato *et al.* 2005, stating that correlations diminish in very high dimensions. We verify that, as in three dimensions, the packings show no signs of crystallization, are isostatic, and have a power-law divergence in $g_2(r)$ at contact. Additionally, we obtain estimates of the freezing and melting densities for the liquid-solid transition.

1.3.4 Jammed Packings of Hard Ellipsoids

Hard sphere packings have been used as model systems extensively primarily because of their simplicity, even though realistic packings are rarely made of truly spherical particles. The computational investigations we present in Chapter 5 demonstrate that adding even small asphericity significantly affects disordered jammed packings, leading to significantly higher densities, contact numbers, and also to a different kind of jamming than that in sphere packings. Experiments validating some of our computational results are described in Chapter 7. Crystal packings of ellipsoids are studied in Chapter 8.

1.3.4.1 Computational and Experimental Studies of Disordered Packings

Packing problems, how densely objects can fill a volume, are among the most ancient and persistent problems in mathematics and science. For equal spheres, it has only recently been proved that the face-centered cubic lattice has the highest possible packing fraction $\varphi \approx 0.74$. It is also well-known that the corresponding amorphous (MRJ) packings have $\varphi \approx 0.64$. We used our MD algorithm to study random (disordered) packings of hard ellipsoids in Ref. [78], as we detail in Chapters 5 and 7. We showed experimentally and computationally that ellipsoids can randomly pack more densely; up to $\varphi = 0.68 - 0.71$ for spheroids with an aspect ratio close to that of M&M'S Candies® , and even approach $\varphi \approx 0.74$ for general ellipsoids. We suggested that the higher density relates directly to the higher number of degrees of freedom per particle and supported this claim by measurements of the number of contacts per particle Z, obtaining $Z \approx 10$ for our spheroids as compared to $Z \approx 6$ for spheres.

Continuing on the computational and experimental work on jammed packings of hard ellipsoids in Ref. [78], in Ref. [74] and Chapter 5 we consider jamming in packings of smooth strictly convex nonspherical hard particles. We explain why the isostatic conjecture, stating that for large disordered jammed packings the

average contact number is twice the number of degrees of freedom per particle, does not apply to nonspherical particles. We also consider packings that are nearly jammed and draw connections to packings of deformable (but stiff) particles. Finally, we consider the jamming conditions for nearly spherical particles and explain quantitatively the behavior we observe in the vicinity of the sphere point.

The simulations in Ref. [78] predicted the ellipsoid shape which gives the highest random packing density, namely, that ellipsoids with axes ratios near $1.25 : 1 : 0.8$ form amorphous packings as dense as the densest crystal packing (FCC) of spheres. Subsequent experimental work confirmed the predictions [79] and is presented in Chapter 7. We demonstrated that such dense packings are realizable by manufacturing ellipsoids using stereolitography and packing them inside spherical containers. The analysis of the packings required better understanding of finite-size effects, and we used both simulations and novel experimental methods to understand and minimizes surface effects and to obtain good estimates of bulk packing densities. We showed that, in a sphere, the radial packing fraction $\phi(r)$ can be obtained from $V(h)$, the volume of added fluid to fill the sphere to height h. We also obtained $\phi(r)$ from a magnetic resonance imaging (MRI) scan. The measurements of the overall density ϕ_{avg}, the radial density distribution $\phi(r)$, and the core density $\varphi_C \approx 0.740 \pm 0.005$ agreed with simulations. This verified that idealized computer packings of frictionless hard particles are relevant to the quantitative understanding of real-life packings of frictional macroscopic particles in a gravitational field.

1.3.4.2 Crystal Packings

Extensive experience with spheres has shown that for reasonably large packings, sufficiently slowing down the growth of the density, so that the hard-particle system remains close to the equilibrium solid branch of the equation of state, leads to packings near the FCC lattice. This however requires impractically long simulation times for large ellipsoid packings. By running the MD packing algorithm for very small unit cells, from 4 to 16 particles per unit cell, we were able to identify crystal packings of ellipsoids significantly denser than the FCC lattice [80], as explained in Chapter 8. Subsequent analytical calculations suggested by the simulation results led us to discover ellipsoid packings with a remarkably high density of $\phi \approx 0.7707$. The family of new packings we discovered are crystal (periodic) arrangements of nearly spherically-shaped ellipsoids, and always surpass the densest lattice packing. The maximum density of $\varphi \approx 0.7707$ is achieved for both prolate and oblate ellipsoids with aspect ratios of $\sqrt{3}$ and $1/\sqrt{3}$, respectively, and each ellipsoid has 14 touching neighbors. These results do not exclude the possibility that even denser crystal packings of ellipsoids could be found, and that a corresponding Kepler-like conjecture could be formulated for ellipsoids.

1.3.5 Thermodynamics of Hard-Particle Systems

The better understanding of nearly jammed hard-particle packings enables the study of ensembles of packings, i.e., the thermodynamics of dense hard-particle systems. Hard-particle packings, although completely dominated by entropic exclusion, provide an incredibly rich phase behavior representative of real phases found in experimental systems. They capture the essential competition for lower free energy between order and disorder, namely, between free volume and degeneracy. This competition drives the thermodynamic liquid-solid phase transitions, and is intimately related to the elusive glass transition. Each of the studies described in previous sections contributes to better understanding of the thermodynamics of hard-particle systems, for example, the identification of the densest packings directly determines the high-density crystal phases. In Chapter 12 we investigate the phase behavior of hard rectangles of aspect ratio two (dominos), using superellipses in our molecular dynamics algorithm. In Chapter 6 we describe a molecular-dynamics algorithm to calculate the free energy of nearly-jammed hard particle packings, and then apply this algorithm to the question of the existence of an ideal glass transition in a binary hard-disk system in Chapter 13.

1.3.5.1 Hard Rectangles (Dominos)

Previous Monte Carlo investigations by Wojciechowski and Frenkel have found two unusual phases in two-dimensional systems of anisotropic hard particles: a tetratic phase of four-fold symmetry for hard squares, and a nonperiodic degenerate solid phase for hard-disk dimers. In Ref. [81], we studied a system of hard rectangles of aspect ratio two, i.e., hard-square dimers (or dominos), and demonstrate that it exhibits a solid phase with both of these unusual properties, as detailed in Chapter 12. The solid shows tetratic, but

not nematic, order, and it is nonperiodic having the structure of a random tiling of the square lattice with dominos. We obtain similar results with both a classical Monte Carlo method using true rectangles and collision-driven molecular dynamics employing superellipses. Although we have not yet performed exact free-energy calculations, we expect that the random domino tiling is thermodynamically stabilized by its degeneracy entropy, well-known to be $1.79k_B$ per particle from previous studies of the dimer problem on the square lattice. Our numerical observations were consistent with a KTHNY two-stage phase transition scenario with two continuous phase transitions, the first from isotropic to tetratic liquid, and the second from tetratic liquid to solid.

1.3.5.2 Free-Energy Calculations via Molecular Dynamics

In Ref. [82] and Chapter 6, we present in considerable detail an event-driven molecular dynamics algorithm for measuring the free energy of nearly jammed hard-particle packings. This Bounding Cell Molecular Dynamics (BCMD) algorithm calculates exactly the free-energy of a single-occupancy cell (SOC) model, as first defined by Kirkwood, in which each particle is restricted to the neighborhood of its initial position using a hard-wall bounding cell. It is based on previous MD algorithms appearing in the literature, however, several small but important modifications enable us to apply it to nonspherical particles as well as to measure the free-energy change during continuous irreversible transformations. Additionally, we point connections to the well-studied problem of computing the volume of convex bodies in high dimensions using random walks. We test and verify the numerical accuracy of the method by comparing against rigorous (asymptotic) results for the free energy of isostatic disordered packings of both hard spheres and ellipsoids, for which the free energy can be calculated directly as the volume of a high-dimensional simplex. We also compare our results to previously published Monte Carlo results for hard-sphere crystals and find excellent agreement.

1.3.5.3 Ideal Glass Transition for Binary Disk Mixtures

We have applied the BCMD algorithm described in Section 6.3 to a model glass former, namely, binary disk mixtures with large-to-small disk stoichiometry of $1:2$ and diameter ratio $\kappa = 1.4$. We summarize the surprising and important findings from our investigations in Ref. [57], and we give considerable details in Ref. [83] and Chapter 13. We have produced bidisperse hard disk glasses by compressing liquids at a wide range of compression rates using event-driven MD. We observe that the density at which the liquid falls out of equilibrium and becomes glassy increases with decreasing compression rates, and find that the jamming density of the glass itself continuously increases as the liquid is given more time to equilibrate. Even at the slowest compression rates equilibration seems unattainable beyond $\phi_g \approx 0.8$ with classical methods. Our free energy calculations give a freezing point of $\phi_F \approx 0.775$, and also show that the estimated configurational entropy near ϕ_g is very close to the entropy of mixing, a fact seen in numerous other studies in the literature. Furthermore, the equilibrated liquids slightly below ϕ_g show micro-clustering of the large particles indicating that supercooled liquids at higher densities would show increased demixing of the two disk species. The configurational entropy has been postulated to go to zero at the density of an amorphous *ideal glass*, however, our numerical observations suggest that the predicted ideal glass is not an amorphous structure but rather a phase-separated crystal. In fact, we provide an explicit construction (based on discretized leveled random Gaussian fields) of an exponential number of binary jammed packings with density ranging from that of the MRJ state to the crystal. This construction clearly demonstrates that there is not a maximally dense amorphous packing, but rather a continuum of structures from most disordered to most ordered. This work shows that the very premise of the proposed theories of an ideal glass transition underlying the kinetic one is flawed.

1.4 Notation

We have tried to develop a clear and consistent notation, however, it is inevitable that some notation has evolved throughout the duration of this project. In particular, the same Latin or Greek letter may be used to denote different physical quantities if those quantities are localized in appearance and the context clarifies their usage. Additionally, the sign convention for positive and negative quantities may be different between certain chapters. We will alert the reader to any such notational inconsistency.

1.4.1 Matrix and Vector Notation

We use a notation that tries to distinguish between scalar, vector and matrix (or tensor) quantities, however, in order to avoid excessive indexing and notation complexity we will often rely on the context for clarity. The notation is kept independent of the dimensionality d whenever possible.

We will use matrix notation extensively, and denote vectors and matrices with bolded letters, and capitalize matrices in most cases. Infinite-dimensional or discrete quantities such as sets or graphs, higher order matrices (linear operators), or other quantities that are neither scalars nor vectors, will typically be denoted with script letters. We will often capitalize the letter denoting a vector to denote a matrix obtained from that vector, for example, a diagonal matrix containing the vector $\boldsymbol{\gamma}$ on its diagonal will be denoted with $\boldsymbol{\Gamma}$ when the context clarifies the meaning. Matrix multiplication is assumed whenever products of matrices or a matrix and a vector appear. We prefer to use matrix notation whenever possible and do not carefully try to distinguish between scalars and matrices of one element. We denote the dot product $\mathbf{a} \cdot \mathbf{b}$ with $\mathbf{a}^T\mathbf{b}$, and the outer product $\mathbf{a} \otimes \mathbf{b}$ with $\mathbf{a}\mathbf{b}^T$. We denote a vector with all entries unity (of the appropriate size for the context) by $\mathbf{e} = 1$, so that $\sum_i a_i = \mathbf{e}^T\mathbf{a}$, and a unit matrix with $\mathbf{I}$.

We consider matrices here in a more general linear operator sense, and they can be of order higher than two (i.e., they do not necessarily have to be a rectangular two-dimensional array). We refer to differentials as gradients even if they are not necessarily differentials of scalar functions. Gradients of scalars are considered to be column vectors and gradients of vectors or matrices are matrices or matrices (linear operators) of higher rank. More consistent notation with derivatives of vectors and matrices can be developed and should in principle be employed in calculations to avoid confusions about the order of matrix multiplications and transpositions [84].

1.4.2 Cross Products

When dealing with non-spherical particles, rotations, and therefore cross products, appear frequently. In three dimensions (3D), the cross product of two vectors is a linear combination of them that can be thought of as matrix-vector multiplication

$$\mathbf{a} \times \mathbf{b} = \mathbf{A}\mathbf{b} = -\mathbf{b} \times \mathbf{a} = -\mathbf{B}\mathbf{a} \tag{1.3}$$

where

$$\mathbf{A} = |\mathbf{a}|_\times = \begin{bmatrix} 0 & -a_z & a_y \\ a_z & 0 & -a_x \\ -a_y & a_x & 0 \end{bmatrix} = -\mathbf{A}^T$$

is a skew-symmetric matrix which is characteristic of the cross product and is derived from a vector. We will simply *capitalize* the letter of a vector to denote the corresponding *cross product matrix* (like $\mathbf{A}$ above corresponding to $\mathbf{a}$), or use $|\mathbf{a}|_\times$ when capitalization is not possible. This capitalization may lead to confusion with other capitalizations of vectors/matrices, however, we hope that in such cases the context will clarify the notation.

In two dimensions (2D), there are two "cross products". The first one gives the velocity of a point $\mathbf{r}$ in a system which rotates around the origin with an angular frequency $\boldsymbol{\omega}$ (which can also be considered a scalar ω),

$$\mathbf{v} = \boldsymbol{\omega} \boxtimes \mathbf{r} = \begin{bmatrix} -\omega\mathbf{r}_y \\ \omega\mathbf{r}_x \end{bmatrix} = \boldsymbol{\Omega}\mathbf{r}, \tag{1.4}$$

where

$$\boldsymbol{\Omega} = |\boldsymbol{\omega}|_\boxtimes = \begin{bmatrix} 0 & -\omega \\ \omega & 0 \end{bmatrix} = -\boldsymbol{\Omega}^T$$

is a cross product matrix derived from $\boldsymbol{\omega}$. The second kind of "cross product" gives the torque around the origin of a force $\mathbf{f}$ acting at a point (arm) $\mathbf{r}$,

$$\boldsymbol{\tau} = \mathbf{f} \times \mathbf{r} = -\mathbf{r} \times \mathbf{f} = [f_x r_y - f_y r_x] = \mathbf{F}^L\mathbf{r}, \tag{1.5}$$

where

$$|\mathbf{f}|_\times^L = \mathbf{F}^L = \begin{bmatrix} -f_y & f_x \end{bmatrix} = -\left(|\mathbf{f}|_\times^R\right)^T = -\left(\mathbf{F}^R\right)^T$$

is another cross product matrix derived from a vector (the L and R stand for left and right multiplication, respectively). Note that in three dimensions all of these coincide, $\mathbf{F}^L = \mathbf{F}^R = \mathbf{F}$, $\boxtimes \equiv \times$, and $\mathbf{AB} = \mathbf{ba}^T - \left(\mathbf{a}^T\mathbf{b}\right)\mathbf{I}$. In two dimensions the two cross products are related via $\mathbf{a} \boxtimes \mathbf{b} = \mathbf{Ab} = -\mathbf{B}^R a$. The wedge product generalizes the cross product in higher dimensions [85].

1.4.3 Representing Rigid-Body Orientations

Nonspherical particles are rigid bodies and representing their position requires representing their spatial orientation in addition to the position of their centroid. Representing the orientation of a rigid body in a computationally convenient way has been a subject of debate in the past [86]. A rigid body has

$$f_R = \frac{d\,(d-1)}{2} = \begin{cases} 1 \text{ if } d = 2 \\ 3 \text{ if } d = 3 \end{cases} \tag{1.6}$$

rotational degrees of freedom, and this is the minimal number of coordinates needed to specify the configuration of a hard nonspherical particle, in addition to the usual d coordinates needed to specify the position of the centroid. In two dimensions orientations are easy to represent via the angle φ between the major semiaxes of the ellipsoid and the x axis. But in three dimensions specifying three (Euler) angles is numerically unstable, and extensive experience has determined that for MD computationally the best way to represent orientations is via *normalized quaternions*, which in fact represent finite *rotations* starting from an initial reference configuration (but see Ref. [87] for a discussion). In the case of ellipsoids this reference configuration is one in which all semiaxes are aligned with the coordinate axes. In two dimensions we use a normalized complex number to represent orientation, but for simplicity we will sometimes use the term "quaternion" in both two and three dimensions. Higher dimensional generalizations are discussed in Ref. [88].

In three dimensions, normalized quaternions consist of a scalar s and a vector $\mathbf{p}$,

$$\mathbf{w} = [s, \mathbf{p}] = \left[\cos\frac{\varphi}{2}, \left(\sin\frac{\varphi}{2}\right)\hat{\boldsymbol{\varphi}}\right], \tag{1.7}$$

where $\hat{\boldsymbol{\varphi}}$ is the unit vector along the axis of rotation and φ is the angle of rotation around this axis, and the normalization condition

$$\|\mathbf{w}\|^2 = s^2 + \|\mathbf{p}\|^2 = 1$$

is satisfied. Therefore in three dimensions we use 4 numbers to represent orientation, which seems like wasting one floating-point number. It is in fact possible to represent the rotation with the oriented angle $\boldsymbol{\varphi} = \varphi\hat{\boldsymbol{\varphi}}$, which is just a vector with 3 coordinates. However, such a representation has numerical problems when $\varphi = 0$, and also the representation is not unique[3]. More importantly, combining rotations (as during rotational motion) does not correspond (as one may expect) to vector addition of the $\boldsymbol{\varphi}$'s, but it does correspond to quaternion multiplication of the $\mathbf{w}$'s, which is fast since there is no need of repeating the trigonometric evaluations. This is the reason why we also use quaternions in two-dimensions, and represent the orientation of a particle in the plane with 2 coordinates (components of a unit complex number),

$$\mathbf{w} = [s, p] = [\cos\varphi, \sin\varphi]\,. \tag{1.8}$$

The orthogonal *rotation matrix* corresponding to the rotation described by the quaternion (1.7) is given with

$$\mathbf{Q} = 2\left[\mathbf{pp}^T - s\mathbf{P} + \left(s^2 - \frac{1}{2}\right)\mathbf{I}\right]$$

in three dimensions, and

$$\mathbf{Q} = \begin{bmatrix} s & p \\ -p & s \end{bmatrix}$$

in two dimensions, corresponding to the complex number (1.8). The resulting orientation after first the rotation $\mathbf{Q}_1$ is applied and then the rotation $\mathbf{Q}_2$ is applied, $\mathbf{Q}_{12} = \mathbf{Q}_2\mathbf{Q}_1$, is represented by the quaternion product

$$\mathbf{w}_{12} = \mathbf{w}_1\mathbf{w}_2 = [s_1 s_2 - \mathbf{p}_1 \cdot \mathbf{p}_2, s_1\mathbf{p}_2 + s_2\mathbf{p}_1 - \mathbf{p}_1 \times \mathbf{p}_2] \tag{1.9}$$

[3]Note that the quaternion representation is also not unique since $-\mathbf{q}$ and $\mathbf{q}$ represent the same orientation.

in three dimensions, and by the complex number product

$$\mathbf{w}_{12} = \mathbf{w}_1\mathbf{w}_2 = [s_1s_2 - p_1p_2, s_1p_2 + s_2p_1] \tag{1.10}$$

in two dimensions.

In this work we are interested in particles which move continuously in time. The rate of rotation of a rigid body is given by the *angular velocity* $\boldsymbol{\omega}$ (which can also be considered a scalar ω in two dimensions), or equivalently, the infinitesimal change in orientation is given by the infinitesimal rotation $d\boldsymbol{\varphi} = \boldsymbol{\omega}dt$. The instantaneous time derivative of the normalized quaternion is given with

$$\dot{\mathbf{w}} = \frac{1}{2}\begin{bmatrix} s & -\mathbf{p} \\ \mathbf{p} & s\mathbf{I} + \mathbf{P} \end{bmatrix}\begin{bmatrix} 0 \\ \boldsymbol{\omega} \end{bmatrix}$$

in three dimensions, and with

$$\dot{\mathbf{w}} = \begin{bmatrix} s & -p \\ p & s \end{bmatrix}\begin{bmatrix} 0 \\ \omega \end{bmatrix}$$

in two dimensions. The time derivative of the corresponding rotation matrix is

$$\dot{\mathbf{Q}} = -\mathbf{Q}\boldsymbol{\Omega},$$

and this result has was used extensively in deriving the various time derivatives related to the contact function for ellipsoids, as will be given shortly. In Algorithm 1 we give a numerical prescription for updating the orientation of a particle rotating with a *constant* angular velocity for a time Δt.

---

1. Calculate the change in orientation $\mathbf{w}_{\Delta t}$ using $\boldsymbol{\varphi}_{\Delta t} = \boldsymbol{\omega}\Delta t$ in eq. (1.7) or (1.8).
2. Update the quaternion, $\mathbf{w} \leftarrow \mathbf{w}\mathbf{w}_{\Delta t}$, using eq. (1.9) or (1.10).
3. If $|\|\mathbf{w}\| - 1| > \epsilon_W$ (due to accumulation of numerical errors), renormalize the quaternion, $\mathbf{w} \leftarrow \mathbf{w}/\|\mathbf{w}\|$.

Algorithm 1: Update the orientation of a particle rotating with a uniform angular velocity $\boldsymbol{\omega}$ after a time step Δt.

1.4.3.1 Rigid-Body Dynamics

Using quaternions, the equation of motion for a spherically symmetric top

$$\ddot{\mathbf{w}} = |\dot{\mathbf{w}}|^2\,\mathbf{w} + \frac{1}{4I}\tilde{\boldsymbol{\tau}},$$

where I is the common moment of inertia (along all major axes), and the quaternion "torque"

$$\tilde{\boldsymbol{\tau}} = -\widetilde{\boldsymbol{\nabla}}_{\mathbf{w}}U$$

is the negative gradient of the potential energy, projected onto the unit quaternion hypersphere.

1.4.4 Particle Packings

A jammed particle packing has a *contact network* indicating the touching pairs of particles $\{i, j\}$. We will sometimes talk about a particular particle i or a particular contact $\{i, j\} \equiv ij$ and we will usually let the context determine what specific particle or contact is being referred to, or, if deemed necessary, put subscripts such as i or ij to make it specific what particle or contact is being referred to. The contact ji is physically the same *undirected* contact as ij, but the two *directed contacts* are considered distinct.

There are two primary kinds of vectors $\mathbf{x}$, *particle vectors* $\mathbf{X} = (\mathbf{x}_i) = (\mathbf{x}_1, \ldots, \mathbf{x}_N)$, which are obtained by concatenating together the vectors $\mathbf{x}_i$ (typically of size of the order of the space dimensionality d) corresponding to each of the N particles, and *contact vectors* $\mathbf{y} = (y_{ij}) = (y_1, \ldots, y_M)$, obtained by concatenating together the (typically scalar) values y_{ij} corresponding to each of the M contacts (numbered in arbitrary

order from 1 to M). Note the capitalization of particle vectors, which we will often do implicitly, to indicate that one can view $\mathbf{X}$ as a matrix where each row corresponds to a given particle. If a contact vector agglomerates a vector quantity attached to each contact, for example, the common normal vector $\mathbf{n}$ at the point of contact of two particles, it too would be capitalized, e.g., $\mathbf{N} = (\mathbf{n}_{ij})$.

Note that we often use "sphere" and "ellipsoid" in any dimension d, but sometimes we will emphasize "disk" and "ellipse" in two dimensions for clarity.

Part I

Jamming: Theory and Algorithms

Chapter 2

Overlap Potentials for Hard Particles

An essential component to any algorithm aimed at generating hard-particle packings is the verification of the nonoverlap conditions between pairs of particles. For spheres, testing for overlap is trivial, since two spheres overlap if and only if the distance between their centers is smaller than the sum of their radii. However, calculating the distance between nonspherical particles is not trivial. In this Chapter we develop in detail *overlap potentials* designed specifically for this purpose. In Section 2.1 we discuss designing overlap potentials for general convex particle shapes, and then we specifically focus on ellipsoids in Section 2.2. A simple but powerful generalization of ellipsoids, superellipsoids, is discussed in Section 2.3.

2.1 Nonoverlap Constraints

A packing is a collection of N hard particles in $\mathcal{R}^d$ such that no two particles overlap. Each particle i has d_f configurational *degrees of freedom* $\mathbf{q}_i$, for a total of $N_f = Nd_f$ degrees of freedom. A packing $\mathcal{Q} = (\mathbf{Q}, \phi)$ is characterized by the *configuration* $\mathbf{Q} = (\mathbf{q}_1, \ldots, \mathbf{q}_N) \in \mathcal{R}^{N_f}$, determining the positions of the centroid and the orientations of each particle, and the *packing fraction* (*density*), ϕ, determining the size of the particles. For spheres $\mathbf{Q} \equiv \mathbf{R}$ corresponds to only the positions of the centroids, and $d_f = d$. For nonspherical particles without any axes of symmetry there are an additional $d(d-1)/2$ rotational degrees of freedom, for a total of $d_f = d(d+1)/2$ degrees of freedom. In actual numerical codes particle orientation is represented using unit quaternions, which are redundant representations in the sense that they use $d(d-1)/2+1$ coordinates to describe orientation. When we focus on displacements of the particles $\Delta\mathbf{Q}$ from a reference jammed configuration $\mathbf{Q}_J$ we will represent particle orientations as a rotational displacement from a reference orientation $\Delta\boldsymbol{\varphi}$. In two dimensions $\Delta\boldsymbol{\varphi} = \Delta\varphi$ simply denotes the angle of rotation in the plane, and in three dimensions the direction of $\Delta\boldsymbol{\varphi}$ gives the axis of rotation and its magnitude determines the angle of rotation. For simplicity, we will sometimes be sloppy and not specifically separate centroid positions from orientations, and refer to $\mathbf{q}_i$ as (a generalized) *position*; similarly, we will sometimes refer to both forces and torques as (generalized) *forces*.

2.1.1 The PW Overlap Potential for Convex Particles

In this section we consider a general particle shape given by the inequality $\zeta(\mathbf{r}) \leq 0$, where the *shape function* ζ is strictly convex and defined through

$$\zeta(\mathbf{r}) = [\mu(\mathbf{r})]^2 - 1,$$

where μ is the unique *scaling factor* by which the particle needs to be resized in order for the point $\mathbf{r}$ to lie on its surface. The unnormalized normal vector to the surface at a given point $\mathbf{r}$, if the particle is rescaled so that it passes through it, is $\mathbf{n}(\mathbf{r}) = \boldsymbol{\nabla}\zeta(\mathbf{r})$. The requirement of having a smooth and convex particle shape is important in this work because it makes the normal vector $\mathbf{n}(\mathbf{r})$ a one-to-one function: There is a unique normal vector at any point on the surface of a particle, and given a normal vector one can determine the point on the surface that corresponds to that normal.

The nonoverlap condition between a pair of particles A and B can be thought of as an inequality between the positions and orientations of the particles. Define the displacement between the particle centroids $\mathbf{r}_{AB} = \mathbf{r}_A - \mathbf{r}_B$, and the unit vector joining the two particle centroids with $\mathbf{u}_{AB} = \mathbf{r}_{AB} / \|\mathbf{r}_{AB}\|$. For this purpose, we measure the distance between the two ellipsoids using the *overlap potential* $\zeta(A, B) = \zeta(\mathbf{q}_A, \mathbf{q}_B)$, whose sign not only gives us an overlap criterion,

$$\begin{cases} \zeta(A,B) > 0 \text{ if } A \text{ and } B \text{ are disjoint} \\ \zeta(A,B) = 0 \text{ if } A \text{ and } B \text{ are externally tangent} \\ \zeta(A,B) < 0 \text{ if } A \text{ and } B \text{ are overlapping,} \end{cases}$$

but which is also at least twice continuously differentiable in the positions and orientations of A and B. An additional requirement is that $\zeta(A, B)$ be defined and easy to compute for *all* positions and orientations of the particles. We will also make use of an overlap potential $\upsilon(A, B)$ for the case when ellipsoid A is completely contained within B (for example, B can be the bounding neighborhood of A, or it can be an ellipsoidal hard-wall container),

$$\begin{cases} \upsilon(A,B) > 0 \text{ if } A \text{ is completely contained in } B \\ \upsilon(A,B) = 0 \text{ if } A \text{ is internally tangent to } B \\ \upsilon(A,B) < 0 \text{ if part or all of } A \text{ is outside } B, \end{cases}$$

and give such a potential below. Such potentials have not been considered before since they do not appear in other algorithms, however, we will make extensive use of them.

First, we focus on the computation of $\zeta(A, B)$. We define and compute the overlap conditions using a procedure originally developed for ellipsoids by Perram and Wertheim [89]. The Perram and Wertheim (PW) overlap potential is defined through

$$\zeta = \mu^2 - 1 = \max_{0 \le \lambda \le 1} \min_{\mathbf{r}_C} \left[\lambda \zeta_A(\mathbf{r}_C) + (1-\lambda)\zeta_B(\mathbf{r}_C) \right].$$

For every multiplier λ, the solution of the inner optimization over $\mathbf{r}_C$ is unique due to the strict convexity of $\mathbf{r}_C$, and satisfies the gradient condition

$$\Delta\mathbf{n} = \lambda \mathbf{n}_A(\mathbf{r}_C) + (1-\lambda)\mathbf{n}_B(\mathbf{r}_C) = 0, \tag{2.1}$$

which shows that the normal vectors are parallel (with opposite directions). The solution of the outer optimization problem over λ is given through the condition

$$\zeta = \zeta_A(\mathbf{r}_C) = \zeta_B(\mathbf{r}_C), \tag{2.2}$$

which means that when the particles are rescaled by a common scaling factor $\mu = 1 + \Delta\mu = \sqrt{1+\zeta}$ they are in external tangency, sharing a common normal direction $\mathbf{n} = \mathbf{n}_A / \|\mathbf{n}_A\|$ (i.e., normalized to unit length and directed from A to B), and sharing a *contact point* $\mathbf{r}_C$. When focusing on one particle we can measure $\mathbf{r}_C$ with respect to the centroid of the particle, or otherwise specifically denote $\mathbf{r}_{AC} = \mathbf{r}_C - \mathbf{r}_A$ and $\mathbf{r}_{BC} = \mathbf{r}_C - \mathbf{r}_B$. This is illustrated for ellipses in Fig. 2.1. If the particles are touching then $\mu = 1$ and the procedure described above gives us the geometric contact point and therefore the common normal vector. In the case of spheres of radius O the PW overlap potential simply becomes

$$\zeta_{AB} = \frac{(\mathbf{r}_A - \mathbf{r}_B)^T (\mathbf{r}_A - \mathbf{r}_B)}{(O_A + O_B)^2} - 1 = \frac{l_{AB}^2}{(O_A + O_B)^2} - 1, \tag{2.3}$$

which avoids the use of square roots in calculating the distance between the centers of A and B, l_{AB}, and is easily manipulated analytically.

The geometrical idea behind the Perram-Wertheim overlap potential, namely, considering scaling the size of the ellipsoids uniformly until they are in external or internal tangency, can be generalized to different situations. It is worthwhile to walk through its definition in a more geometrical manner. Consider for example the case when A and B are disjoint, as illustrated in the leftmost part of Fig. 2.1. If ellipsoid A is scaled by a nonnegative factor $\mu(A)$ such that the centroid of B is still outside it, then there is a corresponding scaling

of B, $\mu(B)$, which brings B into external tangency with A at the *contact point* $\mathbf{r}_C[\mu(A)]$. This scaling is a solution to a simple eigenvalue-like problem involving $\mathbf{X}_A$ and $\mathbf{X}_B$. The normal vectors of A and B at the contact point are of opposite direction, and by changing the ratio of their lengths from 0 to ∞ we get a path of contact points going from the center of A to the center of B. It was a wonderful idea of Perram and Wertheim [89] to parameterize this path with a scalar $\lambda \in [0,1]$, and then look for the $\lambda = \Lambda$ which makes $\mu(A) = \mu(B)$, i.e. look for the common scaling factor μ_{AB} which brings A and B into external tangency at the contact point $\mathbf{r}_C$ (shown in Fig. 2.1), or equivalently, look for the largest common scaling factor which preserves non-overlap. This approach is very well-suited for the case when both A and B are particles and thus should be treated equally. Sometimes, however, ellipsoid B has a special status, for example, it may be the bounding neighborhood of another particle. In this case we look for the scaling factor $\mu_B(A)$ of A which brings A into external tangency with the fixed B (see second subsection in Section II.C in Ref. [90]), or equivalently, the largest scaling of A which preserves non-overlap, as illustrated in the middle part of Fig. 2.1. A similar idea applies to the case when A is contained within B, in which case we look for the largest scaling $\nu_B(A)$ of A which leaves A contained completely within B, or equivalently, which brings A into internal tangency with B.

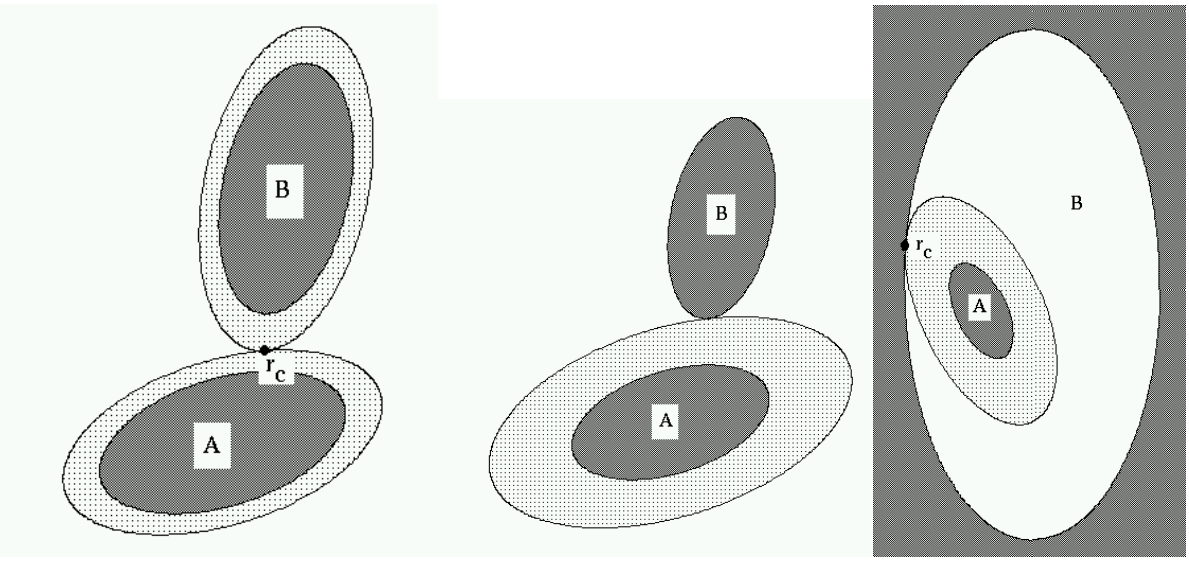

Figure 2.1: Illustration of the scaling μ in the PW contact function: *Left*: The outer tangency potential μ_{AB}. *Middle*: The outer tangency potential $\mu_B(A)$. *Right:* The inner tangency potential $\nu_B(A)$.

Using these scaling factors, we can define several *generalized* PW overlap potentials,

$$\zeta_{AB}(A,B) = \mu_{AB}^2 - 1 \tag{2.4}$$
$$\zeta_{AB}^{(B)}(A,B) = \mu_B^2(A) - 1 \tag{2.5}$$
$$\upsilon_B(A,B) = \nu_B^2(A) - 1, \tag{2.6}$$

which we will refer to as the Perram-Wertheim (PW), the *modified* PW overlap potential, and the *internal* PW overlap potential respectively. Note that $\zeta_{AB}^{(B)}$ and υ_B are not defined for all positions of the ellipsoids, namely, if the center of A is inside B, $\zeta_{AB}^{(B)}$ is not defined, and conversely, if the center of A is outside B, υ_B is not defined. Next we focus on evaluating $\zeta = \zeta_{AB}$ (we omit the subscripts for notational clarity) and its time and configurational derivatives for a general convex particle shape.

2.1.2 Evaluating ζ

Calculating the PW overlap potential can be done by solving for $\mathbf{r}_C(\lambda)$ in Eq. (2.1) and then looking for the λ that solves Eq. (2.2). A rigorous method to solve this problem is to continuously track the solution to Eq.

(2.1) as λ goes from 0 to 1, by solving the ODE obtained by differentiating Eq. (2.1) with respect to λ,

$$\mathbf{M}\frac{d\mathbf{r}_C}{d\lambda} = \left[\lambda\boldsymbol{\nabla}^2\zeta_A + (1-\lambda)\,\boldsymbol{\nabla}^2\zeta_B\right]\frac{d\mathbf{r}_C}{d\lambda} = \Delta\mathbf{g} = \boldsymbol{\nabla}\zeta_B - \boldsymbol{\nabla}\zeta_A,$$

with initial conditions $\mathbf{r}_C = \mathbf{r}_B$ and $\lambda = 0$. We can then use ODE event location [91] to stop when $\zeta_A = \zeta_B$. This method is rigorous in the sense that it can solve for the optimal λ to an arbitrary accuracy, however, it is inefficient since it requires solving an ODE. If a good-enough initial guess can be provided for λ, one can directly use Newton's method on this system of two equations (2.1). The Newton step is the determined by the step

$$\begin{aligned}\Delta\lambda &= \frac{1}{\zeta_{\lambda\lambda}}\left[(\zeta_B - \zeta_A) - \Delta\mathbf{g}^T\mathbf{M}^{-1}\Delta\mathbf{n}\right]\\ \Delta\mathbf{r}_C &= \mathbf{M}^{-1}\left(\Delta\mathbf{g}\Delta\lambda - \Delta\mathbf{n}\right),\end{aligned}$$

where $\zeta_{\lambda\lambda} = \Delta\mathbf{g}^T\mathbf{M}^{-1}\Delta\mathbf{g}$. There is no guarantees that this Newton method will converge, however, if it does, it will identify the unique contact point and the contact λ.

2.1.3 Time Derivatives of ζ

Other than the evaluation of ζ, many of the algorithms described in this work require the evaluation of time derivatives as the two particles in question move with certain (angular) velocities. The time derivatives necessary to track λ, $\mathbf{r}_C$ and ζ when predicting collisions between particles are given with

$$\begin{aligned}\dot{\lambda} &= \frac{1}{\zeta_{\lambda\lambda}}\left[\Delta\mathbf{g}^T\mathbf{M}^{-1}\Delta\mathbf{n}' - \left(\frac{\partial\zeta_B}{\partial t} - \frac{\partial\zeta_A}{\partial t}\right)\right]\\ \dot{\mathbf{r}}_C &= \mathbf{M}^{-1}\left(\Delta\mathbf{g}\frac{d\lambda}{dt} - \Delta\mathbf{n}'\right)\\ \dot{\zeta} &= \frac{\partial\zeta_A}{\partial t} + \left(\boldsymbol{\nabla}\zeta_A\right)^T\frac{d\mathbf{r}_C}{dt},\end{aligned}$$

where

$$\Delta\mathbf{n}' = \frac{\partial\Delta\mathbf{n}}{\partial t} = \lambda\frac{\partial\boldsymbol{\nabla}\zeta_A}{\partial t} + (1-\lambda)\frac{\partial\boldsymbol{\nabla}\zeta_B}{\partial t}$$

and we used the standard dot notation for time derivatives.

2.1.4 Configurational Derivatives of ζ

We will frequently need to consider derivatives of the overlap function with respect to the (generalized) positions of the particles, either first order,

$$\boldsymbol{\nabla}_{\mathbf{q}_i}\zeta = \boldsymbol{\nabla}_i\zeta = \left(\frac{\partial\zeta}{\partial q_i}\right),$$

or second order

$$\boldsymbol{\nabla}^2_{\mathbf{q}_i\mathbf{q}_j}\zeta = \boldsymbol{\nabla}^2_{ij}\zeta = \left[\frac{\partial^2\zeta}{\partial q_i\partial q_j}\right].$$

To first order, the particles can be replaced by their (parallel) tangent planes at the point of contact and the first order derivatives can be expressed in terms of quantities relating to the two tangent planes. To second order, the particles can be replaced by paraboloids that have the same tangent plane, as well as the same principal curvature axes and the same radii of curvatures as the two particles at the point of contact. It is therefore possible to derive general expressions for the derivatives in terms of quantities relating to the normal vectors and surface curvatures of the particles at the point of contact.

The first order derivatives can easily be expressed in terms of the position of the contact point $\mathbf{r}_C$ and the (normalized and outwardly-directed) contact normal vector $\mathbf{n}$. For this purpose, it is easier to measure the distance between two particles in near contact via the Euclidean *interparticle gap* h giving the (minimal)

surface-to-surface distance between the particles along the normal vector. Moving one of the particles by $\Delta\mathbf{q} = (\Delta\mathbf{r}, \Delta\boldsymbol{\varphi})$ displaces the contact point by $\Delta\mathbf{r}_C = \Delta\mathbf{r} + \Delta\boldsymbol{\varphi} \boxtimes \mathbf{r}_C$ and therefore changes the gap by $\Delta h = -\mathbf{n}^T \Delta\mathbf{r}_C = -\mathbf{n}^T \Delta\mathbf{r} - (\mathbf{r}_C \times \mathbf{n})^T \Delta\boldsymbol{\varphi}$, giving the gradient

$$\boldsymbol{\nabla}_{\mathbf{q}} h = -\begin{bmatrix} \mathbf{n} \\ \mathbf{r}_C \times \mathbf{n} \end{bmatrix}.$$

The relation between the (small) Euclidean gap h and the (small) gap as measured by the PW overlap potential ζ can be seen by observing that scaling an ellipsoid by a factor μ displaces the contact point by $\Delta\mathbf{r}_C = \Delta\mu\mathbf{r}_C$. Therefore, the scaling factor needed to close the interparticle gap is

$$\mu \approx \frac{\zeta}{2} \approx \frac{h}{(\mathbf{r}_{BC} - \mathbf{r}_{AC})^T \mathbf{n}} = \frac{h}{\mathbf{r}_{AB}^T \mathbf{n}},$$

giving the gradient of the overlap potential $\boldsymbol{\nabla}_{\mathbf{q}} \zeta = 2 (\boldsymbol{\nabla}_{\mathbf{q}} h) / \left(\mathbf{r}_{AB}^T \mathbf{n}\right)$,

$$\boldsymbol{\nabla}_{A/B} \zeta = \mp \frac{2}{\mathbf{r}_{AB}^T \mathbf{n}} \begin{bmatrix} \mathbf{n} \\ \mathbf{r}_{(A/B)C} \times \mathbf{n} \end{bmatrix}.$$

For spheres the cross product is identically zero and rotations can be eliminated from consideration.

The second-order derivatives are not as easily evaluated for a general particle shape. In two dimensions, or in three dimensions when the principal radii of curvatures at the point of contact are equal, one can replace the particle around the point of contact with a sphere of the appropriate position and radius. However, when the radii of curvatures are different this is not as easy to do. We will give explicit expressions for the second-order derivatives of ζ for ellipsoids in Section 2.2.4.2.

2.2 Hard Ellipsoids

In this section, we focus exclusively on ellipsoidal particles, particularly in two (ellipses) and three (ellipsoids) dimensions. We present all of the necessary tools to evaluate the (three) generalized PW overlap potentials and its derivatives for a pair of ellipsoidal particles. We first give some introductory material in Section 2.2.1, and then we then focus on calculation of these overlap functions and their time derivatives, which is used in Section 3.5 to robustly determine the time-of-collision for two moving ellipsoids. We will attempt to present most of the results so that they generalize to other dimensions as well; however, this is not always possible. Readers looking for more detailed background information are referred to Refs. [43] and [86].

2.2.1 Ellipsoids

An ellipsoid is a smooth convex body consisting of all points $\mathbf{r}$ that satisfy the quadratic inequality

$$(\mathbf{r} - \mathbf{r}_0)^T \mathbf{X} (\mathbf{r} - \mathbf{r}_0) \leq 1, \tag{2.7}$$

where $\mathbf{r}_0$ is the position of the center (centroid), and $\mathbf{X}$ is a characteristic *ellipsoid matrix* describing the shape and orientation of the ellipsoid. The case when $\mathbf{X} = \frac{1}{O^2}\mathbf{I}$ is a diagonal matrix describes a sphere of radius[1] O, which does not require orientation information. In the general case,

$$\mathbf{X} = \mathbf{Q}^T \mathbf{O}^{-2} \mathbf{Q}, \tag{2.8}$$

where $\mathbf{Q}$ is the rotational matrix describing the orientation of the ellipsoid, and $\mathbf{O}$ is a diagonal matrix containing the major semi-axes of the ellipsoid along the diagonal. The time derivative of the matrix (2.8) for an ellipsoid rotating with instantaneous angular velocity $\boldsymbol{\omega}$ is

$$\dot{\mathbf{X}} = \boldsymbol{\Omega}\mathbf{X} - \mathbf{X}\boldsymbol{\Omega}. \tag{2.9}$$

[1]We will use the letters r and R to denote positions of points, and therefore resort to using O when referring to radius, $\mathbf{o}$ when referring to the vector containing the semiaxes, and $\mathbf{O}$ for the diagonal matrix containing the semiaxes on its diagonal.

2.2.2 Ellipsoid Overlap Potentials

The problem of determining whether two ellipsoids A and B overlap (have a common point) or not has been considered previously in relation to Monte Carlo or MD simulations of hard-ellipsoid systems [43]. More than three decades ago, Vieillard-Baron proposed an overlap criterion based on the number of negative eigenvalues of a certain matrix [92], and this criterion has been subsequently rediscovered [93]. It easily generalizes to two dimensions and can be used to obtain an overlap potential. We have implemented and tested this overlap potential but have found it both computationally and theoretically inferior to the overlap potential proposed by Perram and Wertheim [89]. We have therefore completely adapted the PW overlap potential and also extended it to the case of one ellipsoid contained within another. Many other approaches are possible, for example, an approximate measure of the Euclidean distance between the surfaces of the two ellipsoids can be used [94, 95, 96, 97]. However, the advantage of the PW approach is its inherent symmetry, dimensionless character, and most of all, its simple geometric interpretation in terms of scaling factors.

Extensive use of all three of the generalized PW overlap potentials (2.4-2.6) has been made in the implementation of the algorithms described in this work. The original PW overlap potential (2.4) is the most efficient in practice and also has the property that it is symmetric with respect to the interchange of A and B, and is preferred over (2.5) unless $\mu_B^2(A)$ is needed. In this section we address the issue of efficiently and reliably calculating the three PW overlap potentials. We base our discussion on outlines of recipes for calculating ζ_{AB} and $\zeta_{AB}^{(B)}$ in the literature [89, 90, 43], but focus on detail and describe a specific computational scheme based on polynomials. Additionally, contact information such as the point of contact or the common normal vector at the point of contact can be calculated once the overlap potential is found.

2.2.2.1 Evaluating ζ_{AB}

As Perram and Wertheim observed, $\mathbf{r}_C(\lambda)$ can be determined analytically for ellipsoids (see below), and gives

$$\tilde{\zeta}_{AB}(\lambda) = \lambda\zeta_A(\mathbf{r}_C) + (1-\lambda)\zeta_B(\mathbf{r}_C) = \lambda(1-\lambda)\mathbf{r}_{AB}^T\mathbf{Y}^{-1}\mathbf{r}_{AB}, \tag{2.10}$$

where $\mathbf{r}_{AB} = \mathbf{r}_B - \mathbf{r}_A$, and

$$\mathbf{Y} = \lambda\mathbf{X}_B^{-1} + (1-\lambda)\mathbf{X}_A^{-1}. \tag{2.11}$$

It turns out that this function is strictly concave on the interval $[0,1]$ and thus has a unique maximum at $\lambda = \Lambda \in [0,1]$, from which one can directly calculate the overlap potential

$$\zeta_{AB} = \tilde{\zeta}_{AB}(\Lambda) = \max_{0\le\lambda\le1} \tilde{\zeta}_{AB}(\lambda).$$

The maximum of $\tilde{\zeta}_{AB}(\lambda)$ can easily be found numerically using only polynomial manipulations, by making extensive use of matrix adjoints (sometimes called adjugates) and determinants (both of which are polynomials in the matrix elements). First rewrite $\tilde{\zeta}_{AB}(\lambda)$ as a rational function:

$$\tilde{\zeta}_{AB}(\lambda) = \frac{p_{AB}(\lambda)}{q_{AB}(\lambda)} = \frac{\lambda(1-\lambda)\left\{\mathbf{a}_{AB}^T \mathrm{adj}\left[\lambda\mathbf{I} + (1-\lambda)\mathbf{A}_{AB}\right]\mathbf{a}_{AB}\right\}}{\det\left[\lambda\mathbf{I} + (1-\lambda)\mathbf{A}_{AB}\right]}, \tag{2.12}$$

where

$$\mathbf{a}_{AB} = \mathbf{X}_B^{1/2}\mathbf{r}_{AB} \text{ and } \mathbf{A}_{AB} = \mathbf{X}_B^{1/2}\mathbf{X}_A^{-1}\mathbf{X}_B^{1/2}.$$

Note that powers of $\mathbf{X}$ are easy to calculate because of the special form (2.8) and orthogonality of $\mathbf{Q}$. We have made use of the symbolic algebra system Maple® and its code generation abilities to generate inlined Fortran code to form the coefficients of the polynomial $\mathrm{adj}\left[\lambda\mathbf{I} + (1-\lambda)\mathbf{A}\right]$ and $\det\left[\lambda\mathbf{I} + (1-\lambda)\mathbf{A}\right]$ for a given symmetric matrix $\mathbf{A}$, and this has found numerous uses when dealing with ellipsoids, such as in evaluating the coefficients of the polynomials p_{AB} and q_{AB} in eq. (2.12). The unique maximum of $\tilde{\zeta}_{AB}(\lambda)$ can be found by finding the root of its first derivative, which is the same as finding the unique root of the degree-$2d$ polynomial

$$h_{AB} = p'_{AB}q_{AB} - p_{AB}q'_{AB}$$

in the interval $[0,1]$, which can be done very rapidly using a safeguarded Newton method.

A good initial guess to use in Newton's method is the exact result for spheres

$$\Lambda = \frac{\overline{O}_A}{\overline{O}_A + \overline{O}_B},$$

where $\overline{O}$ is the largest semiaxis, i.e., the radius of the enclosing sphere for an ellipsoid. Additionally, one often has a better initial guess for Λ in cases when the relative configuration of the ellipsoids has not changed much from previous evaluations of ζ_{AB}. Finally, a task which appears frequently is to evaluate the overlap potential between two ellipsoids but only if they are closer than a given cutoff, in the sense that the exact value is only needed if $\zeta_{AB} \leq \zeta_{AB}^{(\text{cutoff})}$, or equivalently $\mu_{AB} \leq \mu_{AB}^{(\text{cutoff})}$. This cutoff can be used to speed up the process by terminating the search for Λ as soon as a value $\tilde{\zeta}_{AB}(\lambda) > \zeta_{AB}^{(\text{cutoff})}$ is encountered during Newton's method. Additionally, one can first test the enclosing spheres for A and B with the same cutoff and not continue the calculation if the spheres are disjoint even when scaled by a factor $\mu_{AB}^{(\text{cutoff})}$.

Since almost always the value of $\lambda = \Lambda$ is used, henceforth we do not explicitly denote the special value Λ, unless there is the possibility for confusion. The reader should keep in mind that expressions to follow are to be evaluated at $\lambda = \Lambda$. The subscript C will be used to denote quantities pertaining to the contact point. The contact point $\mathbf{r}_C$ of the two ellipsoids is

$$\mathbf{r}_C = \mathbf{r}_A + (1-\lambda)\mathbf{X}_A^{-1}\mathbf{n} = \mathbf{r}_B - \lambda\mathbf{X}_B^{-1}\mathbf{n}, \tag{2.13}$$

where

$$\mathbf{n} = \mathbf{Y}^{-1}\mathbf{r}_{AB} \tag{2.14}$$

is the unnormalized common *normal vector* at the point of contact (once the ellipsoids are scaled by the common factor μ_{AB}), directed from A to B in this case. Here $\mathbf{r}_{BC} = \mathbf{r}_C - \mathbf{r}_B$ and $\mathbf{r}_{AC} = \mathbf{r}_C - \mathbf{r}_A$ are the "arms" from the centers of the ellipsoids to the contact point. An important value is the curvature of $\tilde{\zeta}_{AB}(\lambda)$ at the special point $\lambda = \Lambda$,

$$\zeta_{\lambda\lambda} = \frac{d^2\tilde{\zeta}_{AB}}{d\lambda^2} = 2\frac{\mathbf{r}_{BC}^T\mathbf{Y}^{-1}\mathbf{r}_{AC}}{\lambda(1-\lambda)} = -2\mathbf{n}^T\mathbf{Z}\mathbf{n} < 0,$$

where

$$\mathbf{Z} = \mathbf{X}_A^{-1}\mathbf{Y}^{-1}\mathbf{X}_B^{-1} = \mathbf{X}_B^{-1}\mathbf{Y}^{-1}\mathbf{X}_A^{-1} = \left[\lambda\mathbf{X}_A + (1-\lambda)\mathbf{X}_B\right]^{-1}.$$

2.2.2.2 Evaluating $\zeta_{AB}^{(B)}$ and υ_B

The evaluation of the modified outer and internal tangency PW overlap potentials $\zeta_{AB}^{(B)}$ and υ_B proceeds in a similar fashion, but with a differing sign in several expressions. Here the upper sign will denote the case of internal tangency (υ_B), and the lower the case of outer tangency ($\zeta_{AB}^{(B)}$). We proceed to give a prescription for evaluation of these potentials without detailed explanations.

Define the parameterized function

$$f_B(\lambda) = \lambda^2\mathbf{r}_{AB}^T\mathbf{Y}^{-1}\mathbf{X}_B^{-1}\mathbf{Y}^{-1}\mathbf{r}_{AB}, \tag{2.15}$$

as well as

$$g_B(\lambda) = (1-\lambda)^2\mathbf{r}_{AB}^T\mathbf{Y}^{-1}\mathbf{X}_A^{-1}\mathbf{Y}^{-1}\mathbf{r}_{AB}, \tag{2.16}$$

where

$$\mathbf{Y} = \lambda\mathbf{X}_B^{-1} \mp (1-\lambda)\mathbf{X}_A^{-1}. \tag{2.17}$$

We then numerically look for the largest $\lambda = \Lambda$ in $[0,1]$ which solves the nonlinear equation

$$f_B(\lambda) = 1, \tag{2.18}$$

and then we have the desired scaling factor

$$\upsilon_B \text{ or } \zeta_{AB}^{(B)} = g_B(\Lambda).$$

Additionally, the contact point is

$$\mathbf{r}_C = \mathbf{r}_A \mp (1-\lambda)\,\mathbf{X}_A^{-1}\mathbf{n} = \mathbf{r}_B \pm \lambda \mathbf{X}_B^{-1}\mathbf{n}, \tag{2.19}$$

where the normal vector $\mathbf{n}$ is as in eq. (2.14). An additional useful value is the slope

$$f_\lambda = \frac{df_B}{d\lambda} = 2\frac{\mathbf{r}_{BC}^T \mathbf{Y}^{-1}\mathbf{r}_{AC}}{(1-\lambda)}.$$

We can again use polynomial algebra to efficiently solve eq. (2.18) using a safeguarded Newton method, by rewriting $f_B(\lambda)$ as a rational function

$$f_B(\lambda) = \frac{\sum_{k=1}^{d}\left[\lambda p_k^{(B)}(\lambda)\right]^2}{q_B^2(\lambda)} = \frac{\lambda^2 \left\|\mathrm{adj}\left[\lambda\mathbf{I} \mp (1-\lambda)\,\mathbf{A}_{AB}\right]\mathbf{a}_{AB}\right\|^2}{\det\left[\lambda\mathbf{I} \mp (1-\lambda)\,\mathbf{A}_{AB}\right]} = 1, \tag{2.20}$$

where $p_k^{(B)}$ and q_B are polynomials, to obtain the equivalent equation

$$\frac{q_B(\lambda)}{\lambda\sqrt{\sum_{k=1}^{d}\left[p_k^{(B)}(\lambda)\right]^2}} = 1,$$

which is better suited for numerical solution. The search interval for Λ in this case should be taken to be $\left[\tilde{\Lambda}, 1\right]$, where $\tilde{\Lambda}$ is the largest root of the degree-d polynomial q_B in $[0,1]$, which can be found exactly in both two and three dimensions using standard algebraic methods for the solution of polynomial equations of degree less than 5. A reasonable initial guess when evaluating υ_B is $\lambda = \tilde{\Lambda}$.

Unlike the evaluation of ζ_{AB} and ζ_{AB}^B, which are both rapid[2] and robust, the evaluation of υ_B poses numerical difficulties due to the presence of the minus sign in eq. (2.17), which can cause $\mathbf{Y}$ to become singular. This happens when $\left\|\mathbf{Y}^{-1}\mathbf{r}_{AB}\right\| \to 0$, which does occur when $\Lambda \to \tilde{\Lambda}$. In this case, since $\mathbf{Y}$ is singular, its adjoint is (almost always) rank-1,

$$\mathrm{adj}\left[\mathbf{Y}\right] \to \mathbf{u}\mathbf{u}^T,$$

where $\mathbf{u}$ is some (eigen)vector, and the problem occurs because $\mathbf{u}^T\mathbf{r}_{AB} \to 0$, yielding an apparently indeterminate 0/0 in eq. (2.20). The limiting value of υ_B is mathematically well-defined even in this case, however, its numerical evaluation is unstable, and has been a constant source of numerical problems in our implementation. One alleviating trick is to avoid explicitly inverting $\mathbf{Y}$ and instead the adjoint should be used, $\mathbf{Y}^{-1} = \mathrm{adj}\left[\mathbf{Y}\right]/\det\left[\mathbf{Y}\right]$, where the determinant of $\mathbf{Y}$ can be calculated by using (2.18),

$$\det\left[\mathbf{Y}\right] = \lambda\sqrt{\tilde{\mathbf{n}}_1^T\mathbf{X}_B^{-1}\tilde{\mathbf{n}}_1},$$

where $\tilde{\mathbf{n}}_1 = \mathrm{adj}\left[\mathbf{Y}\right]\mathbf{r}_{AB}$. Even with such precautions, we have observed numerical difficulties in the calculations involving inner tangency of A and B. It would therefore be useful to explore alternative overlap potentials for the case when ellipsoid A is contained within ellipsoid B, or different ways of calculating υ_B.

2.2.3 Time Derivatives of the Overlap Potentials

When dealing with moving ellipsoids, and in particular, when determining the time-of-collision for two ellipsoids in motion, expressions for the time derivatives of the contact potentials are needed. We give these expressions here without a detailed derivation. We have additionally obtained expressions for second order derivatives, however, these are not needed for the current exposition and are significantly more complicated, and are not presented here.

[2] In our numerical experience ζ_{AB} and its time derivatives can be evaluated significantly faster, which is one of the reasons we have chosen to use it.

2.2.3.1 Derivatives of ζ_{AB}

Consider two ellipsoids moving with instantaneous velocities $\mathbf{v}_A$ and $\mathbf{v}_B$ and rotating with instantaneous angular velocities $\boldsymbol{\omega}_A$ and $\boldsymbol{\omega}_B$. For the purposes of the Lubachevsky-Stillinger algorithm, we also want to allow the ellipsoid semiaxes to change with an expansion/contraction rate $\dot{\mathbf{o}}$. We have the expected result that the rate of change of overlap depends on the *projection* of the relative velocity at the point of contact, $\mathbf{v}_C$, along the common normal vector $\mathbf{n}$,

$$\dot{\zeta}_{AB} = 2\lambda\left(1-\lambda\right)\mathbf{n}^T\mathbf{v}_C, \tag{2.21}$$

where

$$\mathbf{v}_C = \left[\mathbf{v}_B + \boldsymbol{\omega}_B \boxtimes \mathbf{r}_{BC} + \boldsymbol{\Gamma}_B\mathbf{r}_{BC}\right] - \left[\mathbf{v}_A + \boldsymbol{\omega}_A \boxtimes \mathbf{r}_{AC} + \boldsymbol{\Gamma}_A\mathbf{r}_{AC}\right]$$

is the relative velocity of the two ellipsoid surfaces at the point of contact (the *contact velocity*), and we denoted

$$\boldsymbol{\Gamma} = \mathbf{Q}^T\left(\mathbf{O}^{-1}\dot{\mathbf{O}}\right)\mathbf{Q}.$$

One sometimes also needs the time derivative of $\lambda = \Lambda$

$$\dot{\lambda} = -\frac{2}{\zeta_{\lambda\lambda}}\left\{\tilde{\mathbf{n}}_2^T\mathbf{v}_C + \mathbf{n}^T\left[\lambda\boldsymbol{\Gamma}_B\mathbf{r}_{BC} + \left(1-\lambda\right)\boldsymbol{\Gamma}_A\mathbf{r}_{AC}\right] + \lambda\left(1-\lambda\right)\mathbf{n}^T\mathbf{Z}\left[\left(\boldsymbol{\omega}_B - \boldsymbol{\omega}_A\right)\boxtimes\mathbf{n}\right]\right\}, \tag{2.22}$$

where

$$\tilde{\mathbf{n}}_2 = \lambda\mathbf{Y}^{-1}\mathbf{r}_{BC} + \left(1-\lambda\right)\mathbf{Y}^{-1}\mathbf{r}_{AC}.$$

2.2.3.2 Derivatives of $\zeta_{AB}^{(B)}$ and υ_B

For the generalized PW overlap potentials, the corresponding time derivatives are given with

$$\dot{\zeta}_{AB}^{(B)} \text{ or } \dot{\upsilon}_B = \mp 2\left(1-\lambda\right)\mathbf{n}^T\mathbf{v}_C, \tag{2.23}$$

and

$$\dot{\lambda} = -\frac{2\lambda}{f_\lambda}\left\{\mathbf{r}_{BC}^T\mathbf{Y}^{-1}\mathbf{v}_C + \left[\mathbf{r}_{AC}^T\mathbf{Y}^{-1}\boldsymbol{\Gamma}_B\mathbf{r}_{BC} - \mathbf{r}_{BC}^T\mathbf{Y}^{-1}\boldsymbol{\Gamma}_A\mathbf{r}_{AC}\right] \mp \lambda\left(1-\lambda\right)\mathbf{n}^T\mathbf{Z}\left[\left(\boldsymbol{\omega}_B - \boldsymbol{\omega}_A\right)\boxtimes\mathbf{n}\right]\right\}. \tag{2.24}$$

2.2.4 Configurational Derivatives of ζ_{AB}

In Section 2.1.4 we obtained the first-order derivatives of $\zeta = \zeta_{AB}$ with respect to the positions of the particles for general particle shapes. In this section we specialize these to ellipsoids and also give explicit expressions for the second-order derivatives, whose derivation is very tedious and therefore we only present the final results here.

2.2.4.1 First-Order Derivatives

In principle the overlap potential is a function of the normalized quaternions describing the particle orientations, and derivatives of ζ need to be projected onto the unit quaternion sphere, to produce

$$\boldsymbol{\nabla}_{\mathbf{r}_B}\zeta = 2\lambda\left(1-\lambda\right)\mathbf{n}$$

and

$$\begin{aligned}
\widetilde{\boldsymbol{\nabla}}_{s_B}\zeta &= -4\lambda\left(1-\lambda\right)\left[\mathbf{p}_B^T\left(\mathbf{r}_{BC}\times\mathbf{n}\right)\right] \\
\widetilde{\boldsymbol{\nabla}}_{\mathbf{p}_B}\zeta &= 4\lambda\left(1-\lambda\right)\left[s_B\left(\mathbf{r}_{BC}\times\mathbf{n}\right) - \mathbf{p}_B\times\left(\mathbf{r}_{BC}\times\mathbf{n}\right)\right],
\end{aligned}$$

and $\widetilde{\boldsymbol{\nabla}}$ denotes a projected gradient. Of course, similar expressions apply for derivatives with respect to $\mathbf{q}_A$, but with the sign of $\mathbf{n}$ reversed and also λ and $(1-\lambda)$ swapped.

This projection can be avoided if we do not do a traditional Taylor series in the quaternions, namely an additive perturbation $\Delta\mathbf{q}$, but rather consider a multiplicative perturbation to the quaternions in the form

of a small *rotation* from the current configuration $\Delta\boldsymbol{\varphi}$. The gradient of the overlap potential, which enters in the columns of the rigidity matrix, can be shown to be

$$\boldsymbol{\nabla}_B\zeta = -\boldsymbol{\nabla}_A\zeta = \begin{bmatrix} \boldsymbol{\nabla}_{\mathbf{r}_B}\zeta \\ \boldsymbol{\nabla}_{\boldsymbol{\varphi}_B}\zeta \end{bmatrix} = 2\lambda\left(1-\lambda\right)\begin{bmatrix} \mathbf{n} \\ \mathbf{r}_{BC}\times\mathbf{n} \end{bmatrix},$$

as we derived in Section 2.1.4 for a general convex particle shape by using the normalized normal vector $\hat{\mathbf{n}} = \mathbf{n}/\left\|\mathbf{n}\right\|$ [note that $\zeta = \lambda\left(1-\lambda\right)\mathbf{r}_{AB}^T\mathbf{n} - 1 = 0$]. Additionally, it is useful to know the derivatives of λ,

$$\begin{aligned} \boldsymbol{\nabla}_{\mathbf{r}_B}\lambda &= -\frac{2}{\zeta_{\lambda\lambda}}\tilde{\mathbf{n}}_3 \\ \boldsymbol{\nabla}_{\boldsymbol{\varphi}_B}\lambda &= -\frac{2}{f_{\lambda\lambda}}\left[\mathbf{M}_B\mathbf{n}_A - \lambda\left(\mathbf{r}_{BC}\times\mathbf{n}\right)\right], \end{aligned}$$

where

$$\tilde{\mathbf{n}}_3 = \lambda\mathbf{n}_B + \left(1-\lambda\right)\mathbf{n}_A = \lambda\mathbf{Y}^{-1}\mathbf{r}_{AC} + (1-\lambda)\mathbf{Y}^{-1}\mathbf{r}_{BC},$$

and

$$\mathbf{M}_B = \lambda\mathbf{N}^L\mathbf{X}_B^{-1} + \mathbf{R}_{CB}^L.$$

In these equations capitalizations of vectors such as $\mathbf{R}_{CB}^L = \left|\mathbf{r}_{\mathbf{C}}\right|_\times^L$ denote cross-product matrices.

2.2.4.2 Second-Order Derivatives

The explicit expressions for the Hessian of the overlap potential are

$$\boldsymbol{\nabla}_{\mathbf{r}_B}^2\zeta = 2\lambda\left(1-\lambda\right)\mathbf{Y}^{-1} - \frac{4}{f_{\lambda\lambda}}\left(\tilde{\mathbf{n}}_3\tilde{\mathbf{n}}_3^T\right) \succ \mathbf{0}$$

$$\boldsymbol{\nabla}_{\boldsymbol{\varphi}_B\mathbf{r}_B}^2\zeta = 2\lambda\left(1-\lambda\right)\mathbf{M}_B\mathbf{Y}^{-1} + 2\left[\left(\boldsymbol{\nabla}_{\boldsymbol{\varphi}_B}\lambda\right)\tilde{\mathbf{n}}_3^T\right]$$

and finally

$$\begin{aligned} \boldsymbol{\nabla}_{\boldsymbol{\varphi}_B}^2\zeta &= -f_{\lambda\lambda}\left[\left(\boldsymbol{\nabla}_{\boldsymbol{\varphi}_B}\lambda\right)\left(\boldsymbol{\nabla}_{\boldsymbol{\varphi}_B}\lambda\right)^T\right] + 2\lambda\left(1-\lambda\right)\cdot \\ &\quad\{\quad\left[\frac{1}{2}\left(\mathbf{r}_{BC}\mathbf{n}^T + \mathbf{n}\mathbf{r}_{BC}^T\right) - \left(\mathbf{r}_{BC}^T\mathbf{n}\right)\mathbf{I}\right] + \\ &\qquad\lambda\mathbf{N}^L\mathbf{X}_B^{-1}\mathbf{N}^R + \mathbf{M}_B\mathbf{Y}^{-1}\mathbf{M}_B^T\}. \end{aligned}$$

The derivatives with respect to the position and orientation of particle A can be obtained by simply exchanging the roles of particles A and B, however, there are also mixed derivatives involving motion of both particles

$$\begin{aligned} \boldsymbol{\nabla}_{\boldsymbol{\varphi}_B\mathbf{r}_A}^2\zeta &= -\boldsymbol{\nabla}_{\boldsymbol{\varphi}_B\mathbf{r}_B}^2\zeta \\ \boldsymbol{\nabla}_{\boldsymbol{\varphi}_A\mathbf{r}_B}^2\zeta &= -\boldsymbol{\nabla}_{\boldsymbol{\varphi}_A\mathbf{r}_A}^2\zeta \\ \boldsymbol{\nabla}_{\boldsymbol{\varphi}_B\boldsymbol{\varphi}_A}^2\zeta &= -\boldsymbol{\nabla}_{\boldsymbol{\varphi}_B}^2\zeta + \left(\boldsymbol{\nabla}_{\boldsymbol{\varphi}_B\mathbf{r}_B}^2\zeta\right)\mathbf{R}_{AB}^R - \frac{1}{2}\left|\boldsymbol{\nabla}_{\boldsymbol{\varphi}_\mathbf{B}}\zeta\right|_\times. \end{aligned}$$

For efficiency, it may be preferred to express the second-order derivatives for particle A directly from those for particle B,

$$\begin{aligned} \boldsymbol{\nabla}_{\boldsymbol{\phi}_A}^2\zeta &= \boldsymbol{\nabla}_{\boldsymbol{\varphi}_B}^2\zeta - \mathbf{R}_{AB}^L\left(\boldsymbol{\nabla}_{\mathbf{r}_B}^2\zeta\right)\mathbf{R}_{AB}^R + \mathrm{Symm}\left[\left|\boldsymbol{\nabla}_{\mathbf{r}_\mathbf{B}}\mathbf{F}\right|_\times^L\mathbf{R}_{AB}^R - 2\left(\boldsymbol{\nabla}_{\boldsymbol{\varphi}_B\mathbf{r}_B}^2\zeta\right)\mathbf{R}_{AB}^R\right] \\ \boldsymbol{\nabla}_{\boldsymbol{\varphi}_A\mathbf{r}_B}^2\zeta &= -\boldsymbol{\nabla}_{\boldsymbol{\varphi}_B\mathbf{r}_B}^2\zeta - \mathbf{R}_{AB}^L\left(\boldsymbol{\nabla}_{\mathbf{r}_B}^2\zeta\right) + \left|\boldsymbol{\nabla}_{\mathbf{r}_\mathbf{B}}\mathbf{F}\right|_\times^L, \end{aligned}$$

where Symm denotes the operation of symmetrizing a matrix, $\mathrm{Symm}(A) = (A + A^T)/2$.

2.2.5 Surface-Surface Distance

A quantity that is also useful to derive is the true surface-to-surface gap between ellipsoids h, however, an exact expression for it is difficult to calculate even numerically [98]. When the two ellipsoids are close to touching, an estimate of h can be obtained as described in Section 2.1.4. A better, geometrically-motivated, estimate may be obtained by actually calculating the distance between the tangent planes for the ellipsoids that are parallel to the common tangent plane at $\mathbf{r}_C$,

$$h_{\text{lb}} = \mathbf{r}_{AB}^T\widehat{\mathbf{n}} - \sqrt{\widehat{\mathbf{n}}^T\mathbf{X}_A^{-1}\widehat{\mathbf{n}}} - \sqrt{\widehat{\mathbf{n}}^T\mathbf{X}_B^{-1}\widehat{\mathbf{n}}} = \mathbf{r}_{AB}^T\widehat{\mathbf{n}}\left(1-\frac{1}{\mu}\right)$$

which can be shown to be a lower bound on the true surface-to-surface distance. Here we denoted the *unit* normal vector at the point with $\widehat{\mathbf{n}} = \mathbf{n}/\left\|\mathbf{n}\right\|$. An upper bound is the smallest distance between the two surfaces along the direction of $\mathbf{r}_{AB}$, $h_{\text{ub}} = \left\|\mathbf{r}_{AB}\right\|\left(1-\frac{1}{\mu}\right)$ [99].

In Section 2.1.4 we obtained the derivatives of h or equivalently ζ with respect to the positions of the particles, without considering the potential change in particle shape. Uniform scaling of the particle size is easy to add in the general case, however, for ellipsoids one can consider arbitrary changes of shape, so long as the particles remain ellipsoidal. To first order, the change in the gap as one of the ellipsoids changes its axes by $\Delta\mathbf{o}$, is

$$\Delta h = \widehat{\mathbf{n}}^T\left(\mathbf{\Gamma}\Delta t\right)\mathbf{r}_C = \widehat{\mathbf{n}}^T\mathbf{Q}^T\left(\mathbf{O}^{-1}\Delta\mathbf{O}\right)\mathbf{Q}\mathbf{r}_C = \tilde{\mathbf{a}}^T\Delta\mathbf{o},$$

where

$$\tilde{\mathbf{a}} = \boldsymbol{\nabla}_{\mathbf{o}}h = \left(\mathbf{O}^{-1}\mathbf{Q}\widehat{\mathbf{n}}\right)\diamond\left(\mathbf{Q}\mathbf{r}_C\right),$$

and $\diamond$ denotes an element-wise product of two vectors. Of particular interest is the case when the ellipsoids just get scaled in size without changing their shape, $\Delta\mathbf{o} = \delta\mathbf{o}$, which leads to $\Delta h = \left(\widehat{\mathbf{n}}^T\mathbf{r}_C\right)\delta$ from each of the two ellipsoids, leading to a total gap change of

$$\Delta h = \left(\widehat{\mathbf{n}}^T\mathbf{r}_{AB}\right)\delta.$$

2.2.5.1 Flat Walls

As a side note, it is sometimes also useful to consider evaluating the distance between planes (flat walls) and ellipsoids. Consider the tangent plane to an ellipsoid with normal vector $\mathbf{n}$, touching the ellipsoid at the point $\mathbf{r}_C$. We have that

$$\mathbf{r}_C = -\frac{\mathbf{X}^{-1}\mathbf{n}}{d},$$

where $d = \sqrt{\mathbf{n}^T\mathbf{X}^{-1}\mathbf{n}}$ is the distance from the center to the tangent plane. A suitable overlap potential for a plane with normal vector $\mathbf{n}$ a distance d_0 away from the ellipsoid is $\zeta_{\text{plane}} = d_0 - d$, and its time derivative is simply

$$\frac{d\zeta_{\text{plane}}}{dt} = -\mathbf{n}^T\left(\mathbf{v} - \boldsymbol{\omega}\boxtimes\mathbf{r}_C\right).$$

This can be used, for example, to implement molecular dynamics for hard ellipsoids inside a container with hard walls.

2.2.6 Two Near-Spheres (Nearly) Touching

In what follows we will need first-order approximations of the impenetrability constraints between two nearly spherical ellipsoids. Assume there are two spheres A and B of radius $O_{A/B}$ touching. Transform the spheres into ellipsoids with semiaxes $O\mathbf{I} + \Delta\mathbf{O}$, and orientation described by the rotation matrix $\mathbf{Q}$, and denote $\boldsymbol{\epsilon}_O = O^{-1}\Delta\mathbf{O}$. Finally, define the matrix

$$\mathbf{T} = \mathbf{Q}^T\boldsymbol{\epsilon}_O\mathbf{Q},$$

which in the case of turning a disk into an an ellipse with semiaxes O and $O(1-\epsilon)$, i.e., aspect ratio $\alpha = 1+\epsilon$, $\epsilon \ll 1$, becomes

$$\mathbf{T} = -\epsilon\begin{bmatrix} \sin^2\phi & -\sin\phi\cos\phi \\ -\sin\phi\cos\phi & \cos^2\phi \end{bmatrix} = -\epsilon\mathbf{T}_\phi,$$

where θ is the angle of orientation of the ellipse. It can be shown that to first order in ϵ the new distance between the ellipsoids is

$$\Delta\zeta = 2\mathbf{u}_{AB}^T \mathbf{S} \mathbf{u}_{AB}, \tag{2.25}$$

where

$$\mathbf{S} = \frac{O_A}{O_A + O_B}\mathbf{T}_A + \frac{O_B}{O_A + O_B}\mathbf{T}_B.$$

Additionally, the perturbed values of λ, the normal vector $\mathbf{n}$, and the contact vector $\mathbf{r}_C$, for the case $O_A = O_B = O$, are:

$$\begin{aligned} \lambda &\approx \frac{1}{2} - \frac{1}{2}\mathbf{u}_{AB}^T \tilde{\mathbf{S}} \mathbf{u}_{AB} \\ \frac{O}{2}\mathbf{n} &\approx \mathbf{u}_{AB} - 2\mathbf{S}\mathbf{u}_{AB} \\ \frac{\mathbf{r}_{BC}}{O} \text{ or } \frac{\mathbf{r}_{AC}}{O} &\approx \mp\mathbf{u}_{AB} - 2\tilde{\mathbf{S}}\mathbf{u}_{AB} + \left(\mathbf{u}_{AB}^T \tilde{\mathbf{S}} \mathbf{u}_{AB}\right)\mathbf{u}_{AB} \end{aligned}$$

where $\tilde{\mathbf{S}} = \frac{1}{2}(\mathbf{T}_B - \mathbf{T}_A)$.

2.3 Generalized Ellipsoids

Ellipsoids can be generalized to obtain a very versatile particle shape that can interpolate between a sphere and a particle with almost flat faces and edges or sharp corners such as a cylinder or cube. This generalization of quadric shapes known as superquadrics (superellipsoids) have important uses in computer graphics [100, 101], and have also been employed in material modeling as a more general particle shape than spheres and ellipsoids. In this section we give the necessary expressions to handle this particle shape in molecular dynamics or Monte Carlo hard-particle algorithms.

2.3.1 Superellipsoids

Superellipsoids [100, 101] are a spherically-symmetric generalization of ellipsoids given by the shape function

$$\zeta(\mathbf{r}) = g\left[\tilde{\zeta}(\tilde{\mathbf{r}})\right] - 1,$$

where $\tilde{\mathbf{r}} = \mathbf{O}^{-1}\mathbf{Q}(\mathbf{r} - \mathbf{r}_0)$ is the relative position rotated and scaled according to the orientation and shape of the (super)ellipsoid. A single-exponent variant of superellipsoids (one can use two exponents in three dimensions [101]) uses

$$g(x) = x^{1/\epsilon}$$

while

$$\tilde{\zeta}(\tilde{\mathbf{r}}) = \mathbf{e}^T\mathbf{p} = \mathbf{e}^T f(\tilde{\mathbf{r}}).$$

Here f is an element-wise function acting on each of the components of $\tilde{\mathbf{r}}$, $f(x) = |x|^{2\epsilon}$, with derivatives $\mathbf{u} = f'(\tilde{\mathbf{r}})$ and $\mathbf{z} = f''(\tilde{\mathbf{r}})$. The *superexponent* $\epsilon \geq 1$ determines the degree of roundedness of the edges and faces, $\epsilon = 1$ being plain ellipsoids (spheres if $\mathbf{o} = O$), and $\epsilon \to \infty$ being a rectangular prism (cubes if $\mathbf{o} = O$). In three dimensions, it is possible to produce more asymmetric shapes such as cylinders by using two exponents as in

$$\tilde{\zeta}(\tilde{\mathbf{r}}) = \left(r_1^{2\epsilon_1} + r_1^{2\epsilon_1}\right)^{\epsilon_2/\epsilon_1} + r_3^{2\epsilon_2},$$

which only slightly complicates the derivatives below. Note that powers of the form $|x|^\epsilon$ always produce centrally symmetric shapes. To get shapes like triangles or tetrahedra one can add a term of the form $(x + |x|)^\epsilon$ for each of the planar faces/edges, where x is the normal distance to the face/edge (note that this is no longer a strictly convex shape!).

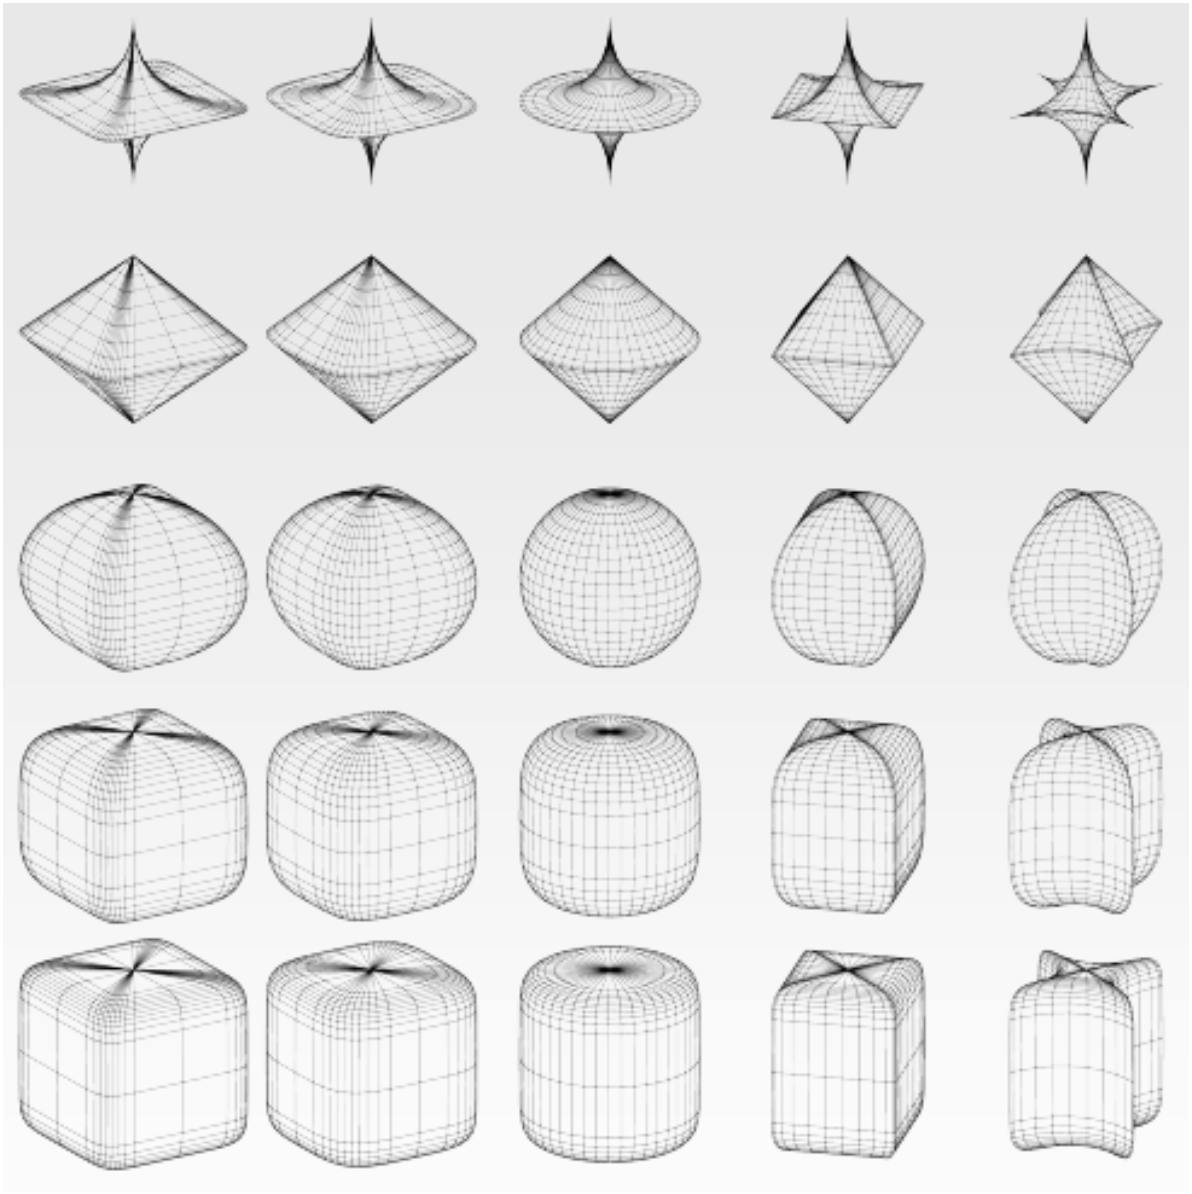

Figure 2.2: Example of superellipsoid shapes that can be produced by changing the exponents ϵ_1 and ϵ_2. Some of these are not convex, and the convex shapes interpolate between a sphere and a cube, a cylinder, a double pyramids, or a double cone. Flat edges and sharp corners are obtained in the limit $\epsilon \to \infty$, however, for finite superexponents the shapes are smooth [Couresy Gordon Kindlmann, see `http://teem.sourceforge.net`].

2.3.2 Geometric Properties

A superellipsoid can be enclosed with a sphere of radius $O_{\max} = \|\mathbf{o}\|$, although one can use a smaller sphere with extra computation, especially close to the sphere shape. The volume of a superellipsoid is

$$V = \left(\prod o_i\right) K,$$

where

$$K_2 = \frac{2}{\epsilon}\mathcal{B}\left(\frac{1}{2\epsilon}, \frac{2\epsilon+1}{2\epsilon}\right) \text{ for } d = 2$$

and

$$K_3 = K_2(\epsilon_1)\frac{1}{\epsilon_2}\mathcal{B}\left(\frac{1}{2\epsilon_2}, \frac{\epsilon_2+1}{\epsilon_2}\right) \text{ for } d = 3,$$

and $\mathcal{B}(x,y) = \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)}$ is the beta function [101].

2.3.3 Evaluating the Overlap Function

The necessary derivatives to evaluate the PW overlap potential as described in Section 2.1.2 are

$$\boldsymbol{\nabla}\zeta = g'(\tilde{F})\left(\boldsymbol{\nabla}\tilde{F}\right)$$

and

$$\boldsymbol{\nabla}^2\zeta = g'(\tilde{\zeta})\left(\boldsymbol{\nabla}^2\tilde{\zeta}\right) + g''(\tilde{\zeta})\left(\boldsymbol{\nabla}\tilde{\zeta}\right)\left(\boldsymbol{\nabla}\tilde{\zeta}\right)^T.$$

Here

$$\boldsymbol{\nabla}\tilde{\zeta} = \left(\mathbf{O}^{-1}\mathbf{Q}\right)^T\mathbf{u}$$

and

$$\boldsymbol{\nabla}^2\tilde{F} = \left(\mathbf{O}^{-1}\mathbf{Q}\right)^T\mathbf{Z}\left(\mathbf{O}^{-1}\mathbf{Q}\right).$$

We implemented the direct Newton method instead of the ODE approach [c.f. Section 2.1.2], since it was simple to add to the existing ellipsoid implementation. It was however necessary to allow the specification of initial guesses for λ *and* $\mathbf{r}_C$ in order to ensure convergence of the Newton method and increase computational efficiency. We found the Newton approach relatively robust for reasonably small values of $\epsilon < 10$. For higher values of the exponent alternative implementations of the overlap potentials would be necessary. We have not yet implemented the inner overlap potential v_B for superellipsoids.

2.3.4 Time Derivatives

The time derivatives used in Section 2.1.3, given a linear motion in time with velocities $\mathbf{v}$, angular velocity $\boldsymbol{\omega}$ and expansion rates $\boldsymbol{\Gamma} = \mathbf{Q}^T\left(\mathbf{O}^{-1}\dot{\mathbf{O}}\right)\mathbf{Q}$, are

$$\begin{aligned}\frac{\partial\zeta}{\partial t} &= -\left(\boldsymbol{\nabla}\zeta\right)^T\mathbf{v}_C \\ \frac{\partial\boldsymbol{\nabla}\zeta}{\partial t} &= \boldsymbol{\omega}\boxtimes\left(\boldsymbol{\nabla}\zeta\right) - \boldsymbol{\Gamma}\left(\boldsymbol{\nabla}\zeta\right) - \left(\boldsymbol{\nabla}^2\zeta\right)\mathbf{v}_C,\end{aligned}$$

where

$$\mathbf{v}_C = \mathbf{v} + \boldsymbol{\omega}\boxtimes\left(\mathbf{r}-\mathbf{r}_0\right) + \boldsymbol{\Gamma}\left(\mathbf{r}-\mathbf{r}_0\right)$$

is the velocity of the surface of the particle at the point of contact.

2.4 Conclusions

In this Chapter we demonstrated that the overlap potential originally developed by Perram and Wertheim for hard ellipsoids can be generalized to any smooth strictly convex shape. We developed in detail the necessary tools for hard ellipsoids and superellipsoids, and also presented explicit formulas for time and configurational derivatives, as used in subsequent chapters. Additionally, we showed how to generalize the PW overlap potential to the case of a particle contained inside another (hard-wall container). This potential, however, exhibits numerical problems and developing a better inner overlap potential is an important future task.

Chapter 3

Molecular Dynamics for Nonspherical Particles

In this chapter we present in considerable detail a collision-driven molecular dynamics algorithm for a system of nonspherical particles, within a parallelepiped simulation domain, under both periodic or hard-wall boundary conditions [11]. The algorithm extends previous event-driven molecular dynamics algorithms for spheres, and is most efficient when applied to systems of particles with relatively small aspect ratios and with small variations in size. We present a novel partial-update near-neighbor list (NNL) algorithm that is superior to previous algorithms at high densities, without compromising the correctness of the algorithm. This efficiency of the algorithm is further increased for systems of very aspherical particles by using bounding sphere complexes (BSC). These techniques will be useful in any particle-based simulation, including Monte Carlo and time-driven molecular dynamics. Additionally, we allow for a non-vanishing rate of deformation of the boundary, which can be used to model macroscopic strain and also alleviate boundary effects for small systems. We then apply the algorithm to systems of hard ellipses and ellipsoids. We describe an algorithm for predicting the time of collision for two moving ellipses or ellipsoids. We present performance results for our implementation of the algorithm, demonstrating that for dense systems of very aspherical ellipsoids the novel techniques of using neighbor lists and bounding sphere complexes, offer as much as two orders of magnitude improvement in efficiency over direct adaptations of traditional event-driven molecular dynamics algorithms.

3.1 Introduction

Classical molecular dynamics (MD) simulations have been used to study the properties of particle systems for many decades. The first MD studies used the simplest multi-particle system, the hard-sphere fluid/solid [102], which is very rich in behavior. Subsequently, methods were developed to follow the dynamics of a system of soft spheres, i.e., particles interacting with a spherically symmetric and continuous interparticle potential, usually with a hard cutoff on the range of the interaction. The algorithms needed to simulate the two types of systems are rather different, and difficult to combine [103]. For soft particles, one needs to integrate a system of ordinary differential equations given by Newton's law of motion. However, for hard particles the interaction potential is singular and the task of integrating the equations of motion becomes a problem of processing a sequence of binary collisions between the particles[1], or collisions of the particles with the hard walls of a container, if any. In other words, for hard particles, one needs to predict and process a sequence of discrete *events* of vanishing duration.

The algorithm for hard particles therefore becomes *event-driven*, as opposed to the *time-driven* algorithm for soft-particle MD in which time changes in small steps and the equations of motion are integrated. Event-driven algorithms have the task of scheduling a sequence of events predicted to happen in the future. The simulation is advanced to the time of the event with the smallest scheduled time (the *impending event*) and that event is processed. The schedule of events is updated if necessary and the same process is repeated.

[1]Multi-particle collisions have zero probability of occurring and will not be considered here.

In molecular dynamics, the primary kind of event are binary collisions, so the simulation becomes *collision-driven*. This kind of collision-driven approach was used for the very first MD simulation of the hard-disk system [102], and has since been extended and improved in a variety of ways, most importantly, to increase the efficiency of the algorithm. All of these improvements, namely, delayed particle updates, the cell method, the use of a *heap* for the event queue, etc., appear in almost any efficient hard-particle MD algorithm. Systems of many thousands of hard disks or spheres can be studied on a modern personal computer using such algorithms, and in the past decade the method has been extended to handle more complex simulations, such as particles in a velocity field [104].

However, we are aware of only one collision-driven simulation in the literature for nonspherical particles, namely, one for very thin rods (needles) [105]. Other molecular dynamics simulations for hard nonspherical particles have used a time-driven approach [106], which is simpler to implement than the event-driven approach but inferior in both accuracy and efficiency (at high densities). Two kinds of smooth shapes are used frequently to model aspherical particles: *spherocylinders* (a cylinder with two spherical caps), and *ellipsoids*. Both can become spheres in a suitable limit. Spherocylinders are analytically much simpler then ellipsoids; however, they are always axisymmetric and cannot be oblate.

The primary reason event-driven algorithms have not yet been used for nonspherical particles is that a high-accuracy collision-driven scheme for nontrivial particle shapes and sufficiently large systems is very demanding. Computers have only recently reached the necessary speeds, and a proper implementation involves a significant level of code complexity. In this work, we present in detail a collision-driven molecular dynamics algorithm for a system of hard nonspherical particles. The algorithm is based on previous event-driven MD approaches for spheres, and in particular the algorithms of Lubachevsky [37] and Sigurgeirsson *et al.* [106]. While in principle the algorithm is applicable to any particle shape, we have specifically tailored it for smooth particles for which it is possible to introduce and easily evaluate continuously differentiable *overlap potentials*. Additionally, it is assumed that the particles have a spherically symmetric moment of inertia, so that in-between collisions their angular velocities are constant. Furthermore, the algorithm is most efficient when applied to relatively dense and homogeneous systems of particles which do not differ widely in size (i.e., the degree of polydispersity is small). We focus in this work on systems with *lattice-based* boundaries, under both periodic or hard-wall boundary conditions. The main innovations and strengths of the proposed algorithm are:

- It specifically allows for non-axially symmetric particles by using quaternions in the representation of orientational degrees of freedom, unlike previous hard-particle algorithms which have been restricted just to needles, spherocylinders or spheroids;

- The particle-shape-dependent components of the algorithm are clearly separated from general concepts, so that it is (at least in principle) easy to adapt the algorithm to different particle shapes; we present in detail the implementation of these components for ellipses and ellipsoids;

- It explicitly allows for time-dependent particle shapes and for time-dependent shape of the boundary cell, which enables a range of nonequilibrium applications, such as the generation of packings using the Lubachevsky-Stillinger algorithm, and also Parinello-Rahman-like [107] constant-pressure molecular dynamics;

- It corrects some assumptions in traditional hard-sphere algorithms that are not correct for nonspherical particles or when boundary deformation is included, such as the *nearest image convention* in periodic systems and the claim that there must be an intervening collision between successive collisions of a given pair of particles;

- It is the first rigorous event-driven MD algorithm to incorporate *near-neighbor lists*, by using the concept of *bounding neighborhoods*. This is a very significant improvement for very aspherical particles and/or at high densities, and has some advantages over the traditional *cell method* even for hard spheres because it allows a close monitoring of the collision history of the algorithm;

- It is the first algorithm to specifically address the problem of efficient near-neighbor search for very elongated or very flat particles by introducing the concept of *bounding sphere complexes*. The algorithm also clearly separates neighbor-search in a static environment (where particle positions are fixed) from

its use in a dynamic environment (where particles move continuously), thus enabling one to easily incorporate additional neighbor search techniques. We emphasize that the developed near neighbor search techniques will improve all particle-based simulations, including Monte Carlo and time-driven molecular dynamics; and

- It is documented in detail with pseudo-codes which closely follow the actual Fortran 95 code used to implement it for ellipses and ellipsoids.

One motivation for developing this algorithm has been to extend the Lubachevsky-Stillinger sphere-packing algorithm [12, 13] to nonspherical particles. We have successfully used our implementation to obtain many interesting results for random and ordered packings of ellipses and ellipsoids [78]. We numerically demonstrate that our novel neighbor-search techniques can speed the simulation by as much as two orders of magnitude or more at high densities and/or for very aspherical particles, as compared to direct adaptations of traditional hard-sphere schemes. Executable versions of the Fortran 95 codes implementing the algorithms described in this chapter, for systems of spheres and ellipsoids in both two and three dimensions and with arbitrary dispersity, can be downloaded at `http://cherrypit.princeton.edu/Packing/Fortran`. Dimension-independent C++ source codes for a simplified version of the collision driven algorithm we describe here, for monodisperse systems of hard spheres, can be downloaded at `http://cherrypit.princeton.edu/Packing/C++`.

We begin by presenting preliminary information and the basic ideas behind the algorithm in Section. 3.2. We then focus on the important task of improving the efficiency of the algorithm by focusing on neighbor search in Section 3.3, and present both the classical cell method and our adaptation of near-neighbor lists. Detailed pseudocodes for all major steps in the algorithm for general nonspherical particles are given in Section 3.4. We then describe the missing particle-shape-dependent pieces of the algorithm in Section 3.5. Some performance results for the algorithm are shown in Section 3.6, particularly focusing on the use of our near-neighbor list and bounding sphere complexes techniques.

3.2 Preliminaries

In this section we give some background information and a preliminary description of the algorithm. First, we discuss the impact the shape of the particles has on the algorithm. Then we briefly describe the two main approaches to hard-particle molecular dynamics, time-driven and event-driven. Finally, we discuss boundary conditions in our event-driven molecular dynamics algorithm and also the possibility of performing event-driven MD in different ensembles. Bold symbols are reserved for vectors and matrices, and subscripts are used to denote their components. Matrix multiplication is assumed whenever products of matrices or a matrix and a vector appear. Subscripts or superscripts are used heavily to add specificity to various quantities, for example, $\mathbf{r}$ denotes position, while $\mathbf{r}_A$ denotes the position of some particle A. We denote the numerical precision with $\epsilon \ll 1$, and use subscripted ϵ's for various user-set (small) numerical tolerances. We often omit explicit functional dependencies when they are clearly implied by the context and it is not important to emphasize them, for example, f and $f(t)$ will be used interchangeably.

3.2.1 Particle Shape

We consider a system of N hard particles whose only interactions are given by impenetrability constraints, although it is easy to allow for additional external fields which are independent of the particles (such as gravity). Many of the techniques developed here are also used to deal with particles interacting with a soft potential if there is a hard cutoff on the potential. We discuss the special case of orientation-less particles, namely spheres, at length, and we use hard ellipsoids to illustrate the extensions to nonspherical particles. We will use the terms sphere and ellipsoid in any dimension, but sometimes we will be more specific and distinguish between disk and ellipse in two dimensions, and sphere and ellipsoid in three dimensions.

Spheres are a very important special case not only because of their simplicity, but also because *bounding spheres* are a necessary ingredient when dealing with aspherical particles. A bounding sphere for a particle is centered at the *centroid* of the particle and has the minimal possible diameter $D_{\max} = 2O_{\max}$ so that it fully encloses the particle itself. Here by centroid we mean a geometrically special point chosen so that the bounding sphere is as small as possible (i.e., it should be chosen to be as close as possible to the midpoint

of the longest line segment joining two points of the particle). For example, for an ellipsoid, the bounding sphere has the same center as the ellipsoid and its diameter is equal to the largest axes of the ellipsoid. The importance of bounding spheres is that they provide a quick and analytically simple way to test for overlap of two particles: Two particles cannot overlap if their bounding spheres do not overlap. Occasionally we make use of *contained spheres*, which are also centered at the centroid of the particle and have the maximal possible diameter $D_{\min} = 2O_{\min}$ so that they are fully within the particle itself. For ellipsoids their diameter is the smallest axes. Note that two particles must overlap if their contained spheres overlap. The efficiency of the EDMD algorithm described in this work is primarily determined by the *aspect ratio* $\alpha = D_{\max}/D_{\min}$. The greater the deviation of α from unity, the worse the efficiency because the bounding/contained spheres become worse approximations for the particles and because of the increasing importance of particle orientations. We propose a novel *near-neighbor list* technique for dealing with very aspherical particles, in addition to the standard *cell method*.

3.2.2 Molecular Dynamics

The goal of our algorithm is to simulate the motion of the particles in time as efficiently as possible, while taking into account the interactions between the particles. For hard-particle systems, the only interactions occur during *binary collisions* of the particles. The goal of hard-particle *molecular dynamics* (MD) algorithms is to correctly predict the time-ordered sequence of particle collisions. Additionally, there may be obstacles such as hard walls with which the particles can collide. Next we briefly introduce the main ideas behind the two main approaches to hard-particle MD, time-driven and event-driven MD. This preliminary presentation will be helpful in understanding the rest of this section. Further details on the event-driven algorithm are given in Section 3.4.

3.2.2.1 Time Driven MD

The *Time-Driven Molecular Dynamics* (TDMD) approach is inspired by MD simulations of systems of soft particles (i.e., particles interacting with a continuous interaction potential). It has been adapted also to the simulation of hard-particle systems, particularly nonspherical particles [106, 108], mainly because of its simplicity. In this approach, all of the particles are displaced *synchronously* in small *time steps* Δt and a check for overlap between the particles is done. If any two particles overlap, time is rolled back until the approximate moment of initial overlap, i.e., the time of collision, and the collision of the particles is processed (i.e., the momenta of the colliding particles are updated), and the simulation continued. The main disadvantage of this approach is that it is *not* rigorous, in the sense that collisions may be missed or the correct ordering of a sequence of successive collisions may be mis-predicted (particularly in dense systems). To ensure a reasonably correct prediction of the system dynamics, a very small time step must be used and this is inefficient. Nonetheless, since only checking for overlap between particles is needed, the simplicity of the method is a very attractive feature. Additionally, such an approach is parallelizable with the same techniques as any other MD algorithm (for example, domain decomposition).

3.2.2.2 Event Driven MD

An alternative *rigorous* approach is to use *Event-Driven Molecular Dynamics* (EDMD), based on a rather general model of discrete event-driven simulation. In EDMD, instead of advancing time independently of the particles as in TDMD, time is advanced from one *event* to the next event, where an event is a binary particle collision, or a collision of a particle with an obstacle (hard wall). Other types of events will be discussed shortly, however, collisions are the central type of event so we label the approach more specifically as *Collision-Driven Molecular Dynamics* (CDMD). We will however continue to use the abbreviation EDMD since the term event-driven is widely used in the literature.

Efficient implementations of EDMD are *asynchronous*: each particle is at the point in time when the last event involving it happened. Each particle predicts what its *impending event* is and when it is expected to happen. All of these events are entered into a priority *event queue* (typically implemented by a heap), which allows for quick extraction of the next event to happen. The positions and momenta of the particles involved in this event are updated, the particles' next event predicted, the event queue updated, and the simulation continued with the next event. Sometimes events may be mis-predicted. For example, a particle

i may predict a collision with particle j, but another (third party) particle m may collide with j before i has time to. A special event called a *check* needs to be introduced, and it amounts to simply (re-)predicting the impending event for a given particle. Given infinite numerical precision, this kind of approach rigorously follows the dynamics of the system.

The computationally expensive step in EDMD is the prediction of the impending event of a given particle i (even though asymptotically the event-queue operations dominate). This typically involves the expensive (especially for nonspherical particles) step of predicting the time of collision between the particle i and a set of other particles j. In the simplest approach, one would predict the time of collision between i and all other particles and choose the smallest one, but a much more efficient approach is described in Section 3.3. For spheres moving along straight lines, predicting the time of collision merely amounts to finding the first positive root (if any) of a quadratic equation, and is very fast. Therefore, for spherical particles EDMD always outperforms TDMD by orders of magnitude, *and* it is rigorous. For nonspherical particles, collision predictions are much more involved, but for algebraically simple smooth particle shapes it is expected that EDMD will still outperform TDMD for a wide range of densities. Furthermore, there are systems for which TDMD is not possible, and one must use EDMD, such as systems of hard line segments [105]. Note however that the efficiency of the EDMD approach is possible only because the motion of the particles between events can be predicted *a priori*, and because binary collisions only affect the two colliding particles. In cases when these assumptions are not true, TDMD may be the only option. Additionally, it is very important to note that the EDMD algorithm is inherently non-parallelizable due its sequential processing of the events. Some attempts have been made to parallelize the method [109] by using the locality of the interactions, and very recently actual implementations have appeared [110]. We will defer any discussion of parallelization to future publications.

3.2.3 Boundary Conditions

In this work we consider MD in a simple bounded simulation domain embedded in a Euclidean space E^d of dimensionality d. In particular, we focus exclusively on *lattice-based boundaries*. This means that the simulation domain, which the particles never leave, is a *parallelepiped* defined by d *lattice vectors*, $\boldsymbol{\lambda}_1, \ldots, \boldsymbol{\lambda}_d$. The simulation domain, or *unit cell*, is a collection of points with d *relative coordinates* $\mathbf{r}$ in the interval $[0, 1]$, and corresponding Cartesian coordinates

$$\mathbf{r}^{(E)} = \sum_{k=1}^{d} r_k \boldsymbol{\lambda}_k = \boldsymbol{\Lambda}\mathbf{r}, \tag{3.1}$$

where $\boldsymbol{\Lambda}$ is a square invertible matrix representing the *lattice*, and contains the lattice vectors as columns. The volume of the unit cell is given by the positive determinant $|\boldsymbol{\Lambda}|$. As illustrated in Fig. 3.1(a), the parameters describing the geometry of a lattice-based boundary separate into components along different "dimensions", meaning along different lattice vectors. For example, there are d perpendicular distances L_k between the "left" and "right" faces of the parallelepiped along $\boldsymbol{\lambda}_k$ (meaning the two $(d-1)$-dimensional faces spanned by all lattice vectors other than $\boldsymbol{\lambda}_k$), one for each dimension $k = 1, \ldots, d$. We assume that in three dimensions the lattice vectors form a right-handed coordinate system, so each lattice vector can be identified as defining the x ($k = 1$), y ($k = 2$), or z ($k = 3$) axis.

Additionally, we allow either *periodic* or *hard-wall* boundary conditions (BCs) to be specified *independently* along each dimension, that is, the "left" and "right" faces of the unit cell along each dimension can either be hard walls or periodic boundaries. The most commonly used BCs are fully periodic, and one can interpret periodic systems as being on a topological torus whose distance geometry is determined by the metric tensor $\mathbf{G} = \boldsymbol{\Lambda}^T\boldsymbol{\Lambda}$ (a "flat" torus), or one can interpret periodic systems as being infinite and covering all of Euclidean space with identical copies of the unit cell and the particles in this unit cell. We will refer to the particles in the unit cell as *original particles* and simply identify them with an integer $i = 1, \ldots, N$. There are infinitely many *image* or *virtual particles* for every original particle, translated from the original by an integer number of lattice vectors $\mathbf{n}_c \in \mathcal{Z}$. We identify such an image of particle i with a pair of integers (i, v), where the *image* or *virtual identifier* v denotes the particular image in question, $\mathbf{n}_c = \mathbf{n}_c(v)$. Traditional hard-sphere algorithms have used the so-called nearest image convention, which assumes that only one (easy-to-identify) image of a given particle j can overlap particle i and thus $\mathbf{n}_c$ need not be explicitly

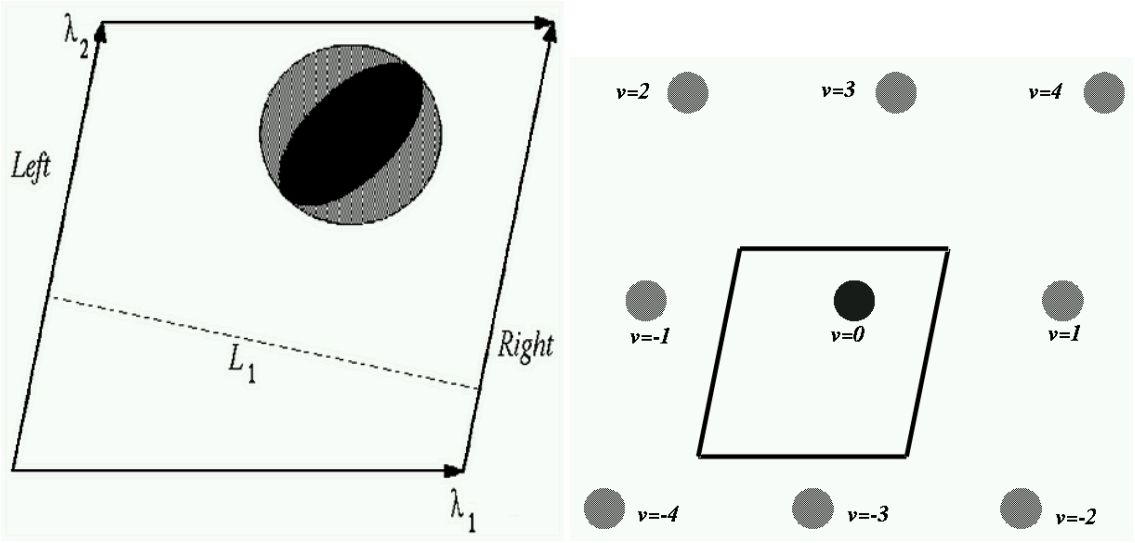

(a) Lattice-based boundary

(b) Periodic BCs

Figure 3.1: Illustration of lattice-based boundaries. The left subfigure shows a unit cell in two dimensions, along with the length of the unit cell along the x direction, L_1, and the left and right "walls" along the x dimension. Also shown is a particle and its bounding sphere (disk). The right subfigure shows a unit cell and an original particle (black), along with the first neighbor images and their image identifiers v.

stored. However, this assumption is wrong for nonspherical particles, where several images of a given particle j can overlap an original particle i.

One almost never needs to worry about any but the 3^d (9 in two, and 27 in three dimensions) images of the unit cell that neighbor the original one (first neighbors, including the cell under consideration). We number these images with $v = -(3^d - 1)/2, ..., (3^d - 1)/2$, so that images with opposite $\mathbf{n}_c$'s have identifiers of equal magnitude but opposite sign, $\mathbf{n}_c(-v) = -\mathbf{n}_c(v)$, as illustrated in two dimensions in Fig. 3.1(b). Note that $(i, 0) \equiv i$. When considering *ordered* pairs of particles $[i, (j, v)]$, one of the particles, i, is always original, while the other one, (j, v), can be an image or an original. Due to our choice of image-numbering, we can alternatively consider every such *unordered* pair to be composed of particles j and $(i, -v)$, $\{i, (j, v)\} \equiv \{j, (i, -v)\}$.

3.2.3.1 Boundary Deformation

In our simulations, we use *relative* coordinates $\mathbf{r}$ (i.e., expressed in terms of the lattice vectors), and likewise relative velocities $\mathbf{v}$, and convert these into their *Euclidean* representations $\mathbf{r}^{(E)}$ and $\mathbf{v}^{(E)}$ when necessary. The relative position of an image particle is $\mathbf{r} + \mathbf{n}_c$. The conversion between the two representations adds computational overhead due the matrix-vector multiplication[2] in (3.1). Some of this overhead can be avoided by only using Euclidean positions, however, we have chosen to express all positions relative to the lattice. The primary reason for this choice is that we allow the lattice to deform, that is, we allow for a *lattice velocity* $\dot{\mathbf{\Lambda}}$. In our algorithms, the lattice can change *linearly* with time,

$$\Delta\mathbf{\Lambda} = \dot{\mathbf{\Lambda}}\Delta t,$$

even though a more correct approach is to have a constant *strain rate*

$$\dot{\boldsymbol{\epsilon}} = \dot{\mathbf{\Lambda}}\mathbf{\Lambda}^{-1}, \tag{3.2}$$

[2]Note however that when the lattice is $\mathbf{\Lambda} = \mathbf{I}$, which is the usual choice unless a special unit cell is needed, all the matrix-vector multiplications become trivial. Our implementation has a special mode for this simple but important case.

that is, to have an exponential time evolution of the lattice,

$$\mathbf{\Lambda}(t) = \exp\left(\dot{\boldsymbol{\epsilon}} t\right) \mathbf{\Lambda}(0).$$

The identification of $\boldsymbol{\epsilon} = (\Delta\mathbf{\Lambda})\,\mathbf{\Lambda}^{-1}$ with the macroscopic strain [111] is explained in Section 4.4.5.1, and we choose it to be symmetric, $\dot{\boldsymbol{\epsilon}} = \dot{\boldsymbol{\epsilon}}^T$, to eliminate rotations of the unit cell.

In our approach, since the positions of the particles are relative to the lattice, the particles move together with the lattice. This is necessary in order to simulate isotropic systems. Namely, had the positions of the particles been independent of the lattice and the lattice deformed, for example, uniformly contracted, the image particles would move with the lattice, but the originals would not, and this would lead to artificial effects at the boundary of the unit cell. However, using relative positions is not without a cost. Consider a particle at relative position $\mathbf{r}$ which moves with constant relative velocity $\mathbf{v}$. It's Euclidean position is a parabola,

$$\mathbf{r}^{(E)}(t) = \left(\mathbf{\Lambda} + \dot{\mathbf{\Lambda}}t\right)(\mathbf{r} + \mathbf{v}t) = \mathbf{\Lambda}\mathbf{r} + \left(\mathbf{\Lambda}\mathbf{v} + \dot{\mathbf{\Lambda}}\mathbf{r}\right)t + \left(2\dot{\mathbf{\Lambda}}\mathbf{v}\right)\frac{t^2}{2}, \tag{3.3}$$

rather than a straight line. We can identify the instantaneous Euclidean position and velocity as well as the acceleration to be

$$\mathbf{r}^{(E)} = \mathbf{\Lambda}\mathbf{r} \tag{3.4}$$

$$\mathbf{v}^{(E)} = \mathbf{\Lambda}\mathbf{v} + \dot{\mathbf{\Lambda}}\mathbf{r} \tag{3.5}$$

$$\mathbf{a}^{(E)} = 2\dot{\mathbf{\Lambda}}\mathbf{v}. \tag{3.6}$$

This complicates, for example, the calculation of the time of collision of two moving spherical particles. Ordinarily a *quadratic* equation needs to be solved, but when the lattice deforms, a *quartic* equation needs to be solved instead. To our knowledge, our algorithm is the first EDMD algorithm to include a deforming boundary.

In our algorithm, the lattice velocity is an *externally* imposed quantity, and our goal is to simulate the motion of the particles as the boundary deforms, for example, in order to study shear banding in systems of ellipsoids [112]. It is the usual case that the boundary deforms slowly compared to the motion of the particles. As the boundary deforms, the unit cell becomes less and less orthogonal, and so in long-time simulations some form of orthogonalization of the unit cell might be necessary (we have not experimented with such techniques). Previously used constant-shear MD techniques [113] do not have this problem, however, they are also not capable of simulating arbitrary shears and are plagued with boundary effects.

Unlike in TDMD, in EDMD it is *not* possible to couple the motion of the boundary to all of the particles, as is done in Parinello-Rahman MD [107]. This is because the efficiency of the method depends critically on the fact that particles move independently between collisions and that particle collisions only affect the colliding particles. However, a pseudo-PRMD approach is possible, in which the lattice velocity is updated after a certain number of particle collisions, and the simulation is essentially restarted with a new lattice velocity. We have some preliminary positive experience with such a method, and in particular, we have used it to implement MD that maintains and approximately constant pressure as the particles change shape. Unit cell dynamics is also needed to properly study anisotropic liquids, which is very important for nonspherical particles [114]. A deforming boundary can also be used to model macroscopic strain in multiscale simulations (for example, to simulate granular flow) in which microscale MD is used to obtain material properties needed for a macroscopic continuum simulation.

3.2.4 EDMD in Different Ensembles

Molecular dynamics is often performed in ensembles different from the NVE one, and in particular, constant temperature and constant pressure are often desired. For this purpose, various thermostats have been developed. However, these are usually designed to be used with time-driven MD and systems of soft particles. We are in fact aware of no work that explicitly discusses thermostats for event-driven MD.

3.2.4.1 Iso-Temperature Simulations

Hard particle systems are inherently athermal due to the lack of energy scale, and the pressure and time scaling are therefore arbitrary. Sophisticated temperature or pressure thermostats are thus not usually

needed. In particular, simple *velocity rescaling* can be used to keep the temperature at the desired value. The average translational kinetic energy E_k per particle can be calculated and then *both* the translational and angular velocities scaled by the factor $s = \sqrt{dk_BT/2E_k}$ and the simulation essentially restarted[3]. This kind of temperature control is needed when, for example, the particles grow or shrink in size, since this leads to nonconservative collision dynamics and an overall heating or cooling of the system. Although velocity rescaling is simple and convenient, it has serious deficiencies [115], and a true canonical thermostat may be needed in some applications. For this purpose, an Andersen thermostat [116] can be included in collision-driven algorithms by considering the thermostat as a possible collision partner (with an appropriate Poisson distribution of collision times). We do not include such a stochastic temperature thermostat in our algorithm explicitly, since we have not used it.

3.2.4.2 Iso-Stress Simulations

A Parinello-Rahman-like isostress (isopressure) thermostat [107] cannot be directly included in collision-driven algorithms, since it implicitly couples the motion of all particles via the deformation of the unit cell and thus destroys the asynchronous efficiency of collision-driven approaches. Such a thermostat is often needed, even for hard particles, in simulations of crystal phases in order to keep the internal stress tensor isotropic and to allow for changes of the crystal unit cell. We have used constant shear boundary deformations, as described in Section 3.2.3.1, to implement a partial isostress thermostat in which the shear rate is periodically updated to reflect the asymmetry of the stress tensor.

More specifically, the MD algorithm obtains an estimate of the collisional contribution to the internal *stress tensor* $\boldsymbol{\sigma}_c$ [107] by averaging over all of the collisions between particles i and j during a time interval Δt,

$$\boldsymbol{\sigma}_c = \frac{1}{V\Delta t}\left\langle \mathbf{r}_{ij}\Delta\mathbf{p}_{ij}^T \right\rangle,$$

where $\mathbf{r}_{ij}$ is the center-to-center distance between the colliding particles, $\Delta\mathbf{p}_{ij}$ is the momentum exchange between the particles during the collision (the *impulse* [117]), and V is the volume of the unit cell. The interpretation of this expression as a stress tensor is discussed in the jamming limit in Sections 4.4.5.3 and 4.4.6.2. Note that there is an additional term in the stress tensor due to temperature [107], however, at high densities, and especially in the jamming limit, the collisional contribution dominates. For an isotropic system at applied external pressure P, the internal stress tensor is $\boldsymbol{\sigma} = P\mathbf{I}$.

The goal of the isostress MD is to keep the collisional stress as close to $P_0\mathbf{I}$ as possible, by changing the lattice vectors in order to accommodate for the particles. We have used a simple scheme in which the strain rate $\dot{\boldsymbol{\epsilon}}$ is updated periodically to be proportional to the stress imbalance (that is, the lattice deforms under the influence of the internal stress),

$$\dot{\boldsymbol{\Lambda}} = \mathbf{W}^{-1}(\boldsymbol{\sigma}_c - P_0\mathbf{I})\boldsymbol{\Lambda},$$

where $\mathbf{W}$ is the mass-matrix for the unit cell, usually chosen to be a multiple of the identity matrix. If the pressure itself is not fixed externally, for example, if it grows due to an increase in the density, then the target pressure can be set to be $P_0 = \text{Trace}(\boldsymbol{\sigma}_c)/d$, so that the MD simulation will try to keep the off-diagonal components of the internal stress tensor to be zero, i.e., it will try to keep the internal stress isotropic 4.4.5.3. This approach has had a mixed success and additional work is required to improve it, especially for anisotropic systems of (aspherical) particles [114].

3.3 Speeding Up the Search for Neighbors

Identifying the near-neighbors of a given particle has the most important impact on efficiency in almost all simulations of particle systems, particularly when the interparticle interactions are short-range. In both MC and TDMD algorithms it is important to quickly identify only the particles that are within the interaction cutoff distance l_{cutoff} (here distance is measured in the metric appropriate for the interaction) from a given particle and only evaluate the force or interaction energy with these particles. Since the number of such near neighbors is typically a small constant (of the order of $5-20$, strongly increasing with increasing

[3]Equipartition of energy is usually maintained by the collision dynamics, and therefore we usually do not use a different scaling for the angular velocities.

dimensionality d), this ensures that the computational effort needed to evaluate the forces on the particles or potential energy scales *linearly* with the number of particles N, as opposed to the quadratic complexity of checking all pairs of particles. In EDMD algorithms, it is useless to predict collisions between all pairs of particles since only nearby particles are actually likely to collide, and in fact $N \log N$ scaling can be obtained in EDMD by only predicting collisions between a given particle and a bounded (small) number of near neighbors.

In this section we describe the traditional cell method for speeding up neighbor search, and its implementation for lattice-based boundaries. We then propose a novel method based on the familiar concept of near-neighbor lists, which offers significant computational savings over the pure cell method for very aspherical particles, as we demonstrate numerically.

3.3.1 The Cell Method

One traditional method for neighbor search in particle systems is the so-called *cell method* (see for example Ref. [86]). It consists of partitioning the simulation domain into N_c disjoint cells and maintaining for each cell a list of all the particles whose centroids are within it. Then, for a given particle i, only the particles j in the neighboring cells (including periodic images of cells) of i's cell are considered neighbors of particle i. The shape of the cells can be chosen arbitrarily, so long as the union of all cells covers the whole simulation domain, and so long as for any given cell c one can (easily) identify all *neighbor cells* c_n that contain a point within Euclidean distance l_{cutoff} from a point in c. This enables a *rigorous* identification of all particles whose centroids are within a given cutoff distance from the centroid of a given particle. The essential aspect of the cell method is that the partitioning into cells is independent of the motion of the particles, so that even as the particles move one can continue to rely on using the cells to rigorously identify neighbors in constant time. This is a unique and necessary strength of the cell method, and *all* of our simulations use the cell method in some form, to ensure correctness while maintaining efficiency.

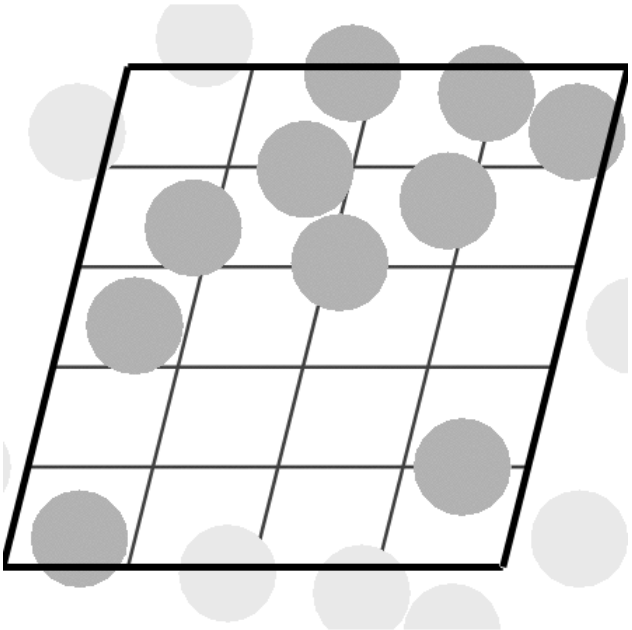

Figure 3.2: *The cell method*: A small disk packing and the associated grid of cells, to be used in searching for possibly overlapping particles. Dark-shaded disks are original particles and light-shaded ones are images.

For maximal efficiency, it seems that it is best to choose the cells as small as possible, but ensuring that only cells which actually share a boundary (i.e., are adjacent) need to be considered as neighbors. While this is obvious for MC or TDMD simulations, it is not so obvious for EDMD. For event-driven algorithms, it can be theoretically predicted and verified computationally that it is best to choose the number of cells to

be of the order of the number of particles [104], and computational experiments suggest that there should be about one particle per cell. For moderately to very dense systems, one should therefore choose the cells so that the maximal Euclidean distance between two points in the same cell, L_c, is as close to the largest enclosing sphere diameter $D_{\max}$ as possible,

$$L_c = (1 + \epsilon_L)\, D_{\max}.$$

We have verified this choice to be optimal in our extensive computational experience and consistently try to maximize the number of cells in all our simulations.

We note that in some simulations the shape of the particles changes. For example, the Lubachevsky-Stillinger algorithm [12] generates dense packings of particles by performing an EDMD simulation while the particles expand uniformly (for example, for spheres the radius O changes linearly with time with a constant expansion rate γ, $O(t) = O + \gamma t$). In such cases one must ensure that a sphere of diameter $D_{\max}$ can always enclose any of the particles in the system. A suitable value for $D_{\max}$ can be found, for example, by assuming that the final packing is the densest possible (or fills space completely if an exact result for the maximal density is not known). It is also important to note that it is sometimes needed to find all particles whose centroids are within a distance larger then $D_{\max}$ from the centroid of a given particle. This is not a problem for the cell method, as one can simply include as many additional cells in the search as needed to guarantee that the search is rigorous. For example, one may need to include neighbor cells of the neighbor cells (i.e., second-neighbor cells).

The need to adjust the partitioning of Euclidean space into cells to the shape of the simulation domain is the most difficult aspect of using the cell-method. Most simulations in the literature have been done with spherical particles and in cubic simulation domains, and the partitioning of the simulation domain is a simple Cartesian grid (mesh) of cells, where each cell is a cube (this is probably an optimal shape of the cells). For other boundary shapes, one has two options:

1. Continue using a partitioning of Euclidean space that is independent of the shape of the boundary. This would likely involve enclosing the simulation domain with a cube and then partitioning the cube into cells (some cells would be outside the domain and thus wasted). It is even possible to use the cell method with an infinite simulation domain if hashing techniques are employed [118].
2. Use a cell shape that conforms to the shape of the boundary in some simple way. The shape of the cells will thus change if the boundary deforms during the simulation.

Both options have their pros and cons. It may not be possible to use the second one for very complex boundary shapes. We have used the first approach to generate packings of ellipsoids in a spherical container (useful in comparing with experimental results), by enclosing the spherical container in a cube and partitioning that cube into cells. However, for lattice-based boundaries, we have chosen to use a partitioning of the unit cell into a possibly non-orthogonal Cartesian grid that conforms to the shape of the unit cell. This is illustrated in two dimensions in Fig. 3.3. The unit cell is partitioned into $N_k^{(c)}$ slabs along each dimension $k = 1, \ldots, d$, to obtain a total of $N_c = \prod_{k=1}^{d} N_k^{(c)}$ identical and consecutively (first along dimension 1, then along dimension 2, etc.) numbered parallelepiped cells. We typically maximize $N_k^{(c)}$ along each dimension such that the distance between the two parallel faces of the cells along any dimension is larger then the extent of the largest particle, $L_c = \min_k L_k > D_{\max}$. Operations in the cell method for lattice-based boundaries basically remain operations on Cartesian grids, just as if the simulation domain had been cubic. Note that each cell has 3^d neighbors (including itself), which is only 9 in two, but 27 in three dimensions. As noted earlier, sometimes more then 3 slabs may need to be checked along certain dimensions, depending on the Euclidean cutoff distance for the neighbor search. For completely periodic boundary conditions, there are other schemes for partitioning into cells which preserve the orthogonality and compactness of the unit cell [119], by using alternative choices of the simulation domain. In simulations where the lattice deforms by large amounts, one can alternatively periodically recompute a well-conditioned basis for the lattice and restart the simulation with a new choice of lattice vectors.

3.3.1.1 The Cell Method in EDMD

It is useful to briefly sketch the basic ideas of how the cell method is integrated in event-driven algorithms. The cell partitioning is used to speed up the prediction of the next event to happen. This event may be a

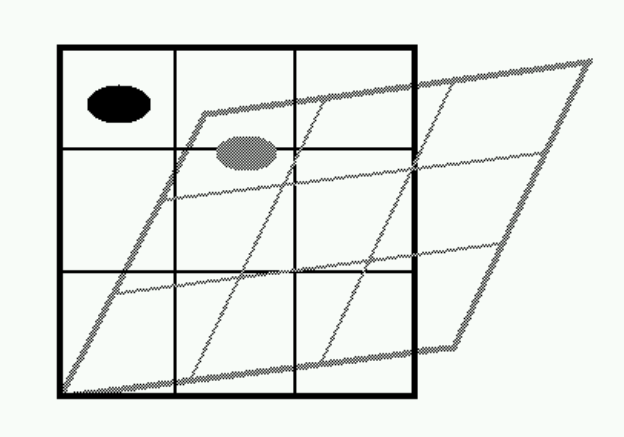

Figure 3.3: The partitioning of a lattice-based simulation box into cells (dark lines). The Cartesian grid of cells deforms in unison with the lattice, as illustrated by a snapshot of the box and its partitioning at a latter time. Particles also move together with the lattice, even if they are at "rest", $\mathbf{v} = 0$, as shown for an ellipsoid in the $(1, 3)$ cell.

boundary event, which can be a collision of a particle with a hard-wall, or a particle leaving the unit cell in a periodic system. The event may also be a binary collision, and each particle predicts collisions only with the particles in the (first) neighboring cells of its current cell. It is clear that a binary collision cannot occur with a particle not in a neighboring cell until the particle leaves its current cell. Therefore, a boundary event may be a *transfer*, where a particle's centroid leaves its current cell and goes into another cell. The algorithm predicts and processes transfers in time order with the other events.

Whenever a particle undergoes a transfer, it must correct its event prediction. In particular, it must predict binary collisions with all the particles in the *new* neighbor cells. If one maintains separately the prediction for the next boundary *and* the next binary collision, then upon a transfer one can reuse the old binary collision prediction and *only* calculate the collision time with particles in the neighbor cells which were not checked earlier [104]. This cuts the number of neighbor cells to process from 3^d to 3^{d-1}, which can save up to $2/3$ in computational effort in three dimensions. In our algorithm we originally maintained the binary collision prediction separately and reused it whenever possible, however, since the neighbor-list method described next is usually superior in practice, we no longer try to reuse previous binary collisions.

3.3.2 The Near-Neighbor List (NNLs) Method

The cell method is the method used in all EDMD algorithms that we are aware of. An exception is the algorithm of Ref. [120], but this algorithm is rather different from the classical EDMD algorithms (and from our algorithm) in more than this respect. There is a preference for the cell method in EDMD because it is very easy to incorporate it into the algorithm, while still maintaining a rigorously provable correct execution of the event sequence, given sufficient numerical precision. For *monodispersed* (equal) spherical particles, particularly at moderate densities, the cell method is truly the best approach. However, for aspherical particles whose aspect ratio is far from 1, the cell method becomes inefficient. This is because one cannot choose the cells small enough to ensure an average of about 1 particle per cell. Instead, due to the large $D_{\max}$, there need to be very few (large) cells which contain many particles and so little computational effort is saved by using the cells. The same is true even for spheres when large polydispersity is present since the cells need to be at least as large as the largest sphere in the system, and therefore there can be many small spheres inside one cell. A more complicated hierarchical cell structure (quadtree or octree) can be used for very polydisperse packings, but such an approach does not directly generalize to nonspherical particles.

In TDMD, a more widely used neighbor search method is the method of *near-neighbor lists* (NNLs) (see for example Ref. [86]). In this method, each particle has a list of its near neighbors, i.e., particles which are in close proximity (for example, within the cutoff for the interaction potential). As the particles move

around the lists need to be updated, and this is often done heuristically. Since the particles displace little from time step to time step in TDMD, the lists need to be updated only after many time steps (especially for dense systems). NNLs are not easy to use within EDMD because of the necessity to ensure correctness of the algorithm rigorously. If the order of events is not predicted correctly, the algorithm will typically fail with error conditions such as endless collision cycles between several particles. However, it is easily recognized that in order to efficiently treat nonspherical particles it is necessary to combine neighbor lists with the cell method. We now describe how this can be accomplished while maintaining a provably correct prediction of the collision sequence.

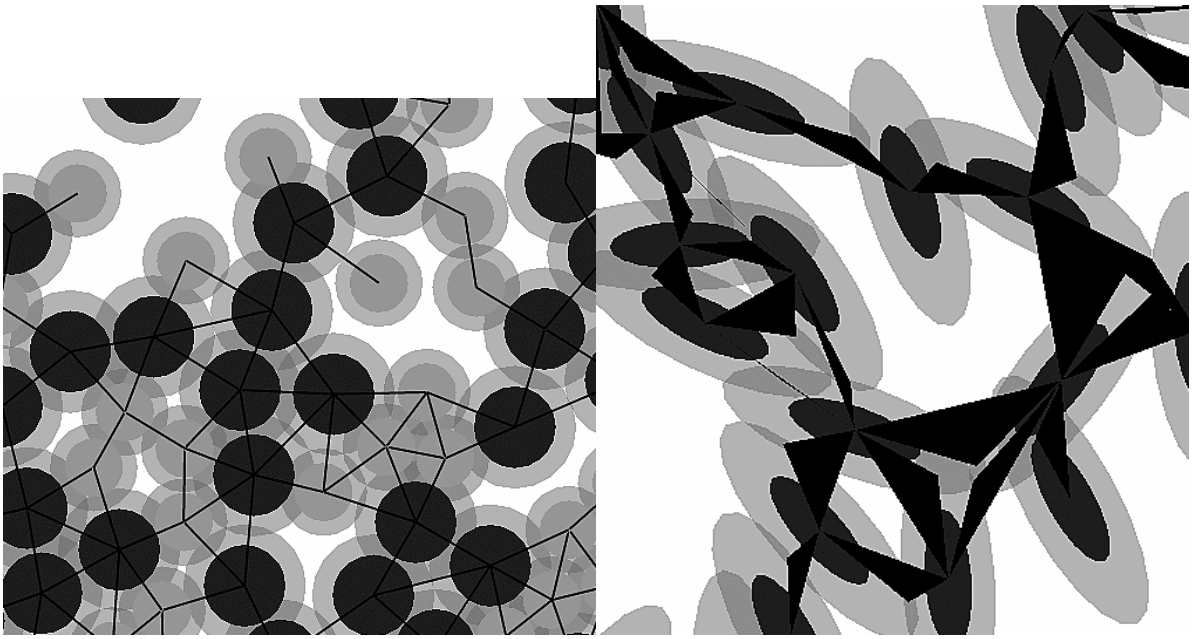

Figure 3.4: Illustration of NNLs for a system of disks (left) and ellipses (right). Particles are darker, and their bounding neighborhoods are lighter (it is easy to see which neighborhood goes with which particle). For disks the system is binary (bidisperse), and the neighborhoods are disks and the pairs of near-neighbors are shown as dark lines. For ellipses the neighborhoods are ellipses themselves and the interactions are shown as dark triangles whose vertices are given by the centroids of the two ellipses and the point of contact of the ellipses.

The main drawback of the cell method is that the shape of the cells is not adjusted to the shape of the particles, for example, elongated or squashed particles, but cubic cells. The main advantage, on the other hand, is that the partitioning into cells is static and independent of the motion of the particles. To correct for the drawback, we must compromise on the advantage: The partitioning into "cells" must be updated from time to time to reflect the motion of the particles, if we are to have any hope of having cells which take into account the shape of the particles. The idea is the following: Surround each particle i with a *bounding neighborhood* $\mathcal{N}(i)$, so that the particle is completely inside its bounding neighborhood, and the shape of the neighborhood is in some sense sensitive to the position and shape of the particle (for example, it should be elongated approximately along the same direction as the ellipsoid). Then, consider any two particles whose neighborhoods overlap to be near neighbors, and only calculate interaction potentials or check for collisions between such pairs. Each particle then stores a list of *interactions* in its near-neighbor list NNL(i), which is equivalent to each bounding neighborhood storing a list of neighborhoods with which it overlaps. This is illustrated for disks by using disks as the bounding neighborhoods in Fig. 3.4. Note that the cell method, as described earlier, *must* be used when (re-)building the NNLs, since overlap between neighborhoods cannot be checked efficiently otherwise. Building and maintaining the NNLs is expensive and dominates the computation for very aspherical particles. Finally, we note that the choice of the shape of the bounding neighborhoods and the exact way one constructs the NNLs is somewhat of a design choice. The

necessary invariant is that each particle be completely contained inside its bounding neighborhood and that there be an interaction in the NNLs for each pair of overlapping neighborhoods.

In this work we describe a specific conceptually simple approach which applies to hard particles of any shape and has worked well in practice. In our algorithm, the shape of $\mathcal{N}(i)$ is the *same* as the shape of particle i, but scaled uniformly with some scaling factor $\mu_{\text{neigh}} > 1$. Additionally, $\mathcal{N}(i)$ has the same centroid as i, at least at the instant in time when $\text{NNL}(i)$ is constructed (after which the particle may displace). This is illustrated for ellipses in Fig. 3.4. One wants to have the bounding neighborhood $\mathcal{N}(i)$ as large as possible so that there is more room for the particle i to move without the need to rebuild its NNL. However, the larger the neighborhood, the more neighbors there will be to examine. We investigate the optimal balance, as determined by the choice of μ_{neigh}, shortly. It is important to note that it would most likely be better to consider $\mathcal{N}(i)$ to be the set of all points that are within a given distance from the surface of particle i, especially for very nonspherical particles. This is because scaling a very elongated particle by a given factor μ produces unnecessarily long neighborhoods, which increases the cost of using the cell method to construct the neighbor lists. However, evaluating point-to-surface or surface-to-surface distances is quite nontrivial even for ellipsoids, and also the geometrical reasoning is obscured. On the other hand, using a bounding neighborhood which has the same shape as the particle is very intuitive and also efficient for ellipsoids.

3.3.2.1 The NNL Method in EDMD

Once the NNLs are built, one no longer needs to use the cell method, so long as all particles are still *completely* contained within their bounding neighborhoods. As time progresses, a particle may protrude outside its neighborhood, and in this case the NNLs need to be updated accordingly, using the fail-safe cell method. Details of this update will be given later. Therefore, when using NNLs, instead of transfers, another kind of event needs to be included: a "collision" with its bounding neighborhood. When using NNLs, transfers do not need to be handled at all. Namely, instead of using the cell method for the particles themselves, it should be used on the bounding neighborhoods. Each cell keeps a list of the bounding neighborhoods whose centroids it contains. Hard-walls are handled by including hard walls as neighbors in the NNLs of the particles whose bounding neighborhoods intersect a hard wall. At present we do not try to reuse any previous binary collisions when rebuilding neighbor lists because dealing with such reuse is rather complicated.

An additional complication when using NNLs arises when the boundary is deforming. Since in our approach all positional coordinates are expressed in relation to the (possibly deforming) lattice, the neighborhoods are not stationary but move together with the boundary. This may lead to originally disjoint neighborhoods overlapping later on. In order to ensure correctness of the neighbor search in such cases, one can add a "safety cushion" around each bounding neighborhood $\mathcal{N}(i)$. Specifically, two particles are to be considered neighbors if their bounding neighborhoods overlap when scaled by a common scaling factor $1+\epsilon_N$, where $\epsilon_N > 0$ is the relative size of the safety cushion. The NNLs need to be rebuilt completely whenever the boundary deformation becomes too large, because of the possibility of new neighborhood overlap. In this context, a measure of how much the boundary has deformed is given by the relative amount that Euclidean distances have changed due to the boundary deformation.

Consider a periodic system and two points with relative displacement $\mathbf{r}$, measured in lattice vectors. The Euclidean distance between them is $l^2 = \mathbf{r}^T\mathbf{G}\mathbf{r}$, where $\mathbf{G} = \mathbf{\Lambda}^T\mathbf{\Lambda}$ is a metric tensor. At a later time Δt, the distance changes, and the largest relative contraction in Euclidean distance between any two points is given by:

$$\min_{\mathbf{r}} \left[\frac{(l+\Delta l)}{l}\right]^2 = \min_{\mathbf{r}_E} \frac{\mathbf{r}_E^T \left(\mathbf{I}+\boldsymbol{\epsilon}\Delta t\right)^T \left(\mathbf{I}+\boldsymbol{\epsilon}\Delta t\right)\mathbf{r}_E}{\mathbf{r}_E^T\mathbf{r}_E^T} = \lambda_{\min}\left[\left(\mathbf{I}+\boldsymbol{\epsilon}\Delta t\right)^2\right],$$

where $\mathbf{r}_E = \mathbf{\Lambda}\mathbf{r}$ and $\lambda_{\min}$ denotes the minimal eigenvalue of a symmetric matrix. Therefore, the Euclidean distance between the centroids of two neighborhoods would not have contracted by more then a factor of $\lambda_{\min}\left[\left(\mathbf{I}+\boldsymbol{\epsilon}\Delta t\right)^2\right]$. In light of this observation, a reasonable heuristic approach is to periodically check the magnitude of the smallest eigenvalue of $\left(\mathbf{I}+\boldsymbol{\epsilon}\Delta t\right)^2$, and rebuild the NNLs completely whenever it deviates from unity by more then a few (as determined heuristically via experimentation) multiples of ϵ_N. Since it is reasonable to assume that the boundary deforms slowly compared to the particles, these kinds of updates will happen infrequently. This approach seems to work well in practice. In EDMD a rigorous approach is also possible, by predicting the first instance in time when two non-overlapping bounding neighborhoods first

overlap, and including this as a special event in the event queue. When this event is at the top of the queue, the simulation is essentially restarted from the current point in time. However, such an approach does not work in TDMD, and we have not found the practical need for such a complicated scheme either.

3.3.3 Very Aspherical Particles

Using the traditional cell method when rebuilding the NNLs is the computational bottleneck for very aspherical particles. To really obtain a fast yet rigorously correct event-driven algorithm for very aspherical particles the traditional cell method needs to be either abandoned or modified. It is clear that any neighbor search mechanism which only uses the centroids cannot be efficient. Although in a sense Ref. [121] studies the worst case of $\alpha \to \infty$ (needles, and similarly for platelets), it does not mention any additional techniques to handle the fact that as many as 50 needles can be in one cell in the reported simulations. This is probably because at that time only small systems ($N = 100 - 500$) could be studied, for which the cell method does not offer big savings even for spheres.

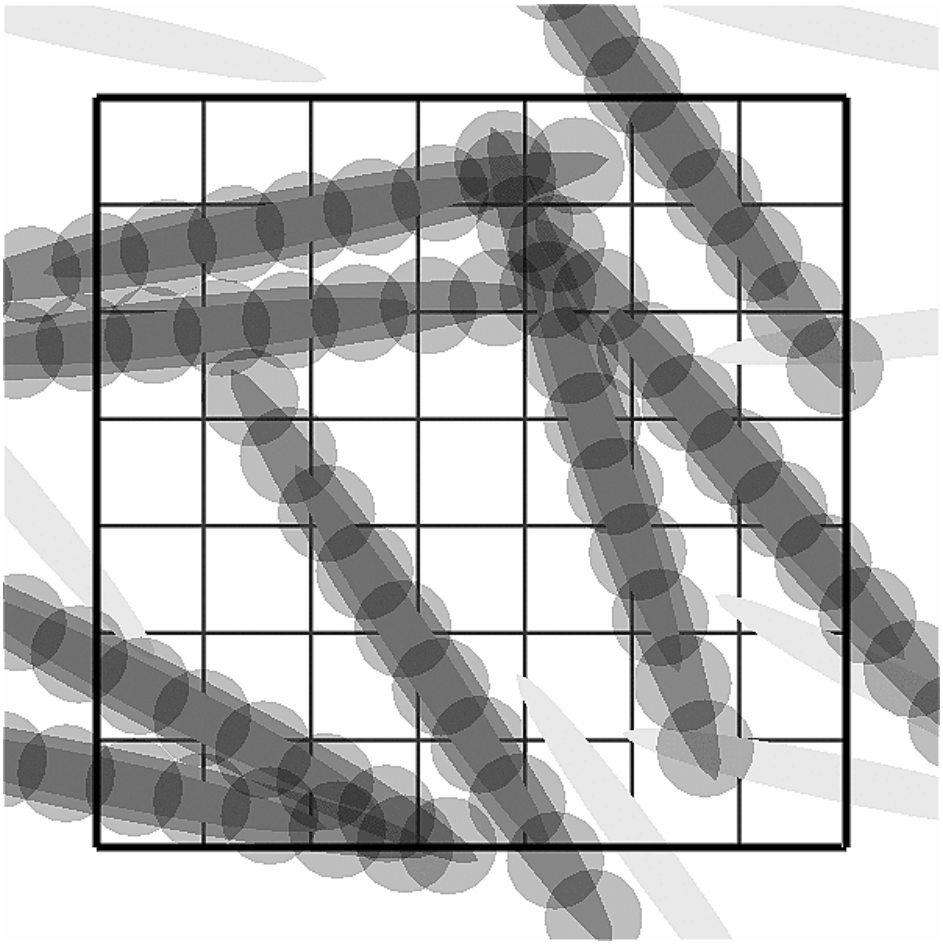

Figure 3.5: A small periodic packing of ellipses of aspect ratio $\alpha = 10$ illustrating the use of bounding sphere complexes. Each particle i (darkest shade) is bounded by its neighborhood (lighter shade) $\mathcal{N}(i)$, which is itself bounded by a collection of 10 disks BSC(i). A bounding neighborhood $\mathcal{N}(j)$ may overlap with $\mathcal{N}(i)$ if some of the bounding disks of particles j and i overlap. Therefore the usual cell grid (also shown) can be used in the search for neighbors to add to NNL(i). Image particles are shown in a lighter shade.

The approach we have implemented is to use several spheres to bound each particle, instead of just one large bounding sphere. We will refer to this collection of bounding spheres as the *bounding sphere complex*

(BSC). For the purposes of neighbor search, we still continue to use the cell method, however, we use the cell method on the collection of bounding spheres, not on the particles themselves. That is, we bin all of the bounding spheres in the cells, and the minimal Euclidean length of a cell is at least as large as the largest diameter of a bounding sphere. By increasing the number of bounding spheres per particle one can make the cells smaller. When searching for the neighbors of a given particle, one looks at all of its bounding spheres and their neighboring bounding spheres, and then checks whether the particles themselves are neighbors. This slightly complicates the search for neighbors, but the search can be optimized so that a given pair of particles is only checked once, rather then being checked for every pair of bounding spheres that they may share. It is hard to maintain the binning of the bounding spheres in cells as particles move. It is therefore essential to combine using BSCs with using NNLs. Each bounding neighborhood $\mathcal{N}(i)$ is bounded by BSC(i), that is, $\mathcal{N}(i)$ is completely contained in the union of the bounding spheres in BSC(i). The binning of the bounding spheres is only updated when NNL(i) is updated, and particle i is free to move inside $\mathcal{N}(i)$ without possibility of overlapping with a particle not in NNL(i). Using BSCs in two dimensions is illustrated in Fig. 3.5.

In our implementation, we use relative positions and radii for the spheres in BSC(i), expressed in a coordinate system in which particle i's orientation is aligned with the global coordinate system and the radius of its bounding sphere is unity. This enables us to not have to update the above quantities as the particle moves and changes shape, and also to share them between particles of identical shapes using pointers. When updating $\mathcal{N}(i)$, we can easily calculate the absolute (Euclidean) positions and radii of the bounding spheres from the relative ones.

In two dimensions, for very elongated objects, it is relatively easy to construct bounding complexes, however, this is not so easy in three dimensions, even though there are general methods (taken from computational geometry) for finding a good approximation to a particle shape with a few spheres [122]. We expect that there will be an optimal number of spheres N_S to use, this number increasing as the aspect ratio increases, however, it is not clear how to construct optimal BSCs. The approach we have implemented is to first bound each ellipse or ellipsoid in an orthogonal parallelepiped (rectangle in two dimensions), and then use a subset of a simple cubic lattice cover (a collection of identical spheres whose union covers all of Euclidean space) to bound (cover) the orthogonal parallelepiped. This kind of approach is far from optimal (for example, the lowest density sphere cover in three dimensions is given by a body-centered lattice of spheres), but it is very simple and works relatively well for sufficiently aspherical particles. This is illustrated in three dimensions for prolate and oblate ellipsoids in Fig. 3.6. As can be seen from the figure, it seems hard, if not impossible, to construct BSCs with few small spheres for flat (oblate) particles. Future research is needed to find a way to speed neighbor search for very oblate particles, and a promising direction to investigate is hierarchical bounding sphere complexes. Later we demonstrate that using BSCs in conjunction with NNLs significantly improve the speed of the EDMD algorithm for very elongated (prolate) particles. Note that using a large number of small bounding spheres (for very aspherical particles) requires a significant increase in the number of cells, and to save memory hashing may need to be used when manipulating the cell partitioning [123].

3.4 EDMD Algorithm

In this section, we describe our EDMD algorithm in significant detail, in the hope that this will prove very useful to other researchers implementing similar methods. Starting from a brief history of the main ideas used in the algorithm and a description of the basic notation, we proceed to give detailed descriptions of each step in the algorithm in the form of pseudo-codes. We first explain the top level event loop and its most involved step of predicting the impending event for a given particle. We then focus on binary collisions and boundary events separately, and finally describe algorithms for maintaining NNLs in a dynamic environment. Some of the steps of the algorithm, such as predicting the time of collision of two particles or processing a binary collision, depend on the particular particle shape in question and are illustrated specifically for ellipsoids in Section 3.5.

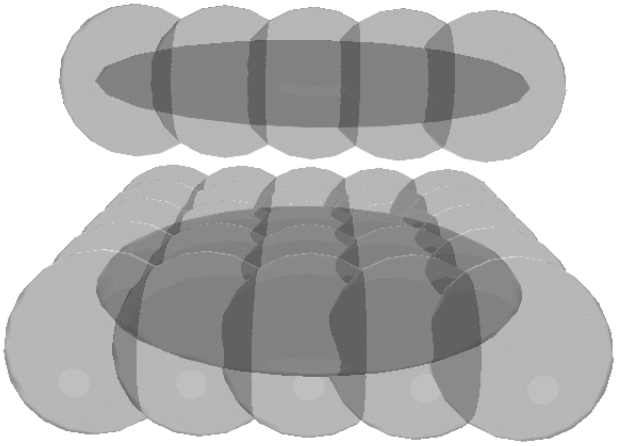

Figure 3.6: Bounding sphere complexes for spheroids of aspect ratio $\alpha = 5$. The prolate particle has 5 bounding spheres, but the oblate one has 25 bounding spheres.

3.4.1 History

We briefly summarize some of the previous work on EDMD algorithms. Although this has been done in other publications, we feel indebted to many authors whose ideas we have used and combined together to produce our algorithm, and would like to acknowledge them.

The very first MD simulation used an event-driven algorithm [102], and since those early attempts the core of an efficient EDMD algorithm for spherical particles, entailing a combination of delayed updates for the particles, the cell method and using a priority queue for the events, has been developed [124, 125]. Our approach borrows heavily from the EDMD algorithm developed by Lubachevsky [37]. We do not use a double-buffering technique as does Lubachevsky, following Ref. [126], and incorporate additional techniques developed by other authors.

One of the controversial questions in the history of EDMD is how many event predictions to retain for each particle i? As Ref. [126] demonstrates, it is best to use a heap (complete binary search tree) for the priority event queue, and we follow this approach. It seems clear that only the impending prediction for each particle should be put in the event queue (i.e., the size of the heap is equal to N), but this prediction may be invalidated later (due to a third-party event, for example). In such cases, it may be possible to reuse some of the other previously-predicted binary collisions for i, for example, the one scheduled with the second-smallest time [127, 128]. This requires additional memory for storing more predictions per particle and adds complexity to the algorithm. We have adapted the conclusion of Ref. [104] that this complexity is not justified from an efficiency standpoint. Ref. [128] makes the important observation that after a transfer fewer cells need to be checked for collisions. The authors of Ref. [104] thus predict and store separately the next binary collision and the next transfer for each particle, and only insert the one with the smaller time into the event heap. More exotic EDMD algorithms, for example, aimed at increased simplicity or ease of vectorization [120], have been developed. We build on these previous developments and combine neighbor-list techniques traditionally used in TDMD to develop a novel EDMD algorithm specifically tailored to systems of nonspherical particles at relatively high densities.

3.4.2 Notation

As explained above, the EDMD algorithm consists of processing a sequence of time-ordered events. Each particle must store some basic information needed to predict and process the events. An *event* (t_e, p_e) is specified by giving the predicted time of occurrence t_e and the partner p_e. A special type of an event is a binary collision $[t_c, (p_c, v_c)]$, determined by specifying the time of collision $t_c \equiv t_e$ and the partner in the event $(p_c \equiv p_e, v_c)$. The primary use of the image (virtual) identifier v_c is to distinguish between images of a

given particle when periodic boundary conditions are used. Note that the *collision schedules* must be kept symmetric at all times, that is, if particle i has an impending event with (j, v), then particle j must have an impending event with $(i, -v)$. Although the cell a particle belongs to can be determined from the position of its centroid, this is difficult to do exactly when a particle is near the boundary of a cell due to roundoff errors (possible tricks to avoid such problems include adding a cushion around each cell and not considering a transfer until the particle is sufficiently outside the cell [126]). We have chosen to explicitly store and maintain the cell that a particle, a bounding neighborhood of a particle, or a bounding sphere, belongs to (as determined by the corresponding centroid).

In summary, for each particle $i = 1, \ldots, N$, we store:

1. The predicted *impending event* (t_e, p_e, v_c) along with any other information which can help process the event or collision more efficiently should it actually happen later.
2. The last update *time* t.
3. The *state* of the particle at time t, including:
   (a) Its *configuration*, including the relative *position* of the centroid $\mathbf{r}$ and any additional configuration (such as orientation) $\mathbf{q}$, as well as the particle *shape* (such as radius, semiaxes, etc.) $\mathbf{O}$. Note that $\mathbf{O}$ may be shared among many particles using pointers (for example, all particles have the same shape at all times in a monodisperse packing) and thus not be updated to time t but still be at time zero[4].
   (b) The particle *motion*, including the relative velocity of the centroid $\mathbf{v} = \dot{\mathbf{r}}$ and additional (such as angular) velocity $\boldsymbol{\omega}$ representing $\dot{\mathbf{q}}$. Also included in the motion is the rate of deformation of the particle shape $\boldsymbol{\Gamma}$ (possibly shared among different particles).
   (c) The particle cell c, to which $\mathbf{r}$ belongs, if not using NNLs.
4. Dynamical parameters, such as particle mass or moment of inertia (possible shared with other particles).
5. If using NNLs, the configuration of the (immobile) bounding neighborhood $\mathcal{N}(i)$, $\mathbf{r}_N$ and $\mathbf{q}_N$, its shape $\mathbf{O}_N$, as well as the cell c to which $\mathbf{r}_N$ belongs.
6. If using BSCs in addition to NNLs:
   (a) The relative positions $\mathbf{r}_j^{BS}$ and relative radii O_j^{BS}, $j = 1, \ldots, N_{BS}$, of its N_{BS} bounding spheres, along with the largest BS radius $O_{\max}^{BS} = \max_j O_j^{BS}$. These are expressed *relative* to the position and size of $\mathcal{N}(i)$.
   (b) The cell c_j^{BS} that $\mathbf{r}_j^{BS}$ belongs to, $j = 1, \ldots, N_{BS}$.

For each of these quantities, we will usually explicitly indicate the particle to which they pertain, for example, $t(i)$ will denote the time of particle i.

3.4.2.1 Event Identifiers

Each particle must predict its impending event, and there are several different basic types of events: binary collisions (the primary type of event), wall collisions (i.e., collisions with a boundary of the simulation domain), collisions with a bounding neighborhood (i.e., a particle leaving the interior of its bounding neighborhood), transfers (between cells), and checks (re-predicting the impending event). Additionally, several different types of checks can be distinguished, depending on why a check was required and whether the motion of the particle changed (in which case old predictions are invalid) or not (in which old predictions may be reused). We consider transfers and wall collisions together as *boundary events* (or boundary "collisions"), since their prediction and processing is very similar (especially for periodic BCs). The exact cell wall through which the particle exits the (unit) cell, or the wall with which the particle collides, is identified with an integer w, which is negative if the event is with a wall of the unit cell (boundary).

In our implementation, the type of a predicted event for a particle i is distinguished based on the *event partner* p (possibly including an image identifier v):

[4]We have implemented a different approach for systems with a few types of particles (monodisperse, bidisperse, etc.), for which we store the particle shape information separately from the particles and share it among them, and polydisperse systems in which each particle has a (potentially) different shape, for which we store the particle shape together with the rest of the particle state.

$0 \leq p \leq N$ A binary collision between particles i and (p, v), where v is the virtual identifier of the partner.

$p = -\infty$ A check (update) after an event occurred that did not alter the motion of i.

$p > 2N$ Transfer between cells, i.e., "collision" with wall $w = p - 2N$, $w > 0$.

$p < -2N$ Wall collision with wall $w = p + 2N$, $w < 0$, which can be a real hard wall or the boundary of the unit cell.

$N < p \leq 2N$ Check after binary collision with partner $(j, -v)$, where $j = p - N$ (the motion of particle i has changed).

$p = 0$ Check for particle i after an event occurred which altered the motion of i.

$p = \infty$ Collision with the bounding neighborhood $\mathcal{N}(i)$.

The range $-2N \leq p < 0$ is reserved for future (parallel implementation) uses. Of course, one can also store the partner as two integers, one indicating the type of event and the other identifying the partner, however, the above approach saves space.

3.4.3 Processing the Current Event

Algorithm 2 represents the main event loop in the EDMD algorithm, which processes events one after the other in the order they occur and advances the global time t accordingly. It uses a collection of other auxiliary steps, the algorithms of which are given in what follows. Note that when processing the collision of particle i with particle (j, v), we also update particle j, and later, when processing the same collision but as a collision of j and $(i, -v)$, we skip the update. Also, note that when using NNLs, there are two options: Completely rebuild the NNLs as soon as some particle i collides with its neighborhood, or, rebuild only the neighbor list NNL(i). We discuss the advantages and disadvantages of each approach and compare their practical performance in latter sections.

Algorithm 2: Process the next event in the event heap.

---

1. Delete (pop) the top of the event queue (heap) to find the next particle i to have an event with $p_e(i)$ at $t_e(i)$.
2. Perform global checks to ensure the validity of the event prediction. For example:
   (a) If the boundary is deforming, and if at time $t_e(i)$ the cell length L_c is not larger then the largest enclosing sphere diameter $D_{\max}$, $L_c\left[t_e(i)\right] \leq D_{\max}\left[t_e(i)\right]$, then restart the simulation:
      i. Synchronize all particles (Algorithm 3).
      ii. Repartition the simulation box to increase the length L_c (for example, for lattice boundaries, increase the appropriate $N_k^{(c)}$).
      iii. Re-bin the particles into the new cells based on the positions of their centroids.
      iv. Reset the event schedule (Algorithm 4).
      v. Go back to step 1.
   (b) If using NNLs and the NNLs are no longer valid (for example, due to boundary deformation), then:
      i. Synchronize all particles.
      ii. Rebuild the NNLs (Algorithm 9).
      iii. Reset the event schedule.
      iv. Go back to step 1.
3. If the boundary is deforming, update its shape. For example, for lattice-based boundaries, set $\mathbf{\Lambda} \leftarrow \mathbf{\Lambda} + \dot{\mathbf{\Lambda}}\left[t_e(i) - t\right]$.
4. Advance the global simulation time $t \leftarrow t_e(i)$.
5. If the event to process is not a check after a binary collision, then update the configuration of particle i to time t (for example, $\mathbf{r}(i) \leftarrow \mathbf{r}(i) + \left[t - t(i)\right]\mathbf{v}_i$), and set $t(i) \leftarrow t$.
6. If using NNLs and event is a collision with a bounding neighborhood, then:
   (a) If completely rebuilding NNLs, then declare NNLs invalid and execute step 2b.

(b) Else, record a snapshot of the current shape of particle i (recall that this may be shared with other particles) in $\mathbf{O}_i$ and rebuild the NNL of particle i (Algorithm 10).

7. If the event is a wall collision or cell transfer, then:
   (a) If $p_e(i) > 0$ then set $w \leftarrow p_e(i) - 2N$ (transfer).
   (b) Else set $w \leftarrow p_e(i) + 2N$ (wall collision).
   (c) Process the boundary event with "wall" w (Algorithm 8).
8. If the event is a binary collision, then:
   (a) Update the configuration of particle $j = p_e(i)$ to time t and set $t(j) \leftarrow t$ and $p_e(j) \leftarrow N + i$ (mark j's event as a check).
   (b) Process the binary collision between i and j (see specific algorithm for ellipsoids 11).
9. Predict the next collision and event for particle i (Algorithm 5).
10. Insert particle i back into the event heap with key $t_e(i)$.
11. Terminate the simulation or go back to step 1.

Because EDMD is asynchronous, it is often necessary to bring all the particles to the same point in time (synchronize) and obtain a snapshot of the system at the current time t. This is done with Algorithm 3. Note that we reset the time to $t = 0$ after such a synchronization step. Another step which appears frequently is to reset all the future event predictions and start afresh, typically after a synchronization. In particular this needs to be done when initializing the algorithm. The steps to do this are outlined in Algorithm 4.

1. If $t = 0$ then return.
2. For all particles $i = 1, \ldots, N$ do:
   (a) Update the configuration of particle i to time t.
   (b) Set $t_e(i) \leftarrow t_e(i) - t$, $t_c(i) \leftarrow t_c(i) - t$ and $t(i) \leftarrow 0$.
3. Update the shapes of all particles to time t.
4. Store the total elapsed time $T \leftarrow T + t$ and set $t \leftarrow 0$.

Algorithm 3: Synchronize all particles to the current simulation time t.

1. Reset the event heap to empty
2. For all particles $i = 1, \ldots, N$ do:
   (a) Set $p_e(i), p_c(i) \leftarrow 0$ and $t_e(i), t_c(i) \leftarrow 0$.
   (b) Insert particle i into the event heap with key $t_e(i)$.

Algorithm 4: Reset the schedule of events.

3.4.4 Predicting The Next Event

The most important and most involved step in the event loop is predicting the next event to happen to a given particle, possibly right after another event has been processed. Algorithm 5 outlines this process. Note that it is likely possible to further extend and improve this particular step by better separating motion-altering from motion-preserving events and improving the reuse of previous event predictions.

Algorithm 5: Predict the next binary collision and event for particle i, after an event involving i happened.

1. If not using NNLs, then:

(a) Initialize $t_w \leftarrow \infty$ and $\tilde{t_w} \leftarrow \infty$ and set $w \leftarrow 0$.

(b) Predict the next boundary event (wall collision or transfer) time t_w and partner "wall" w for particle i, if any, by looking at all of the boundaries of $c(i)$ (Algorithm 7). If an exact prediction could not be made (for example, if a hard wall was involved and the search was terminated prematurely), calculate a time $\tilde{t_w}$ up to which a boundary event is guaranteed not to happen and set $w \leftarrow 0$.

(c) If $w = 0$, then force a check at time $\tilde{t_w}$, $p_e(i) \leftarrow -\infty$ and $t_e(i) \leftarrow \tilde{t_w}$,

(d) else predict $t_e(i) \leftarrow t_w$ and:

i. If $w < 0$ then set $p_e(i) \leftarrow w - 2N$,

ii. else set $p_e(i) \leftarrow w + 2N$.

(e) If a hard-wall prediction was made, store any necessary information needed to process the collision more efficiently later (for example, store λ in the case of ellipsoids).

(f) For all particles (j, v) in the cells in the first neighborhood of $c(i)$, execute step 4,

2. else if using NNLs, then:

(a) Predict the time t_N particle i will protrude outside of (collide with) its bounding neighborhood $\mathcal{N}(i)$, limiting the length of the search interval to $t_e(i)$. If an exact prediction is not possible, calculate a time $\tilde{t_N}$ before which i is completely contained in $\mathcal{N}(i)$.

(b) If t_N was calculated and $t_N < t_e(i)$, then record:

i. Set $p_e(i) \leftarrow \infty$ and $t_e(i) \leftarrow t_N$.

ii. Potentially store any additional information about this collision for particle i,

(c) else if $\tilde{t_N}$ was calculated and $\tilde{t_N} < t_e(i)$ then force a new prediction for particle i at time $\tilde{t_N}$, $p_e(i) \leftarrow -\infty$ and $t_e(i) \leftarrow \tilde{t_N}$.

(d) For all hard walls w in NNL(i), predict the time of collision t_w. If an exact prediction could not be made, calculate a time $\tilde{t_w}$ up to which the collision is guaranteed not to happen.

(e) If t_w was calculated and $t_w < t_e(i)$, then record:

i. Set $t_e(i) \leftarrow t_w$ and $p_e(i) \leftarrow w - 2N$.

ii. Potentially store any necessary information needed to process the wall collision more efficiently later,

(f) else if $\tilde{t_w}$ was calculated and $\tilde{t_w} < t_e(i)$, then force a check $p_e(i) \leftarrow -\infty$ and $t_e(i) \leftarrow \tilde{t_w}$.

(g) For all particles (j, v) in NNL(i), execute step 4,

3. Skip step 4.

4. Predict the time of collision between particles i and (j, v):

(a) Predict if i and (j, v) will collide during a time interval of length $\min[t_e(i), t_e(j)]$ and if yes, calculate the time of collision t_c, or calculate a time $\tilde{t_c} < t_c$ before which a collision will not happen (see specific algorithm for ellipsoids 11).

(b) If t_c was calculated and $t_c < \min[t_e(i), t_e(j)]$, then record this collision as the next predicted binary collision for particle i:

i. Set $p_e(i) \leftarrow j$, $v(i) \leftarrow v$ and $t_e(i) \leftarrow t_c$.

ii. Potentially store any additional information about this collision for particle i (for example, λ in the case of ellipsoids),

(c) else if $\tilde{t_c}$ was calculated and $\tilde{t_c} < t_e(i)$ then force a new prediction for particle i at time $\tilde{t_c}$, $p_e(i) \leftarrow -\infty$ and $t_e(i) \leftarrow \tilde{t_c}$.

5. If $0 < p_e(i) \leq N$ then let $j = p_e(i)$ (a new collision partner was found), and:

(a) If the involved third-party $m = p_e(j)$ is a real particle, $0 < m \leq N$ and $m \neq i$, then invalidate the third party collision prediction, $p_e(m) \leftarrow -\infty$.

(b) Ensure that the collision predictions are symmetric by setting $p_e(j) \leftarrow i$, $v(j) = -v(i)$ and $t_e(j) \leftarrow t_e(i)$. Also copy any additional information about the predicted collision to particle j as well (in the case of ellipsoids, this involves storing $(1 - \lambda)$ for particle j).

(c) Update the key of j in the event heap to $t_e(j)$.

3.4.5 Binary Collisions

The two main steps in dealing with binary collisions is predicting them and processing them. Processing a collision is inherently tied to the shape of the particle. We give a generic specification of how to predict binary collisions between particles in Algorithm 6, and a specific implementation for ellipsoids is given later.

1. Convert v into a virtual displacement of particle j in terms of unit cells, $\Delta\mathbf{r}_j \equiv \mathbf{n}_c$, as discussed in Section 3.2.3.
2. Calculate the current configuration of the particles i and j, for example, the positions of their centroids,

$$\begin{aligned} \mathbf{r}_i &\leftarrow \mathbf{r}(i) + [t - t(i)]\,\mathbf{v}_i \\ \mathbf{r}_j &\leftarrow \mathbf{r}(j) + [t - t(j)]\,\mathbf{v}_j + \Delta\mathbf{r}_j, \end{aligned}$$

   and their current orientations for nonspherical particles.
3. If the shape of the particles is changing, calculate the current shape of i and j.
4. Eliminate any further use of relative positions in the procedure by calculating the current Euclidean positions, velocities and accelerations of the particles using Eqs. (3.4-3.6) and the above $\mathbf{r}_i$ and $\mathbf{r}_j$.
5. Calculate the collision time t_c or $\tilde{t_c}$ of two moving and possibly deforming particles of the given initial shapes and configurations and initial Euclidean positions, velocities and accelerations, assuming a force-free motion starting at time zero. Optionally collect additional information needed to process the collision faster if it actually happens. See the specific algorithm for ellipsoids 11.
6. Correct the prediction to account for the current time, $t_c \leftarrow t_c + t$ or $\tilde{t_c} \leftarrow \tilde{t_c} + t$.

Algorithm 6: Predict the (first) time of collision between particles i and (j, v), t_c. If prediction cannot be verified, return a time $\tilde{t_c}$ before which a collision will not happen. Possibly also return additional information about the collision.

3.4.6 Boundary Events

In this section we focus on lattice-based boundaries and give a prescription for predicting and processing boundary events (transfers and wall collisions).

3.4.6.1 Prediction

When NNLs are not used, one must check all the boundaries of the current particle cell $c(i)$ and find the first time the particle leaves the cell or collides with a hard wall, if any. We do not give details for predicting or processing hard-wall collisions. For lattice based boundaries, the prediction of the next boundary event proceeds independently along each dimension, and then the smallest of the d event times is selected, as illustrated in Algorithm 7.

3.4.6.2 Processing

Processing the boundary events amounts to little work when the event is a transfer from one cell to another. For periodic BCs however, additional work occurs when the particle crosses the boundary of the unit cell (i.e., the simulation domain), since in this case it must be translated by a lattice vector in order to return it back into the unit cell. Considerably more complicated is the processing of collisions with hard walls, especially for nonspherical particles or when the lattice velocity is nonzero, however, we do not give the details of these steps in Algorithm 8.

3.4.7 Building and Updating the NNLs

In our implementation, all of the NNLs are implemented as an optimized form of linked lists. Each interaction $[(j, v), p]$ in NNL(i) stores the *partner* (j, v) and a *priority* p. We usually prescribe a fixed upper bound on the number of neighbors (interactions) N_i that a particle can have (this allows us to preallocate all storage and guarantee that additional memory will not be used unless really necessary), which can vary between particles if necessary. Only the N_i interactions with highest priority are retained in NNL(i). This kind of NNL can be used for a variety of tasks, including finding the first few nearest neighbors of any particle. We allow the NNLs to asymmetric, i.e., just because particle i interacts with particle (j, k), it is not implied that particle j interacts with $(i, -k)$, but rather, the reverse interaction must be stored in NNL(j) if needed. In

1. Convert the cell identifier $1 \leq c(i) \leq N_c$ into a d-dimensional vector giving the positions of the cell in the Cartesian grid of cells, $1 \leq \mathbf{g}^{(c)} \leq \mathbf{N}^{(c)}$.
2. For all dimensions, $k = 1, \ldots, d$, do:
   (a) Predict the time when the particle centroid will cross a wall of $c(i)$ along dimension k:
      i. If $v_k(i) = [\mathbf{v}(i)]_k > 0$ (particle will exit on the "right" side of the bin), then
         A. Set $w_k \leftarrow 2(k-1)+2$ and
         $$t_k \leftarrow \left[\mathbf{g}_k^{(c)} - N_k^{(c)} r_k(i)\right] / \left[N_k^{(c)} v_k(i)\right].$$
         B. If boundary is periodic along dimension k and $\mathbf{g}_k^{(c)} = N_k^{(c)}$, then set $w_k \leftarrow -w_k$,
      ii. else if $v_k(i) < 0$ (particle will exit on the "left" side of the bin), then
         A. Set $w_k \leftarrow 2(k-1)+1$, and
         $$t_k = \left[N_k^{(c)} r_k(i) - \mathbf{g}_k^{(c)} + 1\right] / \left[N_k^{(c)} v_k(i)\right].$$
         B. If boundary is periodic along dimension k and $\mathbf{g}_k^{(c)} = 1$, then set $w_k \leftarrow -w_k$,
      iii. else set $t_k \leftarrow \infty$ and $w_k \leftarrow 0$.
   (b) If boundary is not periodic along dimension k, then also predict the time of collision with the hard wall boundaries along dimension k, assuming that the particle starts from zero time:
      i. If $\mathbf{g}_k^{(c)} = N_k^{(c)}$, predict time of collision with the "right" hard wall along dimension k, $t_k^{(hw)}$. If $t_k^{(hw)} < t_k$, then set $t_k \leftarrow t_k^{(hw)}$ and $w_k \leftarrow -[2(k-1)+2]$.
      ii. If $\mathbf{g}_k^{(c)} = 1$, predict time of collision with the "left" hard wall along dimension k, $t_k^{(hw)}$. If $t_k^{(hw)} < t_k$, then set $t_k \leftarrow t_k^{(hw)}$ and $w_k \leftarrow -[2(k-1)+1]$.
3. Find the dimension $\tilde{k}$ with the smallest t_k and return $t_w = t(i) + t_{\tilde{k}}$ and $w = w_{\tilde{k}}$.

Algorithm 7: Predict the next *wall event* with "partner" w for particle i moving with relative velocity $\mathbf{v}(i)$ and the time of occurrence t_w, for a *lattice-based boundary*. The sign of w determines the type of event: $w > 0$ specifies that particle i leaves its cell $c(i)$ through one of the cell boundaries, while $w < 0$ specifies that the particle collides with one of the hard walls or crosses one of the boundaries of a unit cell *and* leaves its bin, for a periodic system. The value of $|w|$ determines the exact cell boundary or wall.

the particular use of NNLs for neighbor search, the priorities are the negative of the "distances" between the particles, so that only the closest N_i particles are retained as neighbors.

There are two main ways of updating the NNLs after a particle collides with its bounding neighborhood. One is to completely update the NNLs of all particles and start afresh, and the other one is to only update the NNL of the particle in question. We next discuss these two forms of NNL updates, *complete* and *partial*, and compare them practically in Section 3.6 to conclude that it is in general preferable to use partial updates (however, there are situations when it is best to use complete updates). As explained earlier, we focus on the case when the bounding neighborhoods are scaled versions of the particles. In addition to limiting the number of near-neighbors of any particle to N_i, we limit the maximum scaling of the neighborhood with respect to the particle itself to $\mu_{\text{cutoff}} \geq \mu_{\text{neigh}} > 1$, and count as overlapping any neighborhoods which overlap when scaled by an additional factor $(1+\epsilon_\mu)$, where $\epsilon_\mu \geq 0$ is a safety cushion used when the boundary deforms. Henceforth, denote $\mu_{\max} = (1+\epsilon_\mu)\,\mu_{\text{cutoff}}$.

3.4.7.1 Complete Updates

A simpler form of update is after a complete resetting of the NNLs, i.e., building the NNLs from scratch. Algorithm 9 gives a recipe for this. The aim of the algorithm is to try to make the bounding neighborhoods have a scale factor of $\mu_{\max}$ and add all overlapping neighborhoods in the NNLs. This will always be possible

1. From w, find the dimension k along which the event happens and the side ("left" or "right").
2. If this is a boundary event, $w < 0$, then:
   (a) If the boundary is periodic along k, then:
      i. Shift the particle by a unit cell, $r_k(i) \leftarrow r_k(i) + 1$ if particle is exiting its cell to the right, or $r_k(i) \leftarrow r_k(i) - 1$ if exiting to the left.
      ii. Let $j = p_c(i)$. If $0 < j \leq N$, correct the virtual identifiers for the predicted collision between i and j, $v(i)$ and $v(j)$, to account for the shift in step 2(a)i.
      iii. Pretend that this is a simple transfer, $w \leftarrow 2N - w$,
   (b) else, process the collision of the particle with the hard-wall. This will typically involve calculating the Euclidean position and velocity of the particle, calculating the exchange of momentum between the particle and the wall, calculating the new Euclidean velocity of the particle $\mathbf{v}^{(E)}$ (and also $\boldsymbol{\omega}$ if necessary), converting back to relative velocity, and updating the velocity $\mathbf{v}$ (and $\boldsymbol{\omega}$).
3. If this is a transfer (note step 2(a)iii above), $w > 0$, then update the cell of the particle $c(i)$ and move the particle from the linked list of its previous cell to the list of the new cell.

Algorithm 8: Process the boundary event (transfer or collision with a hard-wall) of particle i with wall w, assuming a lattice-based boundary.

if N_i is large enough. However, we allow one to limit the number of near neighbors. This is useful when there is not a good estimate of what a good μ_{cutoff} is.

The algorithm is significantly more complicated when BSCs are used since the search for possibly overlapping bounding neighborhoods needs to be done over pairs of bounding spheres. To avoid checking a given pair of bounding neighborhoods for overlap multiple times, we use an integer mask $M(i)$ for each particle, which we assume is persistent, i.e., stored for each particle between updates. In our algorithm, a hard wall can be a neighbor in $\text{NNL}(i)$ if $\mathcal{N}(i)$ is intersected by a hard-wall boundary. For simplicity, we do not present pseudocode for adding these hard-wall neighbors, however, it is a straightforward exercise to add these steps to the algorithms below.

Algorithm 9: *Completely* update the near-neighbor lists (NNLs) by rebuilding them from scratch. Assume all particles have been synchronized to the same point in time.

1. For all particles, $i = 1, \ldots, N$, reset $\text{NNL}(i)$ to an empty list.
2. For all particles, $i = 1, \ldots, N$, reduce $\mathbf{r}(i)$ to the first unit cell, and if $\mathbf{r}(i)$ is no longer inside $c(i)$, then remove i from the linked list of $c(i)$, update $c(i)$, and insert i in the list of the new $c(i)$.
3. If using BSCs, then initialize the largest (absolute) radius of a bounding sphere $O_{\max}^{(E)} \leftarrow 0$, and for all particles, $i = 1, \ldots, N$, do:
   (a) Set the bounding neighborhood of i to have the same centroid, orientation and shape as i but be scaled by a factor $\mu_{\max}$, $\mathbf{r}_N(i) \leftarrow \mathbf{r}(i)$, $\mathbf{q}_N(i) \leftarrow \mathbf{q}(i)$ and $\mathbf{O}_N(i) \leftarrow \mu_{\max}\mathbf{O}(i)$.
   (b) For all bounding spheres of i, $k = 1, \ldots, N_{BS}(i)$, do:
      i. Remove the sphere from the linked list of cell $c_k^{BS}(i)$.
      ii. Calculate the new absolute position of its center and the cell it is in, update $c_k^{BS}(i)$ accordingly, and insert the sphere into the linked list of $c_k^{BS}(i)$.
      iii. Calculate the absolute radius $O_k^{(E)}$ of the bounding sphere and set $O_{\max}^{(E)} \leftarrow \max\left\{O_{\max}^{(E)}, O_k^{(E)}\right\}$.
   (c) Initialize the mask $M(i) \leftarrow 0$.
4. else let $O_{\max}^{(E)} \leftarrow \mu_{\max}\left\{\max_i\left[O_{\max}(i)\right]\right\}$ be the largest possible radius of an enclosing sphere of a bounding neighborhood.
5. For all particles, $i = 1, \ldots, N$, do:
   (a) If using BSCs, then for all bounding spheres of i, $k = 1, \ldots, N_{BS}(i)$, do:
      i. For all cells c_i in the neighborhood of $c_k^{BS}(i)$ of Euclidean extent $2O_{\max}^{(E)}$, and for all bounding spheres in c_i belonging to some particle (j, v), do:

A. If $j \geq i$ and $M(j) \neq \text{sign}(v)\,(|v|\,N + i)$, then execute step 7,
B. else mark this pair of particles as already checked, $M(j) \leftarrow \text{sign}(v)\,(|v|\,N + i)$.

(b) else, for all cells c_i in the neighborhood of $c(i)$ of Euclidean extent $2O_{\max}^{(E)}$ (note that his may involve higher-order neighbors of $c(i)$), do:

i. For all particles $(j, v) \in c_i$ such that $j \geq i$, execute step 7.

6. Skip step 7.
7. If the largest common scaling factor which leaves i and (j, v) disjoint, $\mu_{ij} \leq \mu_{\max}$, then:
   (a) Calculate μ_{ij} exactly.
      i. Insert the interaction $[(j, v), -\mu_{ij}]$ in NNL(i). Note this may remove some previous entries in NNL(i) if it is already full.
      ii. Insert the interaction $[(i, -v), -\mu_{ij}]$ in NNL(j).
8. For all particles, $i = 1, \ldots, N$, do:
   (a) Initialize the minimal scaling of i which makes it overlap with the bounding neighborhood of a non-neighbor particle, $\mu_{\min}^{\text{non-neigh}} \leftarrow \mu_{\max}$.
   (b) If NNL(i) is not full, then initialize the maximal scaling of i which leaves it disjoint from at least one of the bounding neighborhoods of a neighbor particle, $\mu_{\max}^{\text{neigh}} \leftarrow \mu_{\max}$, otherwise initialize $\mu_{\max}^{\text{neigh}} \leftarrow 0$.
   (c) For all interactions $[(j, v), p]$ in NNL(i), ensure that they are bi-directional:
      i. If $-p > \mu_{\max}^{\text{neigh}}$, then set $\mu_{\max}^{\text{neigh}} \leftarrow -p$.
      ii. If there is no interaction with particle $(i, -v)$ in NNL(j), then:
         A. If $-p > \mu_{\min}^{\text{non-neigh}}$, then set $\mu_{\min}^{\text{non-neigh}} \leftarrow -p$.
         B. Delete the interaction $[(j, v), p]$ from NNL(i).
      iii. If $-p < \mu_{\text{neigh}}$, then set $\mu_{\text{neigh}} \leftarrow -p$.
   (d) Set $\mu_{\text{neigh}} \leftarrow \min\left(\mu_{\min}^{\text{non-neigh}}, \mu_{\max}^{\text{neigh}}\right)$. Note that if NNL($i$) never filled up then $\mu_{\text{neigh}} = \mu_{\max}$.
   (e) If $\mu_{\text{neigh}} < \mu_{\max}$, then set $\mathbf{O}_N(i) \leftarrow \mu_{\max}\mathbf{O}(i)$.
   (f) If using BSCs and $\mu_{\text{neigh}} < \mu_{\max}$, then for all bounding spheres of i, $k = 1, \ldots, N_{BS}(i)$, do:
      i. Remove the sphere from the linked list of cell $c_k^{BS}(i)$.
      ii. Calculate the new absolute position of its center and the cell it is in, update $c_k^{BS}(i)$ accordingly, and insert the sphere into the linked list of $c_k^{BS}(i)$.
9. If the boundary is deforming, record the current shape of the boundary to be used later to verify the validity of the NNLs (see Section 3.3.2.1).

3.4.7.2 Partial Updates

A considerably more complex task is updating NNL(i) while trying to leave the lists of other particles intact, other than possibly adding or deleting an interaction involving i. We give a prescription for this in Algorithm 10, but do not give many details, as understanding each step is not necessary to get an idea of the overall approach. For simplicity, we do not present the case when BSCs are used, as the modifications to allow for bounding complexes closely parallel those in Algorithm 9 and it is a straightforward exercise for the reader to modify the algorithm below accordingly.

Algorithm 10: Update the near-neighbor list of particle i, NNL(i). Assume that the current shape of i is passed in $\mathbf{O}_i$.

---

1. For all interactions with (j, v) in NNL(i), delete the reverse interaction with $(i, -v)$ in NNL(j).
2. Initialize the minimal scaling of i which makes it overlap with the bounding neighborhood of a non-neighbor particle, $\mu_{\min}^{\text{non-neigh}} \leftarrow \mu_{\max}$, as well as the maximal scaling of i which leaves it disjoint from at least one of the bounding neighborhoods of a neighbor particle, $\mu_{\max}^{\text{neigh}} \leftarrow \mu_{\max}$.
3. For all cells c_i in the neighborhood of $c(i)$ of Euclidean extent $O_{\max}^{\text{neigh}} + \mu_{\max}O_i$, where O_i is the radius of the bounding sphere of i and $O_{\max}^{\text{neigh}}$ is the radius of the largest enclosing sphere of a particle neighborhood, do:
   (a) For all bounding neighborhoods $\mathcal{N}_j$ in the list of c_i, $c(j) = c_i$, do:
      i. If the largest scaling factor which leaves i disjoint from the neighborhood of (j, v), $\mu_{ij}^{\text{neigh}} < \mu_{\max}$, then calculate μ_{ij}^{neigh} exactly, else continue with next particle (j, v).

ii. If there is room in NNL(j), then insert the interaction $\left[(j,v), -\mu_{ij}^{\text{neigh}}\right]$ in NNL(i),

iii. else if $\mu_{ij}^{\text{neigh}} < \mu_{\min}^{\text{non-neigh}}$ then set $\mu_{\min}^{\text{non-neigh}} \leftarrow \mu_{ij}^{\text{neigh}}$.

iv. If $\mu_{ij}^{\text{neigh}} < \mu_{\max}^{\text{neigh}}$ then set $\mu_{\max}^{\text{neigh}} \leftarrow \mu_{ij}^{\text{neigh}}$.

4. For all interactions $[(j,v), p]$ in NNL(i), do:
   (a) Insert the interaction $[(i,-v), -p]$ in NNL(j).
   (b) If $-p > \mu_{\max}^{\text{neigh}}$ then set $\mu_{\max}^{\text{neigh}} \leftarrow -p$.
5. If NNL(i) is full, then set $\mu_{\text{neigh}} \leftarrow \mu_{\min}^{\text{non-neigh}}$,
6. else set $\mu_{\text{neigh}} \leftarrow \min\left(\mu_{\min}^{\text{non-neigh}}, \mu_{\max}^{\text{neigh}}\right)$.
7. Set $\mathbf{r}_N(i) \leftarrow \mathbf{r}(i)$, $\mathbf{q}_N(i) \leftarrow \mathbf{q}(i)$ and $\mathbf{O}_N(i) \leftarrow \mu_{\text{neigh}}\mathbf{O}(i)$, and also update $O_{\max}^{\text{neigh}} \leftarrow \max\left[O_{\max}^{\text{neigh}}, \mu_{\text{neigh}}O_i\right]$.

3.4.8 Strengths of the Algorithm

Collision-driven molecular dynamics can be used to generate tightly jammed packings of ellipsoid with densities far surpassing previously achieved ones. Several features of the molecular dynamics algorithm are necessary for the success of this packing protocol. First, provisions need to be made to allow time-dependent particle shapes, and we have explicitly included them in the treatment in Section 2.2.3. Most importantly, a very high accuracy collision resolution is necessary at very high densities, especially near the jamming point. For this reason, a time-driven approach cannot be used to generate jammed packings, and special care needs to be taken to ensure high accuracy of the overlap potentials and the time-of-collision predictions, as is done with Newton refinement in our algorithms. Finally, the use of neighbor lists significantly improves the speed of the algorithm since most computation is expended on the last stages of the algorithm when the particles are almost jammed and the use of neighbor lists is optimal (particularly combined with adaptive strategies for controlling μ_{cutoff}). Including a deforming boundary in the algorithm additionally allows for a Parinello-Rahman-like adaptation of the shape of the unit cell (as described in Section 3.2.4.2), which leads to better (strictly jammed [15, 71]) packing of the particles [80].

In this work, we have mostly focused on using near neighbor lists as a tool to improve the efficiency of the collision-driven algorithm. However, using neighbor lists has additional advantages as well. The most important one is that it allows one to monitor the collision history of each particle or a pair of particles. This is especially useful for dense hard-particle packings near the jamming point. In the very limit of a jammed packing, each particle has a certain number of contacting geometric neighbors, and cannot displace from its current position [71]. The network of interparticle contacts forms the *contact network* of the packing, and this contact network can carry positive contact forces. For packings of soft spheres, interacting with a differentiable potential, it is easy to obtain contact forces near a jamming point and observe the resulting force chains and the distribution of contact forces, which has been noted to have an exponential tail in a variety of models of granular materials [16].

3.5 Ellipses and Ellipsoids

Having developed the necessary tools for dealing with overlap between ellipses and ellipsoids in Chapter 2, we can now complete the description of the EDMD algorithm for ellipses and ellipsoids. We first discuss the fundamental step of predicting collisions between moving ellipsoids, and then explain how to process a binary collision between two ellipsoids.

3.5.1 Predicting Collisions

The central step in event-driven MD algorithms is the prediction of the time-of-collision for two moving particles, as well as the time when a particle leaves its bounding neighborhood. This is also the most time-consuming step, especially for nonspherical particles. Although general methods can be developed for particles of arbitrary shape [129], efficiency is of primary concern to us and we prefer specialized methods which utilize the properties of ellipsoids, in particular, their smoothness and the relative simplicity of the time derivatives of the overlap potentials given in Section 2.2.3. In three dimensions, we restrict consideration

to ellipsoids with a *spherically symmetric moment of inertia*, i.e., ellipsoids with equal moment of inertia around all axes. This is because the force-free motion of general ellipsoids, as well as their binary collisions, are very complex to handle. For example, the angular velocity is *not* constant but oscillates in a complex manner. It is not hard to adapt the algorithms presented here to ellipsoids with several different moments of inertia, at least in principle.

Essentially, predicting the collision time t_c between two moving ellipsoids $A(t)$ and $B(t)$ consists in finding the first non-zero root of the overlap potential $\zeta(t) = \zeta[A(t), B(t)]$, where ζ can be either one of ζ_{AB}, $\zeta_{AB}^{(B)}$ or ζ_B, depending on the type of collision and the choice of the potential. Formally:

$$\begin{aligned} t_c &= \min t \\ \text{such that} \quad \zeta(t) &= 0 \quad \text{and } t \geq 0, \end{aligned} \tag{3.7}$$

where $\zeta(t)$ is a smooth continuously differentiable function of time, as illustrated in Fig. 3.7. This kind of first root location problem has wide applications and has been studied in various disciplines. For a general non-polynomial $\zeta(t)$, its rigorous solution is a very hard problem and requires either interval methods [121] or rigorous under/over estimation of $\zeta(t)$ based on knowledge of exact bounds on the Lipschitz constant of $\dot{\zeta}$ (and possibly of ζ) [130]. These methods are rather complex and are focused on robustness and generality, rather than efficiency. For particular forms of $\zeta(t)$, rigorous algebraic methods may be possible, such as for example the prediction of time of collision of two needles (infinitely thin hard rods) [105], and possibly spherocylinders. However, this requires a considerable algebraic complexity and is not easy to adapt to a new particle shape, especially ellipsoids, for which there is not even a closed-form expression for the overlap potential.

In particular, the very elegant method for determining the time of collision of two needles proposed in Ref. [105] is related to the one proposed in Ref. [130], and at its core is the need to determine a good local or global estimate of the Lipschitz constant of $\dot{\zeta}$ [130], i.e., an upper bound on $\left|\ddot{\zeta}\right|$ [these are used to construct rigorous under- or over-estimators of $\zeta(t)$)]. Such a global upper bound has been derived for the case of needles [c.f. Eq. (20) in [105]], but for ellipsoids the expression for $\ddot{\zeta}$ (which we do not give here) is very complex and we have not been able to generalize the approach in Ref. [105]. As discussed in Ref. [130], significantly better results are obtained when local estimates of the Lipschitz constant of $\dot{\zeta}$ are available (i.e., upper bounds on $\left|\ddot{\zeta}\right|$ over a relatively short time interval), and this seems an even harder task. Nevertheless, it is a direction worth investigating in the future.

For the purpose of EDMD, it is sufficient only to ensure that an interval of overlap $[t_c, t_c + \Delta t_c]$ is not missed if

$$\min_{t_c \leq t \leq t_c + \Delta t_c} \zeta(t) < -\epsilon_F,$$

where ϵ_F is some small tolerance, typically $10^{-4} - 10^{-3}$ in our simulations, or alternatively, if $\Delta t_c > \epsilon_t$. The use of ϵ_F is preferable because it is dimensionless with a scale of order 1. This essentially means that it is permissible to miss *grazing collisions*, i.e., collisions in which two ellipsoids overlap for a very small amount and/or for a very short time. It certainly is not productive to try to decide if two nearly touching particles are actually overlapping more accurately then the inherent numerical accuracy of $\zeta(t)$. The choice of ϵ_F is determined by the relative importance of correctness versus speed of execution, as well as the stability of the simulation. A large ϵ_F can lead to unrecoverable errors in the event-driven algorithm, such as runaway collisions or increasing overlap between particles.

Homotopy methods can be used to solve problems such as (3.7). They typically trace the evolution of the root of an equation (starting from $t = 0$ in this particular case) as the equation is deformed from an initial simple form to a final form which matches $\zeta(t) = 0$ [131]. An ordinary differential equation (ODE) solver can be used for this purpose. An essential component in these methods is *event location in ODE*s, namely, methods which solve an ODE for a certain variable $f(t)$ and determine the first time that $f(t)$ crosses zero [91, 132]. We have tested a (simple) ODE-based homotopy method for solving (3.7), however, since the problem at hand is one-dimensional, one can directly apply ODE event location to $f(t) \equiv \zeta(t)$, using an absolute tolerance of ϵ_F for the ODE solver, and locate the first root t_c directly more efficiently.

The ODE to solve is given by eq. (2.21) or (2.23). However, also needed is $\lambda(t)$, and one has the option of either explicitly evaluating λ at each time step (reusing the old value of λ as an initial guess), or also

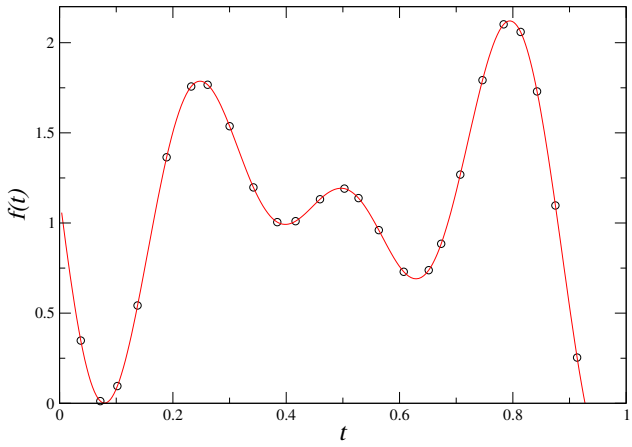

Figure 3.7: The time evolution of the overlap function $\zeta_{AB}(t)$ during an ellipse collision. The overlap function is evaluated on an adaptive grid, which has a smaller time step when ζ_{AB} changes rapidly, and a larger step when it is relatively smooth. The tracing stops when a zero crossing is detected.

including $\lambda(t)$ in a system of ODEs using eq. (2.22) or (2.24). The second option has the advantage that one no longer needs to explicitly evaluate λ (other than at the beginning of the integration), however, it has the additional cost of two variables instead of one in the ODE solver, which additionally leads to smaller time steps in the ODE integrator. Our numerical experiments have indicated that at least in two and three dimensions it is somewhat advantageous to explicitly evaluate λ and only include $\zeta(t)$ in the ODE. This may be reversed for different particle shapes, depending on the relative complexity of evaluating λ versus evaluating $\dot{\lambda}$.

Once λ is evaluated however, very little extra effort is needed to evaluate ζ explicitly, so it seems somewhat pointless to solve an ODE for $\zeta(t)$. We have developed a method with a similar structure to ODE integration, but which uses explicit evaluation of both ζ and $\dot{\zeta}$. The basic idea is to take a small time step Δt, evaluate both ζ and $\dot{\zeta}$ at the beginning and end of the step, and use these values to form a cubic Hermite interpolant $\tilde{\zeta}(t)$ of $\zeta(t)$ over the interval Δt. A theoretically-supported estimate for the absolute error of the interpolant can be obtained by comparing the interpolant $\tilde{\zeta}$ and ζ at the midpoint of the time step, and this error can be used to adaptively increase or decrease the size of the step so as to keep the absolute error within ϵ_F. This is illustrated in Fig. 3.7. When the interpolant crosses the $\zeta = 0$ axes, the first root of the interpolant is used as an initial guess in a safeguarded Newton algorithm to find the exact root t_c. The initial time step Δt needs to be sufficiently small to make the initial error estimate valid, and can easily be obtained by estimating a time-scale for the collision from the sizes and velocities of the particles involved in the collision. Even if this initial guess is conservative, the algorithm quickly increases the step to an appropriate value. We will refer to this algorithm as *trace event location*, since the function $\zeta(t)$ is explicitly traced until a zero-crossing is found.

There are several details one needs to be attentive to when predicting collisions inside an EDMD algorithm. For example, some pairs of particles may already be overlapping by small amounts after having collided. In this case one can look at the sign of $\dot{\zeta}$ to decide whether the two particles are about to have a collision or just had a collision. Two particles *can* have a collision after having collided *without* an intervening collision with third party particles. This can always happen for aspherical particles, as noted in Ref. [105], and it can also happen for spheres when boundary deformations are present (since the particles travel along curved paths), however, it cannot happen for spheres in traditional EDMD algorithms. If the initial ζ is very close to zero, an additional safety measure is to add a small positive correction to ζ to ensure that it is sufficiently far from zero at the beginning of the search [as compared to the accuracy in the evaluation of $\zeta(t)$]. In particular, such precautions are necessary at very high densities (i.e., near jamming).

3.5.1.1 Predicting Collisions for Ellipsoids

We sketch the procedure for predicting collisions between two moving ellipsoids in Algorithm 11. Features in Algorithm 11 include limiting the number of steps in the event location algorithm to avoid wasting resources on predicting collisions that may never happen, as well as allowing the exact prediction to fail. This algorithm first uses collision prediction for the bounding (or contained) spheres, in order to eliminate obvious cases when the particles do not collide, and to identify a short search interval for the event location by calculating the time interval during which the enclosing spheres overlap. Recall that when the boundary is not deforming this step entails solving a quadratic equation, while it involves solving a quartic equation when the lattice velocity is nonzero. In this context not only the first root of this quadratic/quartic equation needs to be determined, but also the second one, giving the interval of overlap. It may be possible to further improve the initial collision prediction step by using a bounding body other than spheres, for example, oriented bounding boxes (orthogonal parallelepipeds) [133, 134]. However, orientational degrees of freedom need to be eliminated since they are too hard to deal with because of the appearance of trigonometric functions. For example, the bounding body can be a cylinder whose axis is the axis of rotation of the ellipsoid and whose radius and length are sufficient to bound the rotating ellipsoid for all angles of rotation.

An important problem we have encountered in practice is the numerical evaluation of ζ_B when predicting the time of collision of an ellipsoid with its bounding neighborhood. Namely, as the particles move, it is highly likely that a point where $\mathbf{Y}$ is singular will be encountered. At such points the evaluation of ζ_B is numerically unstable and often leads to unacceptably small time steps. We have dealt with this problem in an *ad hoc* manner, by simply trying to skip over such points, so that the search for a collision can be continued, but a more robust and systematic approach may be possible.

---

1. For all intervals of overlap of the bounding spheres of A and B (or for all intervals during which the bounding sphere of A intersects the shell between the contained and bounding sphere of its bounding neighborhood B), $[t_{\text{start}}, t_{\text{end}}]$, starting from $t = 0$ and in the order of occurrence, do:
   (a) If $t_{\text{start}} > T$, return reporting that no collision can happen.
   (b) Set $t_{\text{end}} \leftarrow \min\{T, t_{\text{end}}\}$.
   (c) Update the ellipsoids to time t_{start}, and evaluate the initial ζ and λ.
   (d) If $\zeta < \epsilon$ (ellipsoids are overlapping or nearly touching), then evaluate $\dot{\zeta}$ and:
      i. If $\dot{\zeta} \geq 0$ then set $\zeta \leftarrow \epsilon$ (the ellipsoids are moving apart),
      ii. Else return $t_c = t_{\text{start}}$ (the ellipsoids are approaching).
   (e) Use trace event location to obtain a good estimate of the first root of $\zeta(t)$ during the interval $[t_{\text{start}}, t_{\text{end}}]$, putting a limit on the number of time steps (for example, in the range $100 - 250$). If no root crossing is predicted, continue with the next interval in step 1. If the search terminated prematurely, then return the last recorded time $\tilde{t_c} = t$.
   (f) Bracket the estimated first root of $\zeta(t)$ and refine it using a safeguarded Newton's method (this may fail sometimes).
   (g) If the root refinement failed, set $t_{\text{end}} \leftarrow \frac{1}{2}(t_{\text{start}} + t_{\text{end}})$, repeat step 1c and go back to step 1e (attempt to at least find a valid $\tilde{t_c}$).
2. Return reporting that no collision will happen.

Algorithm 11: Predict whether two moving ellipsoids overlap during the time interval $[0, T]$ by more than ϵ_F, and if yes, calculate the time of collision $t_c \in [0, T]$. Essentially the same procedure can be used to determine the time a particle A collides with its bounding neighborhood B. If the prediction cannot be verified, return a time $\tilde{t_c} < t_c$ before which a collision will not happen. Also return λ at the time of collision if desired.

3.5.2 Processing Binary Collisions

The steps necessary to process a binary collision between two hard particles are similar for a variety of particle shapes, and essentially involves exchanging momentum between the two particles. We give a recipe

for colliding ellipsoids with a spherically symmetric moment of inertia in Algorithm 12. To determine things like the pressure it is useful to maintain the collisional contribution to the stress $\boldsymbol{\sigma}_c$ [107], which is a suitable average of the exchange of momentum over all collisions.

---

1. Calculate the Euclidean positions and velocities of particles $\mathbf{r}_A$, $\mathbf{r}_B$, $\mathbf{v}_A$ and $\mathbf{v}_B$, as well as their angular velocities $\boldsymbol{\omega}_A$ and $\boldsymbol{\omega}_B$.
2. Find the contact point $\mathbf{r}_C$ and normal velocity at the point of contact $v_n = \hat{\mathbf{n}}^T \mathbf{v}_C$, using the supplied λ.
3. If $v_n > 0$, then return without further processing this (most likely grazing) mis-predicted collision.
4. Calculate the exchange of momentum between the particles $\Delta \mathbf{p}_{AB} = \Delta p_{AB} \hat{\mathbf{n}}$,

$$\Delta p_{AB} = 2 v_n \left(\frac{1}{m_A} + \frac{1}{m_B} + \frac{\|\mathbf{r}_{CA} \times \hat{\mathbf{n}}\|}{I_A} + \frac{\|\mathbf{r}_{CB} \times \hat{\mathbf{n}}\|}{I_B} \right)^{-1},$$

   where m denotes mass and I moment of inertia.

5. Calculate the new Euclidean velocities of the particles,

$$\begin{aligned} \mathbf{v}_A &\leftarrow \mathbf{v}_A - \frac{\Delta p_{AB}}{m_A} \hat{\mathbf{n}} \\ \mathbf{v}_B &\leftarrow \mathbf{v}_B + \frac{\Delta p_{AB}}{m_B} \hat{\mathbf{n}}, \end{aligned}$$

   as well as the new angular velocities

$$\begin{aligned} \boldsymbol{\omega}_A &\leftarrow \boldsymbol{\omega}_A - \frac{\Delta p_{AB}}{I_A} \left(\mathbf{r}_{CA} \times \hat{\mathbf{n}} \right) \\ \boldsymbol{\omega}_B &\leftarrow \boldsymbol{\omega}_B + \frac{\Delta p_{AB}}{I_B} \left(\mathbf{r}_{CB} \times \hat{\mathbf{n}} \right). \end{aligned}$$

   Optionally update any averages that may need to be maintained (such as average kinetic energy) to reflect the change in the velocities.

6. Record the collisional stress contribution

$$\boldsymbol{\sigma}_c \leftarrow \boldsymbol{\sigma}_c - \Delta p_{AB} \left(\mathbf{r}_{AB} \hat{\mathbf{n}}^T \right).$$

7. If using NNLs, record information about the collision that is being collected for the interaction between i and (j, v), such as an accumulation of the total exchanged momentum for this interaction, total number of collisions for this interaction, etc.

Algorithm 12: Process a binary collision between ellipsoids i and (j, v). Assume that the particles have already been updated to the current time t, and that λ at the point of collision is supplied (i.e., it has been stored for particle i , and also $1 - \lambda$ in particle j, when this collision was predicted).

3.6 Performance Results

In this section we present some results for the performance of the algorithm. Many previous publications have given performance results for EDMD for spheres, and most of these results apply to our algorithm. Exact numbers depend critically on details of the coding style, programming language, compiler, architecture, etc., and are not reported here. Rather, we try to get an intuitive feeling of how to choose the various parameters of the simulation to improve the practical performance. Our main conclusion is that using NNLs is significantly more efficient than using just the cell method for particles with aspect ratio significantly different from one

(greater than 2 or so) or at sufficiently high densities. Additionally, using BSCs offers significant efficiency gains for very prolate particles, for which good bounding sphere complexes can easily be constructed.

As derived in Ref. [104], when only the cell method is used, optimal complexity of the hard-sphere EDMD code is obtained when the number of cells is of the order of the number of particles, $N_c = \Theta(N)$, with asymptotic complexity $O(\log N)$ per collision, which comes from the event-heap operations. In practice however the asymptotic logarithmic complexity is not really observed, and instead to a very good approximation the computational time expanded per processed *binary collision* is constant for a given aspect ratio α at a given density, for a wide range of relative densities (volume fractions) φ. Even though in principle the basic EDMD algorithm remains $O(\log N)$ per collision, our aim is to improve the constants in this asymptotic form, and in particular, their dependence on the shape of the particle (in particular, the aspect ratio α).

It is important to note that on modern serial workstations, the EDMD algorithm we have presented here is almost entirely CPU-limited, and has relatively low memory requirements, even when using NNLs and BSCs. Floating point operations dominate the computation, but memory traffic is also very important. In our implementation, simulating ellipsoids is about an order of magnitude slower than simulating spheres, even for nearly spherical ellipsoids, simply due to the high cost of the collision prediction algorithm (the same observation is reported in Ref. [105]) and increased memory traffic. For example, on a 1666 MHz Athlon running Linux, our Fortran 95 implementation uses about 0.1ms per sphere collision for a wide range of system sizes and densities. With all the improvements described in this work, and in particular, the use of NNLs and BSCs, our implementation uses about 2ms per ellipsoid collisions for prolate spheroids, and about $2-4$ms for oblate spheroids at moderate densities for a wide range $\alpha = 1 - 10$. Including boundary deformations, i.e., solving quartic instead of quadratic equations when predicting binary collisions, slows down the simulation for spheres by about a factor of 2.5 (we use a general quartic solver, and better results may be obtained for a specialized solver).

3.6.1 Tuning the NNLs

We have performed a more detailed study of the performance of the algorithm when NNLs are used, since this is a novel technique and has not been analyzed before. We perform an empirical study rather then a theoretical derivation because such a derivation is complicated by the fact that the neighborhoods evolve together with the particles, and because the numerous constant factors or terms hidden in the asymptotic expansions of the complexity actually dominate the practical performance.

Computationally, we have observed that it is good to maximize the number of cells N_c, even at relatively low density $\varphi \approx 0.1$, especially for rather aspherical particles. This is because binary collision predictions become much more expensive than predicting or handling transfers, and so the saving in not predicting collisions unnecessarily offsets the higher number of transfers handled. Consistent with the results reported in [104], we observe that the number of checks due to invalidated event predictions is comparable and sometimes slightly larger than the number of collisions processed, and this suggests that additional improvements in this area might increase performance noticeably.

We have tested both methods for updating the neighbor lists, the *complete* and the *partial* update. A complete update/rebuild of the near neighbor lists after they become invalid is the traditional approach in most TDMD algorithms appearing in the literature. Since MD is usually performed on relatively homogeneous systems, when one particle displaces by a sufficient amount to protrude outside of its bounding neighborhood, most particles will have displaced a significant amount, and so rebuilding their NNLs is not so much of a waste of computational effort. The main advantage of this approach is that it can be used to build NNLs when a good estimate of μ_{cutoff} is not available, but rather a bound on the number of neighbors per particle N_i is provided (this is very useful, for example, in the very early stages of the Lubachevsky-Stillinger algorithm, when particles grow very rapidly). Additionally, the algorithm for rebuilding the lists is simpler and thus more efficient. Finally, a complete rebuilding of the NNLs yields neighbor lists of higher quality, in the sense that the structure of the network of bounding neighborhoods is better adapted to the current configuration of the system and thus the size of the neighborhoods is maximal.

The algorithm for partial updates on the other hand is more complicated, and to our knowledge has not been used in MD codes. It requires using dynamic linked lists, and it will in general yield smaller average neighborhood size than a complete update, since the particle whose NNL is being updated must adjust its list without perturbing the rest of the NNLs. Note that at the beginning the NNLs must be initialized by using

a complete update. The main advantage of the partial update scheme is that it is more flexible in handling nonisotropic systems or the natural fluctuations in an isotropic one. Just because one particle happened to move fast and leave its neighborhood does not mean that all particles move that fast. This is especially true at lower densities where clustering happens. In clusters particles have more collisions per unit time and thus require fewer updates of the NNLs, but outside clusters particle move large distances without collisions and thus require more frequent updates of their NNLs. In this sense partial updates are local in nature while complete updates are global. We have indeed observed that in most cases it is advantageous to use partial updates, rather then the traditional complete updates.

The first and most important test is to determine whether using near neighbor lists offers any advantages over using just the cell method. The following intuitive arguments seem clear:

- At very high densities, when the system of particles is nearly jammed [15], using NNLs is optimal, regardless of the aspect ratio of the particles. This is because the particles move very little while they collide with nearly the same neighbor particles over and over again (see Fig. 1.1). Therefore near jamming the NNLs are rarely updated and by predicting the collisions only with the particles with which actual collisions happen significant savings can be obtained. However, as the density is lowered, the lists need to be updated more frequently and the complexity of using NNLs becomes significant.

- At very low densities the cell method is faster, even for very aspherical particles. This is because a particle will have many collisions with the bounding neighborhoods before it undergoes a binary collision, so that the cost is dominated by the cost of maintaining the NNLs instead of processing collisions.

- The more aspherical the particles, the more preferable the NNL method becomes compared to the cell method. This is because for very elongated particles at reasonably high densities there will be many particles per cell so using the cell method will require predicting many binary collisions that will never happen, while the NNL method will predict collisions with significantly fewer (truly) neighboring particles. For large α, the dominating cost is that of rebuilding the NNLs (since this step uses the cells), and therefore the primary goal becomes to minimize the number of NNL updates per number of binary collisions processed, as well as to improve the efficiency of the NNL rebuild, i.e., using BSCs.

Our experimental results shown in Fig. 3.8 support all of these conclusions. We show the ratio of the CPU time expanded *per processed binary collision* for the NNL method and for the traditional cell method. We show results for equilibrium systems of prolate spheroids of aspect ratios $\alpha = 1, 3$ and 5 at densities $\varphi = 0.1, 0.3, 0.5$ and 0.6. Note that hard spheres jam in a disordered metastable state at around $\varphi = 0.64$, and that for the case $\alpha = 5$ we use $\varphi = 0.55$, since the jamming density is slightly lower than 0.6 for this aspect ratio [78]. In Fig. 3.8(a) we show the relative *slowdown* caused by using NNLs for the systems for which using the cell method is better. In Fig. 3.8(b) we show the relative *speedup* obtained by using NNLs for the systems for which it is better to use NNLs. Both the results of using partial and complete updates are shown.

As explained earlier, we use two techniques to limit the number of neighbors that enter in the NNLs. The first one is to simply use the upper bound on the number of neighbors (interactions) N_i to choose only the nearest N_i neighbors per particle, and the second one is to choose a relatively small cutoff μ_{cutoff} for the maximal size of the bounding neighborhood $\mathcal{N}(i)$ (compared to the size of the particle i). In practice, only the second approach can be used with partial updates. This is because partial updates must work under the limitations of doing as little change to the NNLs as possible, and this requires that there be enough room to add and remove interactions from the lists as necessary. So when using partial updates one must set μ_{cutoff} to a reasonable value and then set N_i to be larger then the maximal number of neighbors a particle will have given the cutoff μ_{cutoff} and an unlimited N_i. Reasonable values for spheres and not too aspherical particles are to set μ_{cutoff} so that on average each particle has about $5 - 7$ neighbors in two or $11 - 15$ neighbors in three dimensions (the kissing number for spheres is 6 in two and 12 in three dimensions), while setting N_i at about 10 in two and 20 in three dimensions.

Since we wish to compare partial and complete updates, we change μ_{cutoff} and always set N_i to a sufficiently high number (which grows sharply with μ_{cutoff}), and compare partial and complete updates in Fig. 3.8. As expected, there is an optimal value of μ_{cutoff} which is larger for complete updates (for which

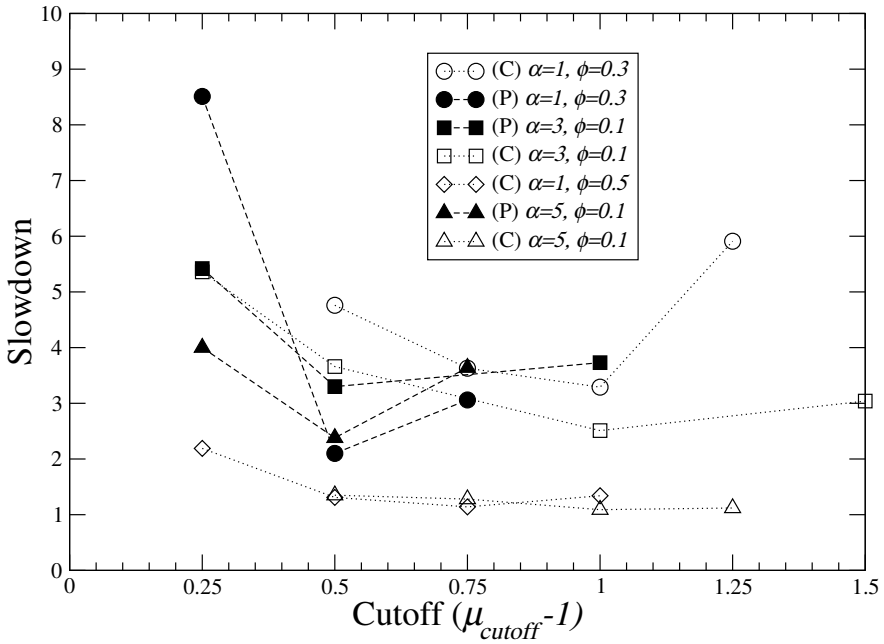

(a) Cell method wins

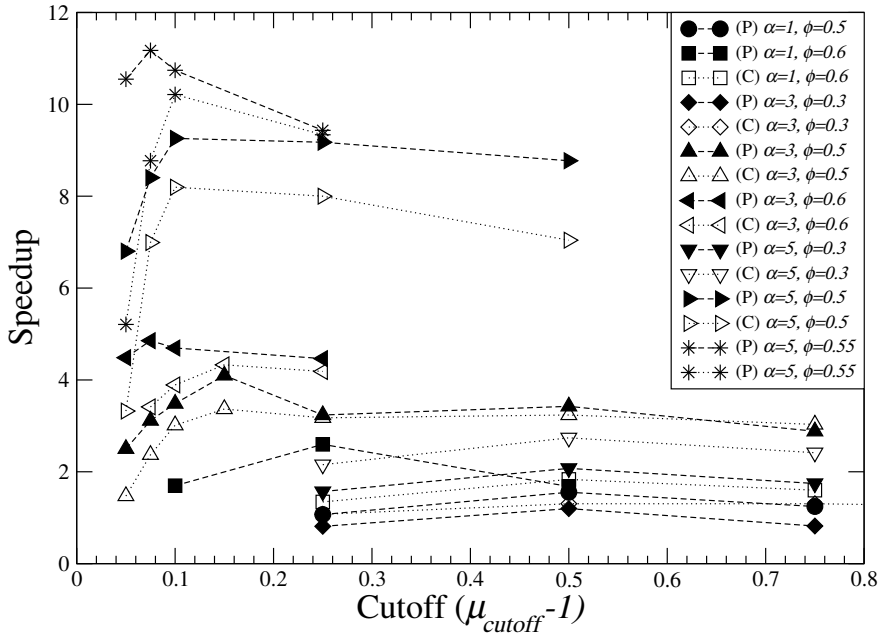

(b) NNL method wins

Figure 3.8: Performance results for using NNLs in addition to the traditional cell method for a variety of aspect ratios α and densities φ for prolate spheroids, with both partial (P) or complete (C) updates.

NNL updates are significantly more expensive) and also at lower densities. Note however that sometimes the computation speed may change discontinuously as μ_{cutoff} is increased because at some point more than first-neighbor cells need to be searched during the NNL update. Important observations to note include the fact that *tuned partial NNL updates almost always outperform tuned complete updates*, and are thus preferred. Another useful observation is that the computation time is not very sensitive to the exact value of μ_{cutoff}. Finally, note that as much as an order of magnitude of improvement is achieved for rather aspherical particles at high densities by using the NNLs. This would be even more pronounced for larger aspect ratios such as $\alpha = 10$.

For large aspect ratios, the dominant cost is that of rebuilding the NNLs, during which many particle pairs need to be tested for neighborhood. Therefore, the most important factor for the speed of the simulation is how many particles need to be examined as potential near neighbors of a given particle i when rebuilding NNL(i). If bounding spheres are not used, this number is proportional to the number of particles that can fit in a cell of length α, i.e., a cube of volume α^3. For prolate spheroids this number is proportional to α^2, but for oblate ones it is only proportional to α. Therefore the simulation of, for example, $\alpha = 10$, is prohibitively expensive for prolates, but not for oblates. As the results in Fig. 3.9 demonstrate, using BSCs significantly increases the speed of processing collisions for very prolate spheroids. In this figure we show the approximate speedup obtained by using BSCs for tuned partial updates for a range of moderate densities and a range of aspect ratios. As expected, at large densities there are very few updates to the NNLs and therefore using BSCs does not offer a large speedup. For oblate spheroids in the same range of aspect ratios, using BSCs does not offer computational savings, and therefore we do not show any performance results. However, it is important to note that when using NNLs (at sufficiently high densities) and BSCs (at sufficiently high aspect ratios) the actual (tuned) processing time per collision is approximately the same for any ellipsoid shape in this range of aspect ratios, somewhere in the range $1 - 5$ms per binary collision in our implementation. Therefore, for practical purposes, we feel that the algorithms presented here can handle a wide range of ellipsoid shapes very well.

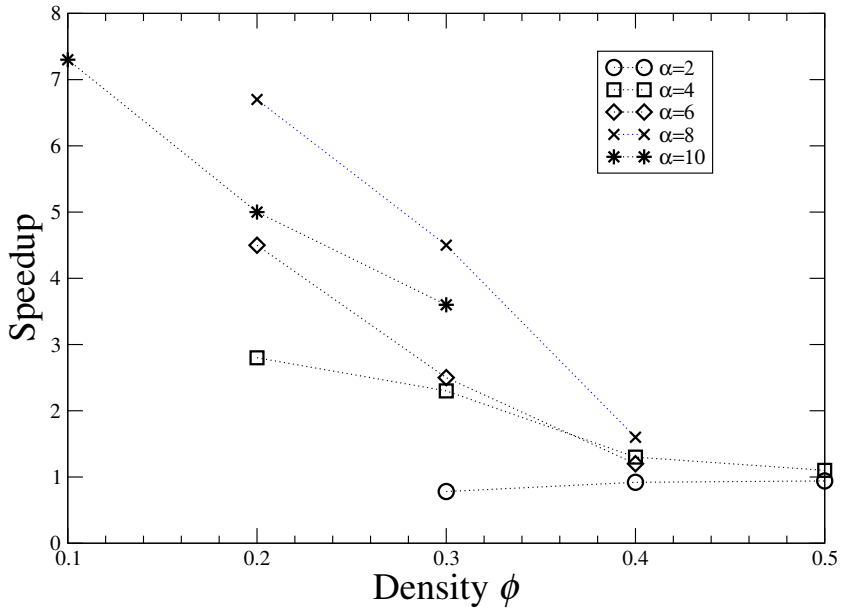

Figure 3.9: Performance results for using BSCs for prolate spheroids at low to moderate densities. Using BSCs does not appear to offer computational savings for oblate spheroids in the same α range.

With the use of BSCs, prolate ellipsoids are handled much better and the scaling reduced to nearly independent of α (note that one needs to examine α bounding spheres per particle, which is much less expensive than looking at neighbor particles, but still not free), as we have demonstrated above. It remains a challenge to find a technique that will also reduce the scaling to nearly independent of α for oblates as

well. Additionally, it is important to develop a theoretical analysis of the performance of the novel steps in the algorithm, and in particular, to give estimates of the number of particles which need to be examined when building the NNLs (per particle), the number of NNL updates which need to be processed per binary collision (per particle).

3.6.1.1 Automatic Tuning of μ_{cutoff}

It is important to note that it is possible to automatically tune μ_{cutoff} during the course of the simulation, at least in a rough way, so that the optimal computation speed is approached. This is very important in the Lubachevsky-Stillinger packing algorithm, since there the density is not constant but rather increases until jamming is reached. Clearly in the beginning a larger μ_{cutoff} is needed, while near the jamming point μ_{cutoff} can be set very close to 1. By monitoring the fraction of events which are collisions with a bounding neighborhood, and other statistics, one can periodically adaptively increase or decrease μ_{cutoff} during the course of the simulation. We have successfully used such techniques to speed the process of obtaining hard-particle packings, for both spheres and ellipsoids, especially for elongated ellipsoids, but do not report details here.

3.7 Conclusions

We presented a serial collision-driven molecular dynamics algorithm for nonspherical particles, with a specific focus on improving the efficiency by developing novel techniques for neighbor search. In particular, we developed a rigorous scheme that incorporates near-neighbor lists into event-driven algorithms, and further improved the handling of very elongated objects via the use of (non-hierarchical) bounding sphere complexes. We gave detailed pseudocodes to illustrate the major steps of the algorithm. We have proposed a general method for determining the time of collision of two particles of any shape for which a smooth overlap potential can be constructed and easily differentiated. The application to ellipses or ellipsoids has been developed in detail.

We have identified a number of important directions for future investigation, which can lead to significantly faster algorithms for very aspherical particles. First, predicting the time of collision of two moving particles can be improved, either by using techniques other than bounding spheres in order to narrow the search intervals during which a collision may happen, or by improving the algorithm to search those intervals for a collision. In particular, a new overlap potential for the case of one small ellipsoid contained within another large ellipsoid is needed. The practical handling of bounding sphere complexes can further be improved. More importantly, it is an open challenge to develop an algorithm to improve the efficiency of building the near-neighbor lists for very oblate particles.

Chapter 4

Jamming in Hard-Sphere Packings

Jamming in hard-particle packings has been the subject of considerable interest in recent years. In a paper by Torquato and Stillinger [15], a classification scheme of jammed packings into hierarchical categories of *locally*, *collectively* and *strictly jammed* configurations has been proposed. They suggest that these jamming categories can be tested using numerical algorithms that analyze an equivalent contact network of the packing under applied displacements, but leave the design of such algorithms as a future task. In this chapter, we present a *rigorous* and *practical* algorithm to assess whether an ideal hard-sphere packing in two or three dimensions is jammed according to the aforementioned categories [71]. The algorithm is based on linear programming and is applicable to regular as well as random packings of finite size with hard-wall and periodic boundary conditions. If the packing is not jammed, the algorithm yields representative multi-particle unjamming motions. Furthermore, we extend the jamming categories and the testing algorithm to packings with significant interparticle gaps. We describe in detail two variants of the proposed randomized linear programming approach to test for jamming in hard-sphere packings. The first algorithm treats ideal packings in which particles form perfect contacts. Another algorithm treats the case of jamming in packings with significant interparticle gaps. This extended algorithm allows one to explore more fully the nature of the feasible particle displacements.

We have implemented the algorithms and applied them to ordered as well as random packings of disks and spheres with periodic boundary conditions [72]. The random packings were produced computationally with a variety of packing generation algorithms, all of which should, in principle, produce at least collectively jammed packings. Our results highlight the importance of jamming categories in characterizing particle packings. One important and interesting conclusion is that the amorphous monodisperse sphere packings with density $\varphi \approx 0.64$ were for practical purposes strictly jammed in three dimensions, but in two dimensions the monodisperse disk packings at previously reported "random close packed" densities of $\varphi \approx 0.83$ were not even collectively jammed. On the other hand, amorphous bidisperse disk packings with density of $\varphi \approx 0.84$ were virtually strictly jammed. This clearly demonstrates one cannot judge "stability" in packings based solely on local criteria. Numerous interactive visualization models are provided online [135].

In this chapter we focus on the special case of hard spheres, for which impenetrability constraints between the particles can be written analytically and it is not necessary to introduce overlap potentials. In the next chapter we will generalize the concept of jamming and extend this work to nonspherical particles. Much of the concepts and details presented here apply in that case as well and we will not repeat them. However, we will see that one of the properties used extensively in this chapter, namely, that first-order expansion of the impenetrability constraints is sufficient, does not extend to nonspherical particles and is responsible for some surprising properties of jammed packings of ellipsoids.

4.1 Introduction

We focus our attention on the venerable idealized hard-sphere model, i.e., the only interparticle interaction is an infinite repulsion for overlapping particles. The idealized hard-sphere model is in a sense the Ising model for modeling a variety of hard-particle physical systems, and the importance of understanding it in detail cannot be overstated. The singular nature of the interaction potential enables us to be precise about

the important concept of "jamming." In particular, a hierarchical classification scheme for jammed packings into *locally*, *collectively* and *strictly jammed* packings was proposed in Ref. [15]. This classification is closely related to the concepts of *rigid* and *stable* packings found in the mathematics literature [136, 137]. The term jamming is used in a different sense in the modeling of granular media, which includes effects such as friction, adhesion, particle deformability, etc., and, by definition, hard-sphere systems do not include these effects. It is also important to note that we do not discuss dynamical effects in hard-particle packings. In the present work, hard-sphere jamming is presented from a rigorous perspective that focuses on the *geometry* of the *final* packed states. The ideas that we employ here are drawn heavily from the mathematics literature [138, 136, 73, 137].

In Section 4.2, we present the conceptual theoretical framework underlying this work. Specifically, we review and expand on the hierarchical classification scheme for jammed packings into locally, collectively and strictly jammed packings proposed in Ref. [15]. In Section 4.3, we present a randomized linear programming algorithm for finding unjamming motions within the approximation of small displacements, focusing on periodic boundary conditions in section 4.4.3. This algorithm is rigorous when applied to *ideal* packings, where interparticle gaps are very small. In section 4.3.3, we extend the concepts of jamming and the randomized linear programming algorithm to packings that have significant interparticle gaps and do not fit well in the rigorous framework suitable for ideal packings. We also introduce a randomized sequential loading algorithm to study nonideal packings. We discuss the two algorithms in detail in Section 4.6, and give some representative illustrations and timing statistics in order to illustrate the utility of the proposed algorithms. In Section 4.7 we discuss the numerical implementation, and provide results for ordered periodic lattice packings and random packings. Through numerical investigations, we show here that several previously used packing algorithms generate collectively jammed packings under appropriate conditions. In particular, we study in detail monodisperse sphere as well as monodisperse and bidisperse disk packings produced by the Lubachevsky-Stillinger packing algorithm [12]. We also tested a sample of monodisperse sphere and bidisperse disk packings produced by the algorithm described in Ref. [31], as well as monodisperse sphere packings produced by the Zinchenko packing algorithm [28], and observed similar behavior as for the Lubachevsky-Stillinger packings.

Our testing of these packings enables us to arrive at several important conclusions. First, we find that the amorphous monodisperse sphere packings (with covering fraction, or density, $\varphi \approx 0.64$) and bidisperse disk packings ($\varphi \approx 0.84$) are practically strictly jammed (though not in the ideal sense). Second, we observe that large monodisperse disk packings are invariably highly crystalline ($\varphi \approx 0.88$) and are only collectively jammed. Previously reported [41] low covering fractions for "random close packed" disks of $\varphi \approx 0.82 - 0.84$ were not even found to be collectively jammed. This conclusion clearly demonstrates that the distinctions between the different jamming categories are important and one cannot judge "stability" in packings based solely on local criteria, as has been done extensively in the literature [139, 140, 141].

4.2 Jamming in Hard-Sphere Packings

The physical intuition behind the word jamming is strong: It connotes that a given configuration is "frozen" or "trapped". Two main approaches can be taken to define jamming, *kinematic* or *static*. In the kinematic approach, one considers the motion of particles away from their current positions, and this approach is for example relevant to the study of flow in granular media (in particular, the cessation of flow as jamming is approached). The term *jammed* seems most appropriate here. In the static approach, one considers the mechanical properties of the packing and its ability to resist external forces (in particular, the infinite elastic moduli near jamming). The term *rigid* is often used among physicists in relation to such considerations. However, due to the correspondence between kinematic and static properties, i.e. strains and stresses, these two different views are largely equivalent. The different perspectives on jamming are revisited in Chapter 5.

We largely adopt a kinematic approach, as we focus on the geometry of packings, but the reader should bear in mind the inherent ties to static approaches. We first give a general approach to jamming in hard-particle packings in Section 4.2.1, and then focus on the fundamental and rigorous case of packings with ideal interparticle contacts (i.e., no interparticle gaps) in Section 4.2.2, studied both in the physics and mathematics literature. Finally, we connect these definitions to the kinematic concept of unjamming motion in section 4.2.3, and also to static concepts in section 4.2.4. Since we are attempting to bring together several

apparently different approaches and terminologies, as well as generalize to packings with interparticle gaps, the exposition will be gradual and more detailed discussion, illustrations and proofs are delayed to later parts of this Chapter.

4.2.1 Jamming as Isolation in Configuration Space

A *hard-particle packing* $\mathcal{P}(\mathbf{R})$ is characterized by the positions and orientations of N nonoverlapping particles, which give the *configuration* $\mathbf{R}$. In particular, a *sphere packing* in a finite region in d-dimensional Euclidean space $\Re^d$ is characterized only by the positions of the sphere centers $\mathbf{R} = (\mathbf{r}_1, \ldots, \mathbf{r}_N)$,

$$\mathcal{P}(\mathbf{R}) = \left\{ \mathbf{r}_i \in \Re^d, \;\; i = 1, \ldots, N \; : \; \|\mathbf{r}_i - \mathbf{r}_j\| \geq \frac{D_i + D_j}{2} \; \forall j \neq i \right\},$$

where the diameter of the i^{th} sphere is D_i. Two configurations are identical if all interparticle distances are the same, i.e., if the configurations are related via a rigid-body motion (and possibly a mirror inversion in addition). We focus here on monodisperse (i.e., $D_i = D = \text{const.}$) hard-sphere packings for simplicity, but the conclusions are applicable to polydisperse sphere packings as well.

Our perspective on jamming focuses on the set $\mathcal{J}_{\Delta\mathbf{R}}$ of configurations around a particular initial configuration $\mathbf{R}$ reachable via continuous displacements $\Delta\mathbf{R}$ of the spheres, subject to nonoverlapping constraints and certain boundary conditions. An illustration of this is provided in Fig. 4.1 for a very simple case in which only one disk is free to move, i.e., there are only two degrees of freedom. If $\mathcal{J}_{\Delta\mathbf{R}}$ is isolated in configuration space, we call it a *jamming basin*, and the configuration $\mathbf{R} \in \mathcal{J}_{\Delta\mathbf{R}}$ determines a *jammed packing* $\mathcal{P}(\mathbf{R})$. To relate this to the physical intuition of jamming, we must further ask that the extent of $\mathcal{J}_{\Delta\mathbf{R}}$ be small, in the sense that only small *continuous* displacements of the particles from their initial configurations are possible for all $\mathbf{R} \in \mathcal{J}_{\Delta\mathbf{R}}$. The natural length scale defining the meaning of "small" is the typical size of the particles, and also the typical size of the interparticle gaps. A more strict mathematical definition of jamming considers packings that have M perfect interparticle contacts, which we will call *ideal packings*. For a jammed ideal packing $\mathbf{R}$ is an isolated *point* in configuration space, i.e., $\mathcal{J}_{\Delta\mathbf{R}} = \{\mathbf{R}\}$, so that the particles cannot at all be displaced continuously from their current configuration (modulo trivial rigid-body motions). We focus first on ideal packings, and we will return to the issue of interparticle gaps later. By changing the boundary conditions, we get several different categories of jamming, namely local, collective and strict jamming.

4.2.2 Three Jamming Categories

First we repeat, with slight modifications as in Ref. [111], the definitions of several hierarchical jamming categories as taken from Ref. [15], and later we make them mathematically specific and rigorous for several different types of sphere packings. When defining jamming, one must be very specific about the type of boundary conditions imposed on the packing, for example, the packing may be contained inside a hard-wall container. For now we simply assume that some boundary conditions are imposed, and we specialize the meaning of the terms boundary and boundary deformation for specific types of packings in the next section.

A finite system of spheres is:

Locally jammed Each particle in the system is locally trapped by its neighbors, i.e., it cannot be translated while fixing the positions of all other particles. This definition is analogous to the definition of *1-stability* in Ref. [136]. Because of its simplicity, this definition has been overused to obtain theoretical estimates of the density of random packings [22, 139].

Collectively jammed Any locally jammed configuration in which no subset of particles can simultaneously be continuously displaced so that its members move out of contact with one another and with the remainder set. An equivalent definition is to require that all finite subsets of particles be trapped by their neighbors. Compare this to the definition of *finite stability* in Ref. [136].

Strictly jammed Any collectively jammed configuration that disallows all globally uniform volume-nonincreasing deformations of the system boundary. Note the similarity with collective jamming but with the additional provision of a deforming boundary. This difference and the physical motivations behind it should become clearer in section 4.4.5. Compare this to the definition of *periodic stability* in Ref. [136] for packings with periodic boundary conditions.

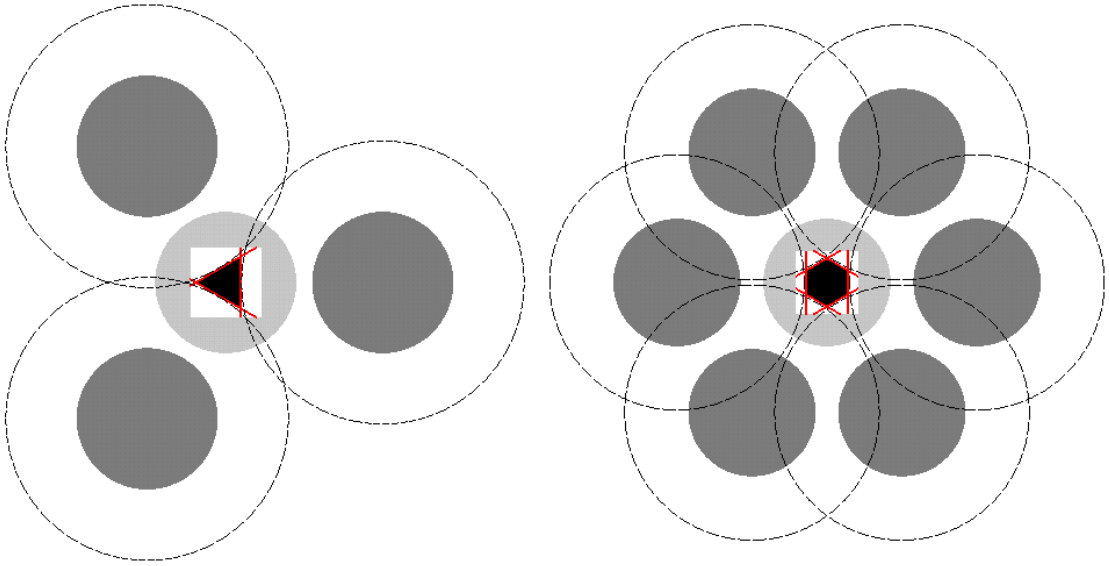

Figure 4.1: The region of allowed displacements $\mathcal{J}_{\Delta\mathbf{R}}$ for a locally jammed disk (light gray) trapped among three (left) or six (right, as in the triangular lattice) fixed disks (dark gray). The exclusion disks (dashed lines) of diameter twice the disk diameter are drawn around each of the fixed disks, delimiting $\mathcal{J}_{\Delta\mathbf{R}}$. Also shown are the linearized versions of these constraints (red), which are simply tangents to the circles at the point of closest approach, as well as the jamming polytope $\mathcal{P}_{\boldsymbol{\Delta}\mathbf{R}}$ they enclose (black). For the isostatic case on the left this polytope is a triangle (a *simplex* in two dimensions), and a hexagon for the *hyperstatic* case on the right.

Observe that these are ordered *hierarchically*, with local being implied by collective and similarly collective being implied by strict jamming. We point out that these do not exhaust all possibilities and various intricacies can arise, especially when considering infinite packings [136].

4.2.3 Unjamming Motions

Note that the mathematics literature often uses the term *rigid* or *stable packing* for what we call a jammed packing in section 4.2.2. It can be shown [73] that to assess jamming for a given sphere packing, one need only look for the existence of *analytic* continuous displacements of the particles from their current configuration[1]. An *unjamming motion* $\Delta\mathbf{R}(t) = (\Delta\mathbf{r}_1(t), \ldots, \Delta\mathbf{r}_N(t))$, where t is a time-like parameter, $t \in [0,1]$, is a continuous analytic displacement of the spheres from their current position along the path $\mathbf{R} + \Delta\mathbf{R}(t)$, starting from the current configuration, $\Delta\mathbf{R}(0) = 0$, and ending at the final configuration $\mathbf{R} + \Delta\mathbf{R}(1)$, while observing all relevant constraints along the way. This means that impenetrability and any other particular (boundary) conditions must be observed, i.e. $\mathcal{P}(\mathbf{R} + \Delta\mathbf{R}(t))$ is a *valid* packing for all $t \in [0,1]$. If such an unjamming motion does not exist, we say that the packing is *jammed*. By changing the (boundary) constraints we get different categories of jamming, such as local, collective and strict.

It can be shown (see references in Ref. [73]) that an equivalent definition[2] is to say that a packing is jammed if it is isolated in the allowed configuration space, i.e., there is no valid packing within some (possibly small) finite region around $\mathbf{R}$ that is not equivalent (congruent) to $\mathcal{P}(\mathbf{R})$. In the language of section 4.2.1, $\mathcal{J}_{\Delta\mathbf{R}} = \{\mathbf{R}\}$.

Furthermore, it is a simple yet fundamental fact that we only need to consider first derivatives $\mathbf{V} = \frac{d}{dt}\Delta\mathbf{R}(t)$, which can be thought of as velocities, and then simply move the spheres in the directions $\mathbf{V} = (\mathbf{v}_1, \ldots, \mathbf{v}_N)$ to obtain an unjamming motion $\Delta\mathbf{R}(t) = \mathbf{V}t$. Therefore, henceforth special consideration will be given to the final displacement $\Delta\mathbf{R}(1)$, so that we will most often just write $\Delta\mathbf{R} = \Delta\mathbf{R}(1)$. The formal

[1] This is the third definition (definition *c*) in section 2.1 of Ref. [73].
[2] This is the first definition (definition *a*) in section 2.1 of Ref. [73].

statement is that *a sphere packing is rigid if and only if it is infinitesimally rigid*, see Refs. [138, 142] and Section 4.3.1. Although the proofs of this statement published in the mathematics literature consider packings of equal spheres in a hard-wall container, the proof carries directly to the case of collective jamming with periodic boundary conditions (i.e., packings on a flat torus), as well as packings of unequal spheres. As discussed in Section 4.4.5, the statement is also true for strict jamming with periodic boundary conditions.

A sphere packing is not jammed if and only if one can give the spheres velocities $\mathbf{V} \neq \mathbf{0}$ such that no two contacting spheres i and j, $\|\mathbf{r}_i - \mathbf{r}_j\| = \|\mathbf{r}_{ij}\| = r_{ij} = D$, have a relative speed v_{ij} toward each other[3],

$$v_{ij} = (\mathbf{v}_i - \mathbf{v}_j)^T \mathbf{u}_{ij} \geq 0, \tag{4.1}$$

where

$$\mathbf{u}_{ij} = \frac{\mathbf{r}_i - \mathbf{r}_j}{\|\mathbf{r}_i - \mathbf{r}_j\|}$$

is the unit vector connecting the two spheres. Of course, some special and trivial cases like rigid body translations ($\mathbf{V} =$ constant) or rigid body rotations need to be excluded since they do not really change the configuration of the system. We will elaborate on this "linearized" perspective in the context of packings with interparticle gaps in section 4.3.1.

We will plot unjamming motions as "velocity" fields, and occasionally supplement such illustrations with a sequence of frames from $t = 0$ to $t = 1$ showing the unjamming process. Note that the lengths of the vectors in the velocity fields have been scaled to aid in better visualization. For the sake of clear visualization, only two-dimensional examples will be used, however, all of the techniques described here are fully applicable to three-dimensional packings as well. Interactive Virtual Reality Modeling Language (VRML) animations which are very useful in getting an intuitive feeling for unjamming mechanisms in sphere packings can be viewed on our webpage [135].

4.2.4 Jamming and Forces

We have defined jamming above using kinematic concepts and focused on the positions of the particles, i.e., on the *geometry* of the packings. It is very instructive to discuss briefly the relations between contact forces and applied loads in the context of jamming. This is crucial because of the physical importance of statical considerations in the study of granular materials and the preponderance of force-based discussions in the physics literature. Furthermore, forces play a very important role in the analysis of the configuration-based definitions given above as dual variables associated with impenetrability constraints [143], and have appeared prominently in the mathematics literature as well [73].

Consider a configuration belonging to a basin of jamming, $\mathbf{R} \in \mathcal{J}_{\Delta\mathbf{R}}$, and an applied load $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_N)$ on the particles. In the case of spheres, $\mathbf{b}_i$ is just the total force acting on sphere i (for example, due to thermal or mechanical vibrations or externally-applied fields). In the case of nonspherical particles, it would also contain the total torque acting on each particle. Assume for simplicity that this load is independent of the configuration. Under this load, the particles will displace to a new configuration of minimal energy,

$$\begin{aligned} & \max_{\Delta\mathbf{R}} \mathbf{B}^T \Delta\mathbf{R} && \text{for virtual work} \\ \text{such that} \quad & \mathbf{R} + \Delta\mathbf{R} \in \mathcal{J}_{\Delta\mathbf{R}} && \text{for impenetrability.} \end{aligned} \tag{4.2}$$

Since the packing is jammed, the program (4.2) will have a bounded solution which lies on the boundary of $\mathcal{J}_{\Delta\mathbf{R}}$, i.e., some particles will be in contact in the new configuration. The Lagrange multipliers associated with the impenetrability constraints are in fact the reaction contact forces which resist the applied load $\mathbf{B}$.

We thus see the meaning of the three jamming categories in the static context: In a locally jammed packing each particle i can support *any* load $\mathbf{b}_i$ if its neighbors are fixed. A collectively jammed packing can resist (support) *any* loading without rearrangements of the particles as long as the boundary is held fixed externally. Strictly jammed packings on the other hand can support *any* load with a compressive global (boundary) component (i.e., positive macroscopic pressure). Note however that a packing may be able to

[3] See section 2.2 of Ref. [73].

support all compressive global loads even though it is not strictly jammed, as it may be unstable due to the existence of collective unjamming mechanisms[4].

4.2.4.1 Jammed Subpackings

It should be mentioned that jammed random particle packings produced experimentally or in simulations typically contain a small population of *rattlers*, i.e., particles trapped in a cage of jammed neighbors but free to move within the cage. For present purposes we shall assume that these have been removed before considering the (possibly) jammed remainder. This idea of excluding rattlers can be further extended to rattling clusters of particles, i.e., groups of particles that can be displaced collectively even though the remainder of the packing is jammed. In fact, one can consider any packing which has a jammed subpacking (collectively or strictly as defined above, with *identical* boundary conditions) to be jammed. This subpacking is the force-carrying *backbone* of the packing.

The physical meaning and mathematical basis for such a modified approach is more evident from the static perspective. Specifically, as long as there is a jammed subpacking, this subpacking will resist (support) global loads (stresses), and furthermore, this jammed subpacking is also able to resist local loads, such as, for example, induced by vibrations (shaking) in granular materials, therefore making the whole packing stable and rigid.

4.3 Linear Programming Algorithm to Test for Jamming

Given a sphere packing, we would often like to test whether it is jammed according to each of the categories given above, and if it is not, find one or several unjamming motions $\Delta\mathbf{R}(t)$. We now describe a simple algorithm to do this that is exact for *gap-less (ideal) packings*, i.e., packings where neighboring spheres touch exactly, and for which the definitions given earlier apply directly. However, in practice, we would also like to be able to study *packings with small gaps*, such as produced by various heuristic compression schemes like the Lubachevsky-Stillinger algorithm [12], and we will consider these along with ideal packings. In this case the meaning of unjamming needs to be modified so as to fit physical intuition. We do this using what Roux [143] calls the *approximation of small displacements* (ASD), and propose an algorithm based on linear programming that can test whether a finite packing is jammed.

We believe that computer-generated packings which are almost ideal are often actually very close in configurational space to an ideal packing. This cannot be verified exactly in most cases, though some support for this claim can be obtained by setting the numerical precision of the generation algorithm to higher and higher values (for example, by increasing the number of collisions per particle in the Lubachevsky-Stillinger algorithm [12]) and verifying that the interparticle gaps monotonically decrease toward zero [75]. It is possible though that gaps are natural and *essential* in certain applications, such as for example, the study of particle rearrangement in granular materials, and we therefore separately study packings which need not be (close to) ideal, using mathematical programming as the fundamental tool. When the configuration is known exactly, often the case for small ordered packings, jamming may be analyzed analytically.

4.3.1 Approximation of Small Displacements

As already explained, an unjamming motion for a sphere packing can be obtained by giving the spheres suitable velocities, such that neighboring spheres do not approach each other. Here we focus on the case when $\Delta\mathbf{R}(t) = \mathbf{V}t + \mathcal{O}(t^2)$ are small finite displacements from the current configuration. We will drop the time designation and just use $\Delta\mathbf{R}$ for the displacements from the current configuration $\mathbf{R}$ to the new configuration $\widetilde{\mathbf{R}} = \mathbf{R} + \Delta\mathbf{R}$. We defer discussion of packings with significant interparticle gaps to section 4.3.3.

In this ASD approximation, we can linearize the impenetrability constraints

$$\|\widetilde{\mathbf{r}}_i - \widetilde{\mathbf{r}}_j\|^2 = \|(\mathbf{r}_i - \mathbf{r}_j) + (\Delta\mathbf{r}_i - \Delta\mathbf{r}_j)\|^2 \geq D^2 \tag{4.3}$$

[4] An example is the Kagomé lattice disk packing, which can support all compressive global loads (sometimes called "loads at infinity" in the engineering literature), but is not collectively jammed with periodic boundary conditions.

by expanding to first order in $\Delta\mathbf{R}$, to get the condition for the existence of a (first-order) *feasible displacement* $\Delta\mathbf{R}$,

$$\Delta\mathbf{r}_{ij}^T\mathbf{u}_{ij} \geq -h_{ij} \text{ for all } \{i,j\}, \tag{4.4}$$

where $\{i,j\}$ represents a *potential contact* between nearby spheres i and j, and $h_{ij} = r_{ij} - D$ is the *interparticle gap* (or *interstice*). The set of contacts $\{i,j\}$ that we include in (4.4) form the *contact network* of the packing, and they correspond to a subclass of the class of fascinating objects called *tensegrity frameworks*, namely *strut frameworks* (see Ref. [73] for details, and also [144] for a treatment of more general packings). We only consider potential contacts $\{i,j\}$ between nearby, and not all pairs of spheres, that is we only consider a contact if

$$r_{ij} \leq (1+\delta)\, D, \tag{4.5}$$

where $\delta \ll 1$ is a chosen *gap tolerance*. Figure 4.8 shows a random bidisperse (containing two different kinds of particles) disk packing and the associated contact network.

For a gap-less packing, we have $\Delta l = 0$ and the condition (4.4) reduces to (4.1), and the packing is jammed if and only if the only nontrivial solution to (4.4) is $\Delta\mathbf{R} = 0$. For packings with finite but small gaps though, condition (4.4) is only a first-order approximation. By transforming Eq. (4.3) we obtain the nonlinear analog of Eq. (4.4):

$$\Delta\mathbf{r}_{ij}^T\mathbf{u}_{ij} + \frac{\Delta\mathbf{r}_{ij}^T\Delta\mathbf{r}_{ij}}{2r_{ij}} \geq -\frac{r_{ij}^2 - D^2}{2r_{ij}} = -h_{ij} + \frac{h_{ij}^2}{2r_{ij}} \text{ for all } \{i,j\}. \tag{4.6}$$

We rewrote the impenetrability constraints in this form so that it can be seen more clearly that any displacement feasible under the linearized constraints (4.4) will also be feasible under the full nonlinear impenetrability constraints (4.6). To see this, consider how far one can move the particles along a unit $\widehat{\Delta\mathbf{r}_{ij}}$ before being blocked by the linearized constraint, $\Delta\mathbf{r}_{ij} = -\Delta r_{ij}^{(L)}\widehat{\Delta\mathbf{r}_{ij}}$,

$$\Delta r_{ij}^{(L)} = -\frac{h_{ij}}{\widehat{\Delta\mathbf{r}_{ij}}^T\mathbf{u}_{ij}} = \frac{h_{ij}}{\cos\theta_{ij}} \geq h_{ij}.$$

In reality however, the particles can displace more, since the true nonlinear constraint only blocks the motion after a displacement

$$\Delta r_{ij}^{(NL)} = r_{ij}\cos\theta_{ij} - \sqrt{(r_{ij}\cos\theta_{ij})^2 + h_{ij}^2 - 2h_{ij}r_{ij}} \geq \Delta r_{ij}^{(L)}.$$

In other words, within the ASD, $\mathcal{J}_{\Delta\mathbf{R}}$ is approximated with the *inscribed* polyhedral set $\mathcal{P}_{\Delta\mathbf{R}} \subseteq \mathcal{J}_{\Delta\mathbf{R}}$ of feasible (linearized) displacements, as determined by the system of linear inequalities (4.4), as illustrated in Fig. 4.2. This is a very special and most useful property of sphere packings which does not generalize to other convex particle shapes. The formal statement is that a sphere packing is rigid if and only if it is infinitesimally rigid [138]. Note also that to first order in h_{ij}, $\Delta r_{ij}^{(NL)} \approx \Delta r_{ij}^{(L)}$, a fact that was used in Ref. [145] to conclude that the volume of $\mathcal{J}_{\Delta\mathbf{R}}$ asymptotically tends to the volume of $\mathcal{P}_{\Delta\mathbf{R}}$ in the jamming limit (we will use this in subsequent chapters).

Also note that for any nontrivial (i.e., $\Delta\mathbf{r}_i \neq \Delta\mathbf{r}_j$) solution to (4.4) the nonlinear inequality (4.6) is satisfied as a *strict* inequality, which means that particles i and j lose contact, even if the inequality is active to first order. This is an important property which shows that by scaling a first-order feasible displacement $\Delta\mathbf{R}$ appropriately one can always obtain a nontrivial feasible displacement which separates some of the contacting particles. This property does not directly generalize to other smooth strictly convex particle shapes, and in particular, it does not apply to packings of ellipsoids. Comparison of (4.4) and (4.6) also suggests that the linearized constraints become too strict as the magnitude of the displacements becomes comparable to the size of the particles $\|\Delta\mathbf{r}_{ij}\| \approx D$. The complicated issue of how well the ASD approximation works when the gaps are not small enough is illustrated in Fig. 4.2.

By putting the $\mathbf{u}_{ij}$'s as columns in a matrix of dimension $[Nd \times M]$, we get the important *rigidity matrix*[5] of the packing $\mathbf{A}$. This matrix is sparse and has two blocks of d non-zero entries in the column corresponding

[5]This is in fact the normalized negative transpose of what is usually taken to be the rigidity matrix, and is chosen to fit the notation in Ref. [146], and also because it resembles the *node-arc incidence matrix* of the (directed) graph corresponding to the contact network.

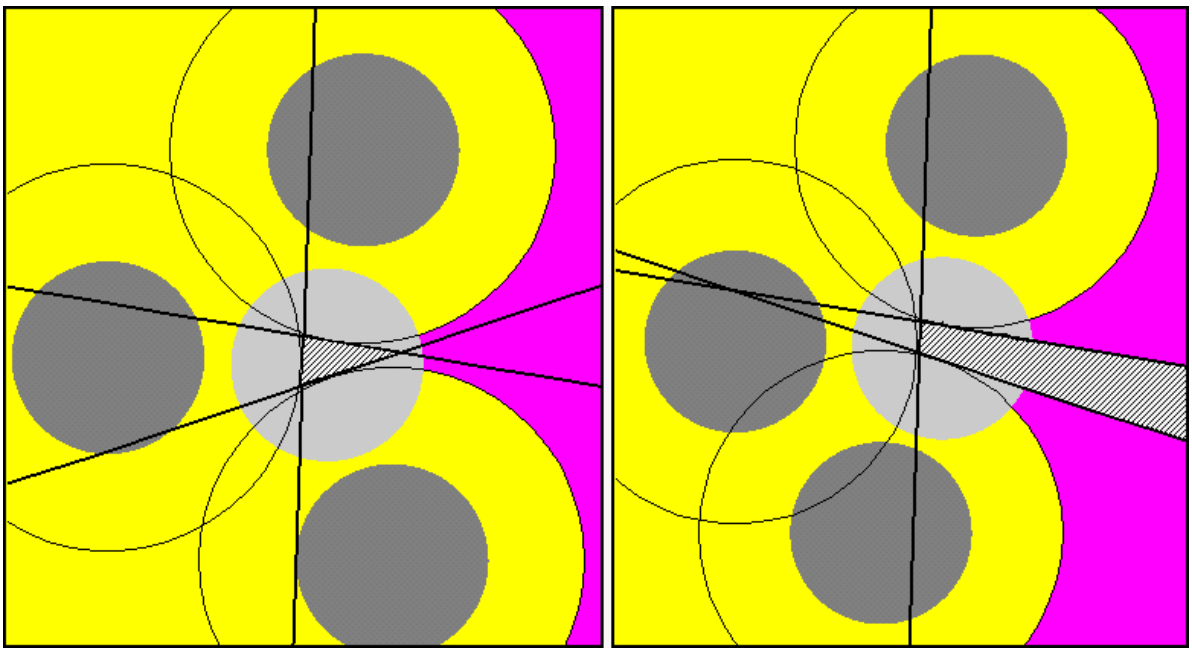

Figure 4.2: *Feasible displacements polyhedron.* The figures show three stationary (dark gray) disks surrounding a mobile disk (light gray), as in the left part of Fig. 4.1. The (nonconvex) region of configuration space excluded to the mobile disk $\mathcal{J}_{\mathbf{\Delta R}}$ (purple) is the complement of the union of disks of radius D surrounding each stationary disk (dark circles), and is surrounded by unreachable regions of configuration space (yellow). Also shown are the linearized versions of these constraints (dark lines), as well as the region $\mathcal{P}_{\mathbf{\Delta R}}$ (shaded gray) they bound.
This region is a polyhedral set, and in the left figure it is bounded, meaning that within the approximation of small displacements (ASD) the mobile disk is locally jammed (trapped) by its three neighbors, while on the right it is unbounded, showing the cone of locally unjamming motions (escape routes). Notice that with the true nonlinear constraints, the mobile disk can escape the cage of neighbors in both cases, showing that the ASD is not exact. However, it should also be clear that this is because we have relatively large interparticle gaps here and a near flat angle between the two nearly vertical contacts (contrast this to Fig. 4.1).

to the particle contact $\{i, j\}$, namely, $\mathbf{u}_{ij}$ in the block row corresponding to particle i and $-\mathbf{u}_{ij}$ in the block row corresponding to particle j. Represented schematically:

$$\mathbf{A} \quad = \quad \begin{array}{l} \\ i \rightarrow \\ \\ j \rightarrow \\ \\ \end{array} \begin{array}{c} \{i,j\} \\ \downarrow \\ \begin{bmatrix} \vdots \\ \mathbf{u}_{ij} \\ \vdots \\ -\mathbf{u}_{ij} \\ \vdots \end{bmatrix} \end{array}$$

For example, for the four-disk packing shown in Fig. 4.2, and with the numbering of the disks depicted in

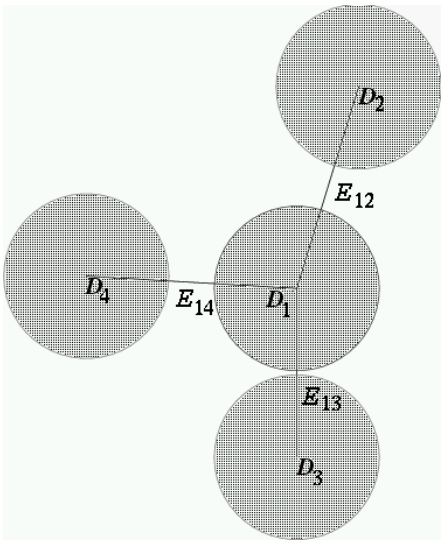

Figure 4.3: The packing from Fig. 4.2 shown again with a numbering of the disks. D_i denotes particle i and E_{ij} denotes the contact between the i-th and j-th particles, i.e., the contact $\{i, j\}$.

Fig. 4.3, we have the following rigidity matrix:

$$\mathbf{A} = \begin{array}{c} \\ D_1 \\ D_2 \\ D_3 \\ D_4 \end{array} \begin{array}{c} \begin{array}{ccc} E_{12} & E_{13} & E_{14} \end{array} \\ \begin{bmatrix} \mathbf{u}_{12} & \mathbf{u}_{13} & \mathbf{u}_{14} \\ -\mathbf{u}_{12} & & \\ & -\mathbf{u}_{13} & \\ & & -\mathbf{u}_{14} \end{bmatrix} \end{array}.$$

Using this matrix, we can rewrite the linearized impenetrability constraints as a simple system of linear inequality constraints

$$\mathbf{A}^T \Delta\mathbf{R} \geq -\mathbf{h}. \tag{4.7}$$

4.3.1.1 The Jamming Polytope $\mathcal{P}_{\Delta\mathbf{R}}$

The polyhedral set $\mathcal{P}_{\Delta\mathbf{R}}$ is necessarily bounded for a jammed configuration, i.e., it is a convex polytope (see the more detailed discussion in Section 4.3.2.3), which we refer to as the *jamming polytope*. The jamming polytope will take a very prominent place in the theories developed in this work.

The fact that $\mathcal{P}_{\Delta\mathbf{R}}$ is a bounded convex polytope implies that $\mathbf{A}$ is of full rank [71], and that the number of faces bounding $\mathcal{P}_{\Delta\mathbf{R}}$, i.e., the number of interparticle contacts M, is at least one larger than the dimensionality N_f of the configuration space[6], $M \geq N_f + 1$. For collective jamming the boundary conditions are fixed and with periodic boundary conditions there are d trivial translational degrees of freedom, so $N_f = (N-1)d$. If hard wall boundary conditions are employed then $N_f = Nd$ and one should also count contacts with the hard walls among the M constraints. For strict jamming the boundary is also allowed to deform and this introduces additional degrees of freedom. For example, with periodic boundary conditions a symmetric non-expansive macroscopic strain tensor is added to the configuration parameters, as discussed in detail in Section 4.4, giving $N_f = (N-1)d + d(d-1)/2 + (d-1)$ degrees of freedom. *Isostatic* packings are jammed packings which have the minimal number of contacts, namely, for collective jamming

$$M = \begin{cases} 2N-1 \text{ for } d=2 \\ 3N-2 \text{ for } d=3 \end{cases} \tag{4.8}$$

[6]The additional +1 comes because we are considering inequality constraints, rather than equalities. One can also think of this extra degree of freedom as representing the density, i.e., the size of the particles. For example, looking at the left panel of Fig. 4.1 we see that at least 3 linear inequalities are necessary to bound a polytope in 2 dimensions.

and for strict jamming

$$M = \begin{cases} 2N+1 \text{ for } d=2 \\ 3N+3 \text{ for } d=3 \end{cases}, \tag{4.9}$$

with periodic boundary conditions. For isostatic packings, the jamming polytope is a *simplex*[7], and this will be used frequently in theoretical calculations since simplices are much simpler to manipulate analytically than general polytopes. Packings having more contacts than necessary are *hyperstatic*, and packings having less contacts are *hypostatic* (we will see that hypostatic sphere packings cannot be jammed). For the trivial example of local jamming and $N = 1$, all particles but one are frozen in place and the free particle must have at least $d + 1$ contacts. Figure 4.1 shows the polytope $\mathcal{P}_{\mathbf{\Delta R}}$ for a locally jammed disk, for both an isostatic and a hyperstatic case.

4.3.2 Randomized Linear Programming (LP) Algorithm

The question of whether an ideal packing is jammed, i.e., whether the system (4.7) is feasible for some $\mathbf{\Delta R} \neq 0$, can be answered rigorously by using standard linear programming (LP) techniques, as described in Section 4.3.2.1. If a packing is jammed, then this LP test is enough. However, for packings that are not jammed, it is more useful to obtain a *representative* collection of unjamming motions, rather then use a binary classification into packings which are jammed and ones which are not jammed. A random collection of such unjamming motions is most interesting, and can be obtained easily by solving several linear programs with a random cost vector.

We adapt such a *randomized LP algorithm* to testing for jamming, namely, we solve the following LP in the *displacement formulation*

$$\begin{array}{rll} & \max_{\mathbf{\Delta R}} \mathbf{B}^T \mathbf{\Delta R} & \text{for virtual work} \\ \text{such that} & \mathbf{A}^T \mathbf{\Delta R} \geq -\mathbf{h} & \text{for impenetrability} \\ & |\mathbf{\Delta R}| \leq \Delta R_{\max} & \text{for boundedness,} \end{array} \tag{4.10}$$

for a *random load* $\mathbf{B}$, where $\Delta R_{\max} \gg D$ is used to prevent unbounded solutions and thus improve numerical behavior[8]. The physical interpretation of $\mathbf{B}$ as an external load was elucidated in section 4.2.4. Trivial solutions, such as uniform translations of the packing $\mathbf{\Delta R} = \text{const.}$ for periodic boundary conditions, can be eliminated *a posteriori*, for example by reducing $\mathbf{\Delta R}$ to zero mean displacement. Alternatively, trivial motions can be handled by introducing extra constraints in (4.10), for example, by fixing the position of one of the spheres, though we have found this less attractive, particularly for packings with gaps. Finally, trivial components of $\mathbf{\Delta R}$ can also be avoided by carefully choosing $\mathbf{B}$ to be in the null-space of $\mathbf{A}$, which usually means it needs to have zero total sum and total torque (see chapter 15 in Ref. [146]). We will discuss numerical techniques to solve (4.10) in Section 4.6.

The reason we have included possibly non-zero gaps $\mathbf{h}$ in (4.10) is that computer-generated packings, which we analyze, are never ideal and there are always small interparticle gaps between some particles[9], typically much less than a percent of the typical particle size D. One can safely consider such packings purely within the ASD. However, we need to modify our definition of jamming to allow for very small particle rearrangements at the application of the load $\mathbf{B}$, i.e., we consider a solution to (4.10) an unjamming motion only if some particle is displaced a significant distance:

$$\exists i \text{ such that } \|\mathbf{\Delta r}_i\| \geq \Delta r_{\text{large}} \gg \overline{h},$$

where $\overline{h} \ll D$ is the typical size of the interparticle gap. Even though any solution to $\mathbf{A}^T \mathbf{\Delta R} \geq 0$ is also a solution to $\mathbf{A}^T \mathbf{\Delta R} \geq -\mathbf{h}$, the latter may have other solutions with large components, corresponding to elongated corners of the polyhedron of feasible displacements $\mathcal{P}_{\mathbf{\Delta R}}$ (see Fig. 4.2), which should also be treated as unjamming motions. Therefore, the primary purpose of including the exact interparticle gaps in

[7] A simplex is a closed convex polytope that has $\nu + 1$ faces and $\nu + 1$ vertices in ν-dimensional space (i.e., a triangle in two or a tetrahedron in three dimensions).

[8] In our tests we usually set $\Delta R_{\max} \sim 100D$.

[9] These gaps may be an inherent and essential feature of disordered packings in general.

(4.10) is to ensure proper handling of degenerate cases, such as a near-180-degree angle between two contacts in 2D (see Fig. 4.3).

In summary, we treat any solution $\Delta\mathbf{R}$ to (4.10) with components significantly larger than $\overline{h}$ as an unjamming motion. For each $\mathbf{B}$, if we fail to find an unjamming motion, we apply $-\mathbf{B}$ as a loading also, for reasons we detail in Section 4.3.2.3. We stress that despite its randomized character, this algorithm is almost rigorous when used as a test of jamming, in the sense that it is *strictly rigorous* for gap-less packings, and also likely to work well if the interparticle gaps are sufficiently small, as explained in more detail in Section 4.3.2.3. We will discuss more complicated adaptations of the randomized LP algorithm to nonideal, i.e., packings with larger gaps, in section 4.3.3.

4.3.2.1 Non-Randomized LP Testing for Jamming

Determining whether the linear system of inequalities (4.7) with $\mathbf{h} = 0$ has nontrivial solutions is an interesting mathematical programming problem. One approach is the following: Solve the following linear program aimed at maximizing the sum of the (positive) gap dilations,

$$\begin{aligned} & \max_{\Delta\mathbf{R}} \left[\sum_{\{i,j\}} (\mathbf{A}^T\Delta\mathbf{R})_{ij} = (\mathbf{A}\mathbf{e})^T \Delta\mathbf{R}\right] \\ \text{such that} \quad & \mathbf{A}^T\Delta\mathbf{R} \geq \mathbf{0}, \end{aligned} \tag{4.11}$$

and if this returns $\Delta\mathbf{R} = \mathbf{0}$ as one of the optimal solutions, test the rigidity matrix $\mathbf{A}$ for rank-deficiency, i.e., look for nontrivial solutions of $\mathbf{A}^T\Delta\mathbf{R} = \mathbf{0}$. If this also fails to find an unjamming motion, the packing is jammed. Notice that this will usually produce a single unjamming motion.

4.3.2.2 Kinematic/Static Duality

The subject of kinematic/static duality and its physical meaning and implications have been discussed in numerous previous works [143, 73, 138, 137, 15]. The dual of the displacement formulation LP (4.10) (excluding the additional practical safeguard constraint $\Delta\mathbf{R} \leq \Delta\mathbf{R}_{\max}$), is the *force formulation* LP

$$\begin{aligned} & \min_{\mathbf{f}} \mathbf{h}^T\mathbf{f} && \text{for virtual work} \\ \text{such that} \quad & \mathbf{A}\mathbf{f} = -\mathbf{B} && \text{for equilibrium} \\ & \mathbf{f} \geq \mathbf{0} && \text{for repulsion only,} \end{aligned} \tag{4.12}$$

and gives the *interparticle repulsive*[10] *force* f_{ij} between spheres i and j as the dual variable associated with the impenetrability constraint (4.4). The displacement- and force-based LP's are of great importance in studying the *stress-strain* behavior of granular materials, and since they are equivalent to each other, we can call them the *ASD* stress-strain LP. We have emphasized the displacement formulation (4.10) simply because we based our discussion of jamming on a kinematic perspective, but a parallel static interpretation can easily be given. For example, a random $\mathbf{B}$ used in the randomized LP algorithm that finds an unbounded unjamming motion physically corresponds to a load that the packing cannot support, i.e., the force formulation (dual) LP is infeasible, implying that the displacement formulation (primal) LP is unbounded.

In general the stress-strain LP will be highly degenerate and its primal and/or dual solution not unique. However, as Roux points out [143], the existence of small gaps in random packings is very important in this context. Namely, if $\mathbf{h}$ is random and nonzero (even if small), and $\mathbf{B}$ is also random, both the primal and dual solutions will likely be non-degenerate (see the references in Ref. [146]), and we have indeed observed this in practice for random packings. A non-degenerate (basic or vertex) solution to (4.12) corresponds to an *isostatic* force-carrying contact network [137, 143].

4.3.2.3 The Geometry of the Set of Unjamming Motions

The linearized impenetrability constraints $\mathbf{A}^T\Delta\mathbf{R} \geq -\mathbf{h}$ define a polyhedral set $\mathcal{P}_{\Delta\mathbf{R}}$ of *feasible (linearized) displacements*. Every such convex polyhedron consists of a finite piece $\mathcal{P}^{\text{hull}}_{\Delta\mathbf{R}}$, a convex polytope given by the convex hull of its extreme points, and possibly an unbounded piece $\mathcal{C}_{\Delta\mathbf{R}}$, a finitely generated *polyhedral*

[10] Note that a negative sign has been used for repulsive forces in the mathematical literature [73].

cone. In some cases this cone will be empty (i.e., $\mathcal{C}_{\Delta\mathbf{R}} = \{\mathbf{0}\}$), but in others it will not, as can be seen in Fig. 4.2. The full nonlinear impenetrability constraints given by (4.6) define the true set of feasible displacements $\mathcal{P}^{\mathrm{NL}}_{\Delta\mathbf{R}} = \mathcal{J}_{\Delta\mathbf{R}} - \{\mathbf{R}\}$, which always relaxes the linearization, $\mathcal{P}_{\Delta\mathbf{R}} \subseteq \mathcal{P}^{\mathrm{NL}}_{\Delta\mathbf{R}}$. A mathematically well defined definition of jamming is to take any ray in the cone $\mathcal{C}_{\Delta\mathbf{R}}$ as an unjamming motion, and exclude others, however, as Fig. 4.2 shows, the elongated corners of $\mathcal{P}_{\Delta\mathbf{R}}$ are in fact very likely to be unbounded in the true non-linear feasible set of displacements $\mathcal{P}^{\mathrm{NL}}_{\Delta\mathbf{R}}$, so we prefer to take any "long" direction in $\mathcal{P}_{\Delta\mathbf{R}}$ as an unjamming motion.

We note that the randomized LP algorithm proposed here strictly answers the question of whether the polyhedral set of feasible displacements contains an unbounded ray (i.e., whether $\mathcal{C}_{\Delta\mathbf{R}} \neq \{\mathbf{0}\}$) just by applying two (nonzero) loads $\mathbf{b}$ and $-\mathbf{b}$. This is because an attempt to find such a ray will be unsuccessful only if $-\mathbf{b} \in \mathcal{C}^*_{\Delta\mathbf{R}}$, where $\mathcal{C}^*_{\Delta\mathbf{R}}$ is the dual (conjugate) cone of $\mathcal{C}_{\Delta\mathbf{R}}$, and in this case $\mathbf{b} \notin \mathcal{C}^*_{\Delta\mathbf{R}}$, so that using the load $-\mathbf{b}$ will find a ray if such a ray exists. We note that one cannot hope to fully characterize the cone of first-order unjamming motions $\mathcal{C}_{\Delta\mathbf{R}}$ (i.e. find its convex hull of generating rays), as this is related to the hard problem of fully enumerating the vertices of a polyhedron. Our randomized approach essentially finds a few sample rays in $\mathcal{C}_{\Delta\mathbf{R}}$.

It would be interesting to consider the behavior of the randomized LP quantitatively for the case when $\mathcal{P}_{\Delta\mathbf{R}}$ is a closed polytope, and determine the probability that a certain vertex of the polytope is found by the LP for a random load $\mathbf{B}$.

4.3.3 Dealing with Interparticle Gaps

We originally motivated our perspective on jamming in section 4.2.1 by looking at the set of available (reachable) configurations $\mathcal{J}_{\Delta\mathbf{R}}$ around a particular initial configuration $\mathbf{R}$, and have since focused mostly on ideal packings, though allowing for sufficiently small interparticle gaps. For these packings, $\mathcal{J}_{\Delta\mathbf{R}}$ is very localized around $\mathbf{R}$, and this makes it possible to define the three jamming categories meaningfully and rigorously, and also allows for a simple randomized linear programming testing algorithm. A somewhat ambitious but desirable goal is to efficiently obtain a grasp on the character and extent of $\mathcal{J}_{\Delta\mathbf{R}}$, and use this to judge whether the packing should be considered jammed or not. However, since $\mathcal{J}_{\Delta\mathbf{R}}$ is a very high-dimensional and nonconvex set, it is a very complex object to describe or understand. We discuss the full geometry of $\mathcal{J}_{\Delta\mathbf{R}}$ in more detail in Section 4.5.

However, the either-or character of the jamming criteria for ideal packings is often too restrictive or specialized when analyzing large disordered packings with possibly larger interparticle gaps, where particle displacements may be comparable to the typical particle size. Furthermore, even very close to the jamming point, unjamming motions exist for very large packings. This is easy to appreciate by noting that one can push one half of the particles to the left, and the other to the right, eliminating the small (but non-zero) interparticle gaps, and forming a large void in the packing where arbitrary particle rearrangements can occur. In other words, for sufficiently large packings large density fluctuations (formation of large voids) is always possible (though unlikely) at densities below the jamming point.

We could focus instead on trying to judge the *extent* of $\mathcal{J}_{\Delta\mathbf{R}}$ by trying to displace the spheres away from their current position by as much as possible. This can be done with a *sequential random loading* algorithm: Repeatedly solve the LP (4.10), displace the spheres in the direction of $\Delta\mathbf{R}$ by as much as possible while still avoiding overlap, until the particles rearrange and form contacts that actually support the applied load $\mathbf{B}$. This should be repeated for several random loads, in the hope of exploring $\mathcal{J}_{\Delta\mathbf{R}}$ along several directions. We give an outline of an algorithm to do this along with representative results in Section 4.6.2. The important point here is that for packings which are almost jammed, mathematical programming is needed in order to efficiently find a direction in which the particles can be displaced by significant amounts. Traditional heuristics such as Monte Carlo schemes in which particles are displaced one-by-one simply get trapped easily, and algorithms which search for collective particle rearrangements are needed.

4.3.3.1 Shrink-And-Bump Heuristic

The following heuristic test for collective jamming has been suggested in Ref. [12]: Shrink the particles by a small amount δ and then start the Lubachevsky-Stillinger molecular dynamics algorithm with random velocities, and see if the system gets unjammed. One would also slowly enlarge the particles back to their original

size while they bump around, so as to allow finite termination of this test (within numerical accuracies). We call this the *shrink-and-bump heuristic.* The idea is that the vector of velocities takes on random values in velocity space and if there is a direction of unjamming, it will be found with a high probability and the system will unjam. Animations of this process can be found at Ref. [135], and applications to testing for jamming in very large packings of hard spheres are described in Chapter 9.

This kind of heuristic has the advantage of being very simple and thus easy to implement and use (and also incorporates nonlinear effects), and it is also efficient, though still significantly slower than the linear programming algorithm since typically many collisions per particle are needed to significantly displace the particles due to the high density[11]. By incorporating deformations of the lattice in the Lubachevsky-Stillinger algorithm, one can also use this test to test for strict jamming [72]. Its disadvantages are its non-rigorous character and indeterminacy, artificial introduction of dynamics into a geometrical problem, and most of all, its strong dependence on the exact value of δ. For example, animations showing how the Kagomé lattice inside a container made of fixed spheres (as in Fig. 4.5) can be unjammed with a large-enough δ, even though it is actually collectively jammed under these boundary conditions, can be found at our webpage [135]. In fact, many jammed *large* packings will appear unstable under this kind of test, as motivated with the notion of *uniform stability*, defined in Ref. [136]. The reader should also note the discussion in Section 4.5.

4.4 Boundary Conditions

Boundary conditions are very important when considering jamming, since some kind of container is necessary to make a finite packing jammed. In fact, we can consider the different jamming categories to be variations in the boundary conditions applied when looking for unjamming motions. First, the configuration, i.e., the positions of all the particles, and the shape of the container, need to be specified. Then the allowed motions of the particles and the container need to be prescribed. We will usually work with periodic boundary conditions (in order to minimize finite-size effects), which we formally describe next, although sometimes we use hard-wall containers (in order to compare to experiments).

4.4.1 Repetitive (Periodic) Packings

Large or infinite packings are most easily created by periodically repeating a certain finite (and possibly small) known packing. A *repetitive packing* $\widehat{\mathcal{P}}(\mathbf{R})$ is generated by replicating a finite *generating packing* $\mathcal{P}(\widehat{\mathbf{R}})$ on a lattice $\mathbf{\Lambda} = \{\boldsymbol{\lambda}_1, \ldots, \boldsymbol{\lambda}_d\}$, where $\boldsymbol{\lambda}_i$ are linearly independent *lattice vectors* and d is the spatial dimensionality. The positions of the spheres are generated by

$$\mathbf{r}_{\widehat{i}(\mathbf{n}_c)} = \widehat{\mathbf{r}}_i + \mathbf{\Lambda}\mathbf{n}_c \text{ and } \mathbf{n}_c \text{ is integer, } \mathbf{n}_c \in Z^d, \tag{4.13}$$

where we think of the *lattice* $\mathbf{\Lambda}$ as a matrix with d^2 elements having the lattice vectors as columns and $\mathbf{n}_c$ is the number of replications of the unit cell along each basis direction. Note that in subsequent chapters the letter $\mathbf{n}$ will be used heavily to denote normal vectors at the point of contact, not to be confused with number of unit cells as used in this Chapter. The sphere $\widehat{i}\,(\mathbf{n}_c)$ is the familiar *image sphere* of the *original sphere* $i \equiv \widehat{i}\,(0)$, and of course for the impenetrability condition only the nearest image matters. Notice that condition (4.13) only gives the positions of the spheres, and additional boundary conditions need to be specified before applying the jamming definitions from Section 4.2.2.

4.4.2 Boundary Conditions for Unjamming Motions

As previously mentioned, the boundary conditions imposed on a given packing are very important, especially in the case of strict jamming. Here we consider two main types of boundary conditions, hard-wall and periodic boundary conditions, focusing on finite packings. Infinite packings are discussed in Ref. [136].

Hard-wall boundaries The packing $\mathcal{P}(\mathbf{R})$ is placed in an impenetrable concave *hard-wall container* K (see Ref. [138]). Figure 4.4 shows that the honeycomb lattice can be unjammed inside a certain hard-wall

[11] In Chapter 9 we monitor the pressure in order to more efficiently detect unjamming motions, visible through the decay in the pressure.

container. We can also make an effective container out of a certain number of *fixed spheres* whose positions cannot change. This is because it is often hard to fit a packing into a simple container such as a square box, while it is easy to surround it with other fixed spheres, particularly if a periodic lattice is used to generate the packing. Specifically, one can take a finite sub-packing of an infinite repetitive packing and freeze the rest of the spheres, thus effectively making a container for the sub-packing. An example is depicted in Fig. 4.5. Note that hard-wall containers do not allow any trivial unjamming motions.

Periodic boundaries Periodic boundary conditions are often used to emulate infinite systems, and they fit the algorithmic framework of this work very nicely. To obtain a *periodic packing* we wrap a repetitive packing $\widehat{\mathcal{P}}(\mathbf{R})$ around a flat torus, i.e. we ask that whatever happens to a sphere i also happens to all of the image spheres $\widehat{i}\,(\mathbf{n}_c)$, with the additional provision that the lattice may also change by $\Delta\mathbf{\Lambda}$,

$$\Delta\mathbf{r}_{\widehat{i}(\mathbf{n}_c)} = \Delta\widehat{\mathbf{r}}_i + (\Delta\mathbf{\Lambda})\mathbf{n}_c. \tag{4.14}$$

When the lattice is fixed ($\Delta\mathbf{\Lambda} = \mathbf{0}$), periodic boundary conditions allow for trivial rigid body translations of the packing, but trivial rotations only exist if the lattice is allowed to change. The trivial translations can be most easily eliminated by fixing (freezing) the position of one of the particles. If the lattice also changes, then we consider the configuration of the packing $\mathbf{Q} = (\mathbf{R}, \mathbf{\Lambda})$ to also include the lattice vectors. By imposing a suitable condition on the deformation of the lattice $\Delta\mathbf{\Lambda}$, as described in Section 4.4.5, one can eliminate the trivial rigid-body rotations of the packing.

4.4.2.1 Using Simple Lattices to Generate Packings

Familiar lattices with a simple basis (unit cell), such as the triangular, honeycomb, Kagomé and square in two dimensions, or the simple cubic (SC), body-centered cubic (BCC), face-centered cubic (FCC) and hexagonal-close packed (HCP) in three dimensions, can be used to create a (possibly large) packing taking a subsystem of size $\mathbf{N}_c$ unit cells along each dimension from the infinite lattice packing. The properties of the resulting system can be studied with the tools developed here, provided that we restrict ourselves to finite $\mathbf{N}_c$. Moreover, it is important to specify which lattice vectors are to be used. We will usually take them to be primitive vectors (for which there is one particle per unit cell), but sometimes it will be more convenient to use conventional ones, as used in the physics literature (usually representing a cubic unit cell having more then one particle per unit cell for variations on the cubic lattice).

For hard-wall boundary conditions, we can take an infinite packing generated by these simple lattices and then freeze all but the spheres inside the window of $\mathbf{N}_c$ unit cells, thus effectively obtaining a hard-wall container. Figure 4.5 illustrates an unjamming motion for the honeycomb lattice under these conditions.

For periodic boundary conditions, the generator $\mathcal{P}(\widehat{\mathbf{R}})$ can itself be generated using $\mathbf{N}_c$ unit cells of a simple lattice[12]. In this case the lattice $\mathbf{\Lambda}$ is a *sub-lattice* of the underlying (primitive) lattice $\widetilde{\mathbf{\Lambda}}$, i.e., $\mathbf{\Lambda} = \widetilde{\mathbf{\Lambda}}\text{Diag}\{\mathbf{N}_c\}$, where $\text{Diag}\{\mathbf{N}_c\}$ denotes a diagonal matrix whose diagonal is $\mathbf{N}_c$. This is not only a convenient way to generate simple finite periodic packings, but it is in general what we mean when we ask, for example, to analyze the jamming properties of the Kagomé lattice under periodic or hard-wall boundary conditions. Figure 4.6 shows a periodic unjamming motion for the Kagomé lattice. Notice though that the jamming properties one finds depend on how many neighboring unit cells $\mathbf{N}_c$ are used as the "base" region (i.e., the generating packing), and therefore, we will usually specify this number explicitly. Some properties may be independent of $\mathbf{N}_c$ (for example, the triangular lattice packing is strictly jammed for all $\mathbf{N}_c$) and tailored mathematical analysis can be used to show this [142, 136]. More systematic approaches based on Bloch wave (Fourier) decompositions of the set of feasible motions should be investigated for repetitive packings[13]. We will not consider these issues in detail here, but rather focus on algorithmic approaches tailored for *finite* and *fixed* systems (i.e., $\mathbf{N}_c$ is fixed and finite), which is important when studying disordered particle packings, i.e., packings where the generator $\mathcal{P}(\widehat{\mathbf{R}})$ is itself a large disordered packing.

[12] This closely resembles the Born-von Karman boundary conditions used in solid-state physics models of lattice vibrations.

[13] The difficulty with using standard Bloch wave (Fourier) decompositions is that here we are dealing with inequalities, rather than equalities, which cannot easily be decomposed into inequalities for each of the wavevectors using standard orthogonality relations. However, as will see in Chapter 5, both the *self stresses* and the *floppy modes* in a packing are solutions to systems of equalities and they can be decomposed into Fourier components.

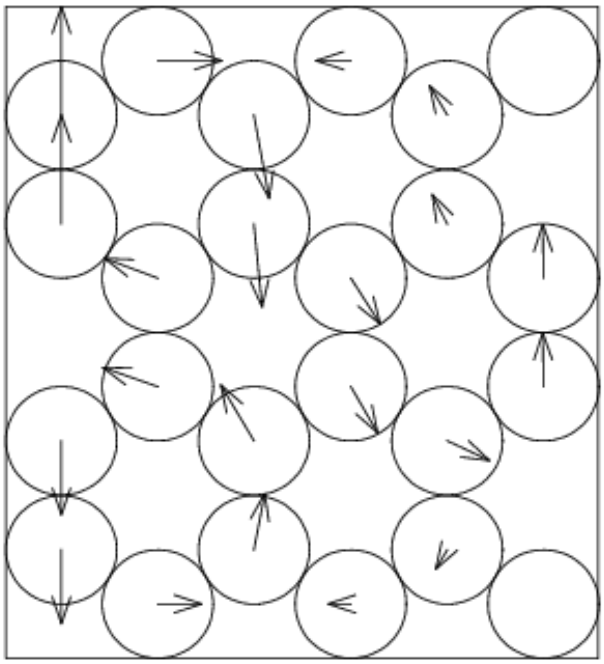

Figure 4.4: *Unjamming the honeycomb lattice.* A subpacking of size $\mathbf{N}_c = (3, 2)$ unit cells of the infinite honeycomb lattice disk packing is placed inside a hard-wall rectangular container. The arrows in the figures given here show the direction of motion of the spheres $\mathbf{V}$ in the linear unjamming motion, scaled by some arbitrary constant to enhance the figure.

4.4.3 Rigidity Matrix for Periodic Boundary Conditions

In this section we give more details on using the randomized linear programming approach to test for local, collective and strict jamming in ideal packings with periodic boundary conditions. An outline of the actual computational algorithm along with representative results is given in Section 4.6.1.

The boundary conditions enter in our formulation through the rigidity matrix $\mathbf{A}$, which is in general of dimension $[N_f \times M]$, where N_f is the total number of degrees of freedom (including boundaries), and M is the total number of constraints. Fixing (freezing) particles (degrees of freedom) is easily done by simply deleting the rows of the rigidity matrix corresponding to those particles (degrees of freedom). For hard-wall boundaries, we add a potential contact to the contact network from each sphere close to a wall to the closest point on the wall, and fix the endpoint on the wall. Such fixed points of contact and fixed spheres j, called *fixed nodes* in tensegrity terminology, are simply handled by transferring the corresponding term $\Delta\mathbf{r}_j^T\mathbf{u}_{ij}$ to the right-hand side of the constraints in (4.7).

Periodic BCs are best handled by considering the unit cell of the packing and considering each contact between an original i and an image particle $\widehat{j}\,(\mathbf{n}_{ij})$ to be a contact between particles i and j. The vector of d integers $\mathbf{n}_{ij}$ specifies how many unit cells the contact $\{i, j\}$ "crosses" over (see also Section 3.2.3). If we think of the periodic packing as being embedded in a flat torus defined by the lattice $\mathbf{\Lambda}$, the integer data $\mathbf{N} = (\mathbf{n}_{ij})$ is now to be considered part of the contact network (i.e., the *connectivity* information for the graph representing the contact network), and specifies how the network wraps around the torus. The entries in the rigidity matrix in column $\{i, j\}$ corresponding to particles i and j are just as for two particles at a relative position

$$\mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j + \mathbf{\Lambda}\mathbf{n}_{ij}.$$

If the lattice $\mathbf{\Lambda}$ changes, as when considering strict jamming, there are additional rows in the rigidity matrix corresponding to those degrees of freedom, and there is a column corresponding to the condition that the boundary deformation be volume non-increasing. We discuss these in details shortly.

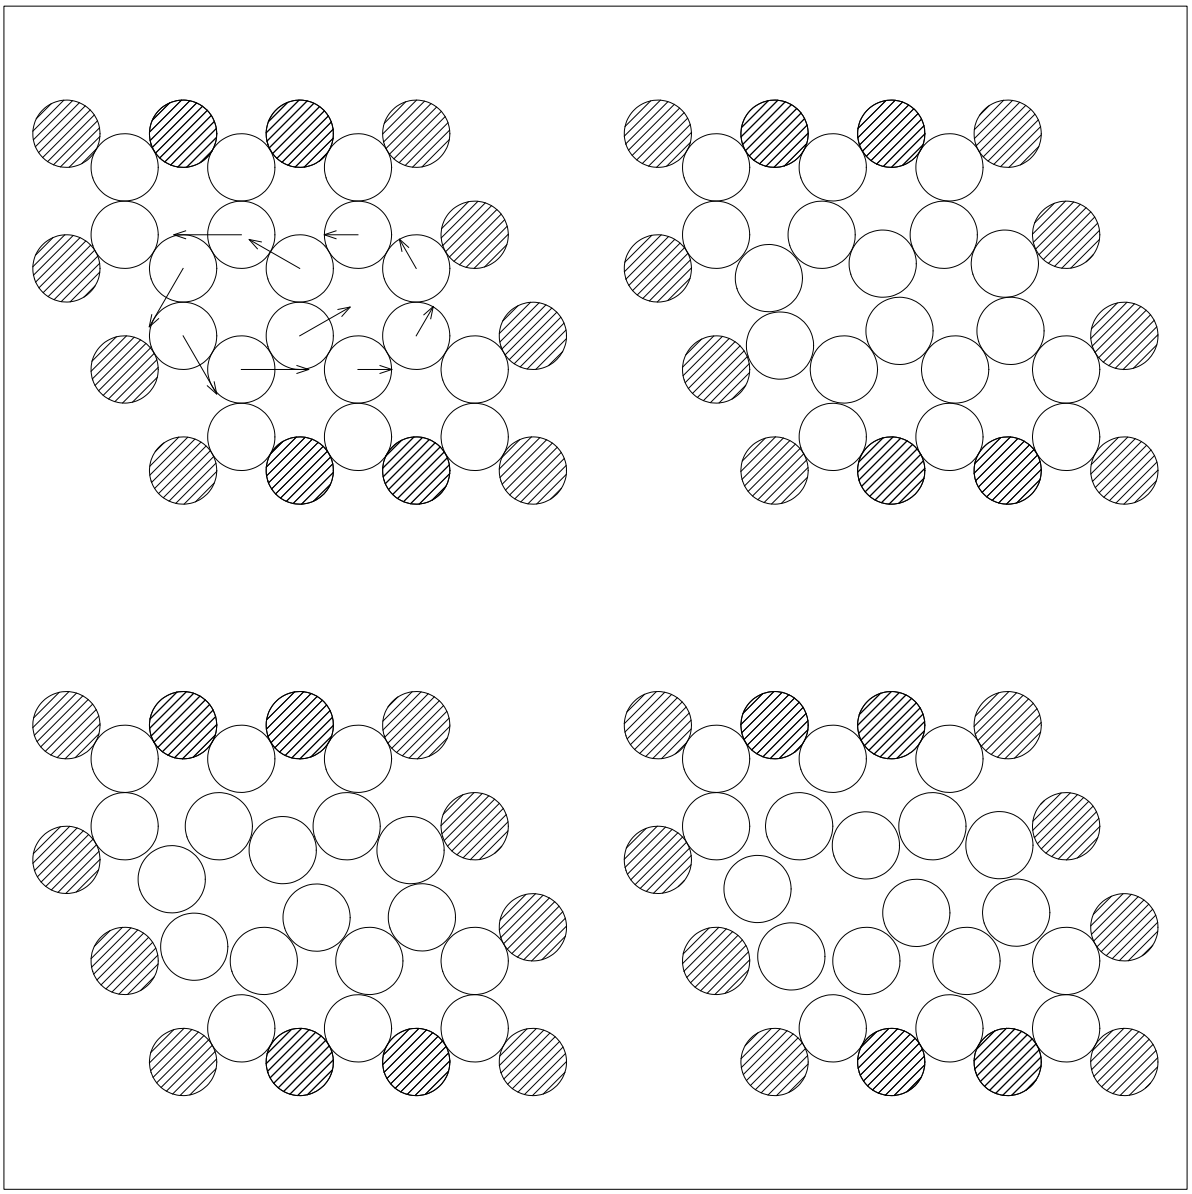

Figure 4.5: *Unjamming the honeycomb lattice.* A subpacking of size $\mathbf{N}_c = (3, 3)$ of an infinite honeycomb packing is pinned by freezing all neighboring image disks. A representative unjamming motion is shown as a sequence of several frames between times $t = 0$ and $t = 1$ (in the order top left, top right, bottom left and bottom right). The unshaded disks represent the particles in the generating packing $\mathcal{P}(\widehat{\mathbf{R}})$, while the shaded ones are image disks that touch one of the original disks.

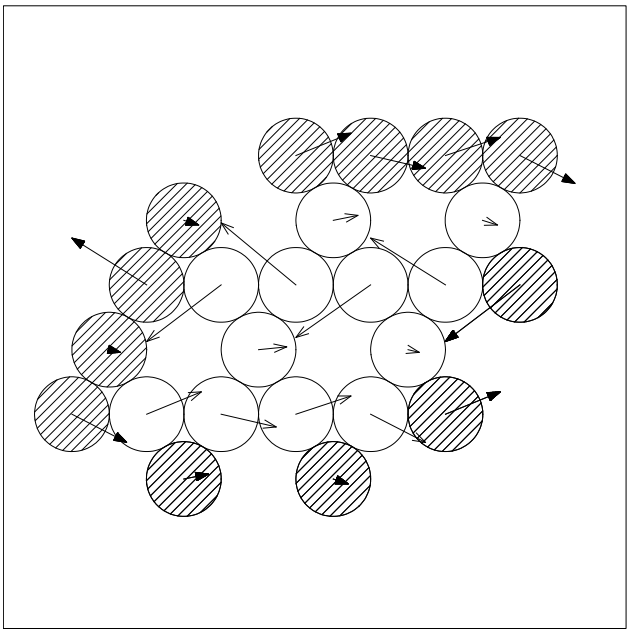

Figure 4.6: *Unjamming the Kagomé lattice.* Periodic boundary conditions are used with $\mathbf{N}_c = (2, 2)$.

4.4.3.1 Local Jamming

Recall that the condition for a packing to be locally jammed is that *each particle be fixed by its neighbors*. This is easy to check. Namely, each sphere has to have at least $d + 1$ contacts with neighboring spheres, not all in the same d-dimensional hemisphere. This can be tested in any dimension by solving a small linear program, and in two and three dimensions one can use more elementary geometric constructions.

We prefer the LP approach because it is in the spirit of this work and because of its dimensional independence, and so we present it here. Take a given sphere i and its set of contacts $\{\mathbf{u}_{i,*}\}$, and put these as rows in a matrix $\mathbf{A}_i^T$. Then solve the local portion of (4.11) (using the simplex algorithm),

$$\begin{aligned} &\max_{\Delta\mathbf{r}_i} (\mathbf{A}_i\mathbf{e})^T \Delta\mathbf{r}_i \\ \text{such that} \quad &\mathbf{A}_i^T \Delta\mathbf{r}_i \geq -\mathbf{h}_{\mathcal{N}(i)}. \end{aligned} \tag{4.15}$$

The local load $\mathbf{b}_i = \mathbf{A}_i\mathbf{e}$ can be replaced with two random loads of opposite direction, which is more suitable when larger gaps are present. When testing for jamming in ideal packings, we remove the rattlers from the packing before proceeding with tests for collective or strict jamming. Notice that checking each sphere for local jamming using (4.15) only once is not enough under this removal scheme. Namely, once a rattling sphere is removed, this removes some contacts from the packing and can make other spheres not locally jammed. We have observed that sometimes, particularly in two-dimensional systems, *all* disks can be removed on the basis of just the local jamming testing.

Of course we can define higher orders of local jamming by asking that *each set of n spheres be fixed by its neighbors*, called *n-stability* in Ref. [136]. However, for $n > 1$ it becomes combinatorially too difficult to check for this because the number of subsets to be tested grows exponentially. Computationally, we have found testing for local jamming using (4.15) to be quite efficient and simple.

4.4.4 Collective Jamming

The randomized LP algorithm was designed to test for collective jamming in large packings, and in this case the linear program (4.10) that needs to be solved is very large and sparse. Notice that boundary conditions are only involved when making the list of contacts in the contact network and deciding if certain spheres or contact points are fixed. In the case of periodic boundary conditions, we simply add the usual contacts

between original spheres near the boundary of the unit cell and any nearby periodic image spheres, and also fix $\Delta\mathbf{\Lambda} = 0$ in Eq. (4.14).

4.4.5 Strict Jamming

To extend the notion of collective jamming to strict jamming we introduced deformations of the boundary. In the case of periodic packings, *the lattice* $\mathbf{\Lambda}$ *is the boundary*. Therefore, the only difference with collective jamming is that we will now allow the lattice to change while the spheres move, i.e., $\Delta\mathbf{\Lambda} \neq 0$ in (4.14). The *lattice deformation* $\Delta\mathbf{\Lambda}$ will become part of the unknowns in (4.10), but since it too enters linearly in (4.14), we still get a linear program, only with coefficient matrix $\mathbf{A}$ augmented with new (denser) rows. These rows have nonzero entries in the columns corresponding to contacts across the periodic boundary, i.e., contacts with nonzero $\mathbf{n}_{ij}$ [147]. Note that one could express the positions of the particles relative to the unit cell, as we did in Chapter 3, however, this makes the motions of the particles parabolic when the lattice deforms, and we are focusing here on linear programming and linear (straight-line) displacements of the particles.

4.4.5.1 Macroscopic Strain

Obviously, we cannot allow the volume of the unit cell to enlarge, since the unit cell is in a sense the container holding the packing together. Therefore, we only consider *volume-non-increasing* continuous lattice deformations $\Delta\mathbf{\Lambda}(t)$,

$$\left|\widetilde{\mathbf{\Lambda}}\right| = \left|\mathbf{\Lambda} + \Delta\mathbf{\Lambda}(t)\right| \leq \left|\mathbf{\Lambda}\right|, \text{ for } t > 0, \tag{4.16}$$

where $||$ denotes determinant. We now think of $[\Delta\mathbf{R}(t), \Delta\mathbf{\Lambda}(t)]$ as an unjamming motion and focus on linear motions $\Delta\mathbf{\Lambda}(t) = \mathbf{W}t$, $\mathbf{W} = \text{const.}$ and the final small deformations $\Delta\mathbf{\Lambda} = \Delta\mathbf{\Lambda}(1)$, and consider first-order linearizations of the non-expansion nonlinear constraint (4.16).

The linearized version of (4.16) is

$$\text{Trace}[(\Delta\mathbf{\Lambda})\mathbf{\Lambda}^{-1}] \leq 0, \tag{4.17}$$

and this is just one extra linear constraint to be added to the linear program (4.10). The matrix inside the trace can be interpreted as a (macroscopic) *strain tensor* $\boldsymbol{\varepsilon}$,

$$\boldsymbol{\varepsilon} = (\Delta\mathbf{\Lambda})\,\mathbf{\Lambda}^{-1}. \tag{4.18}$$

To see this, note that the deformation of the lattice causes a displacement of the lattice point $\mathbf{p}$ (this is a vector of integer lattice coordinates) positioned at $\mathbf{r_p} = \mathbf{\Lambda p}$ of

$$\Delta\mathbf{r_p} = (\Delta\mathbf{\Lambda})\,\mathbf{p} = \left[(\Delta\mathbf{\Lambda})\left(\mathbf{\Lambda}^{-1}\right)\right]\mathbf{r_p},$$

which gives the strain (tensor)

$$\boldsymbol{\varepsilon} = \boldsymbol{\nabla}_{\mathbf{r}}\,(\Delta\mathbf{r}) = (\Delta\mathbf{\Lambda})\,\mathbf{\Lambda}^{-1}.$$

The strain tensor should be symmetric, $\boldsymbol{\varepsilon}^T = \boldsymbol{\varepsilon}$. It turns out that this condition eliminates rotations of the lattice, since rotations of the lattice produce skewsymmetric strains. Because rotations of the lattice belong to the category of trivial motions, which we want to eliminate from the onset, we will use the strain $\boldsymbol{\varepsilon}$ as a variable instead of the deformed lattice $(\mathbf{I} + \boldsymbol{\varepsilon})\,\mathbf{\Lambda}$. This is only strictly valid for infinitesimal lattice deformations; finite lattice deformations in this model are to be considered as an integral of infinitesimal deformations with symmetric strain.

With this in mind, it is best to use the symmetric strain tensor as a variable, that is, to represent motions with $\Delta\mathbf{Q} = (\Delta\mathbf{R}, \boldsymbol{\varepsilon})$. In order to simplify matrix algebra later on, we will need to represent the strain as a vector $\widehat{\boldsymbol{\varepsilon}}$ with $d\,(d+1)\,/2$ coordinates containing only the lower or only the upper triangle of the strain components. How we order the triangle into the vector is immaterial and a matter of convention (e.g., ordering by diagonals starting from the main diagonal or ordering by columns). This ordering establishes a correspondence $s \equiv (p, q)$ between component $\widehat{\boldsymbol{\varepsilon}}_s$ and component $\boldsymbol{\varepsilon}_{p,q} = \boldsymbol{\varepsilon}_{q,p} = \widehat{\boldsymbol{\varepsilon}}_s$. The usual convention (in

three dimensions) is to use the column vector of strains

$$\widehat{\boldsymbol{\varepsilon}} = \begin{bmatrix} \varepsilon_{1,1} \\ \varepsilon_{2,2} \\ \varepsilon_{3,3} \\ 2\varepsilon_{2,3} \\ 2\varepsilon_{1,3} \\ 2\varepsilon_{1,2} \end{bmatrix},$$

which contains additional factors of 2 (see section 4.4.5.3). Using this notation, the distance between the centroids of two particles as the particles move and the lattice deforms is given with

$$l_{ij} = \|\mathbf{r}_{ij} + (\Delta\mathbf{r}_i - \Delta\mathbf{r}_j) + \boldsymbol{\varepsilon}\boldsymbol{\Lambda}\mathbf{n}_{ij}\| = \|\mathbf{r}_{ij} + \mathbf{T}_{ij}\Delta\mathbf{r} + \mathbf{S}_{ij}\widehat{\boldsymbol{\varepsilon}}\|,$$

where $\mathbf{T}_{ij}$ is a $[d \times Nd]$ matrix (with simple structure) and $\mathbf{S}_{ij}$ is a $\left[d \times \frac{d(d+1)}{2}\right]$ matrix. The condition (4.17) simply becomes

$$\sum_{s=1}^{d} \widehat{\boldsymbol{\varepsilon}}_s \leq 0, \tag{4.19}$$

and can be included as an extra column (constraint) in the rigidity matrix.

The motivation for the category of strict jamming and its above interpretation in the periodic case should be clear: Changing the lattice in a volume non-increasing way models *macroscopic non-expansive strain* (i.e., a compressive macroscopic stress) and is therefore of great relevance to studying the macroscopic mechanical properties of random packings (see Ref. [111]). We also again point out that strict jamming is (significantly) stronger than collective jamming for periodic boundary conditions, particularly in two-dimensional packings. This point is illustrated in Fig. 4.7, which shows an unjamming motion involving a deformation of the lattice, even though this lattice packing is collectively jammed. Periodic boundary conditions are often used to model infinite systems, in the hope that a jammed periodic packing will produce a "jammed" infinite packing (for example, in the sense of uniform stability [136]) when periodically replicated in all directions. A simple counting argument demonstrates that isostatic collectively jammed periodic packings *cannot* generate "jammed" infinite packings because they have too few contacts, but isostatic strictly jammed periodic packings *can* since they have enough contacts (due to the inclusion of additional degrees of freedom for the deforming lattice). We give those counting arguments at the end of this section.

4.4.5.2 Rigidity Matrix

The column of $\mathbf{A}$ corresponding to the contact $\{i, j\}$ is nothing more than the gradient of the distance between particles i and j with respect to the degrees of freedom in the packing

$$\mathbf{A}_{ij} = \boldsymbol{\nabla}_{\mathbf{Q}}(l_{ij}) = \begin{bmatrix} \mathbf{A}_{ij}^{(\mathbf{R})} \\ - \\ \mathbf{A}_{ij}^{(\boldsymbol{\epsilon})} \end{bmatrix}$$

The first piece of this is the corresponding column of the usual rigidity matrix

$$\mathbf{A}_{ij}^{(\mathbf{R})} = \boldsymbol{\nabla}_{\mathbf{R}}(l_{ij}) = \begin{matrix} \\ i \rightarrow \\ \\ j \rightarrow \\ \end{matrix} \begin{bmatrix} \vdots \\ \mathbf{u}_{ij} \\ \vdots \\ -\mathbf{u}_{ij} \\ \vdots \end{bmatrix},$$

and the second piece is due to the periodicity of the contact network

$$\mathbf{A}_{ij}^{(\boldsymbol{\epsilon})} = \boldsymbol{\nabla}_{\widehat{\boldsymbol{\varepsilon}}}(l_{ij}) = \left\{\boldsymbol{\nabla}_{\widehat{\boldsymbol{\varepsilon}}}\left[\boldsymbol{\varepsilon}\left(\widehat{\boldsymbol{\varepsilon}}\right)\boldsymbol{\Lambda}\mathbf{n}_{ij}\right]\right\}\mathbf{u}_{ij},$$

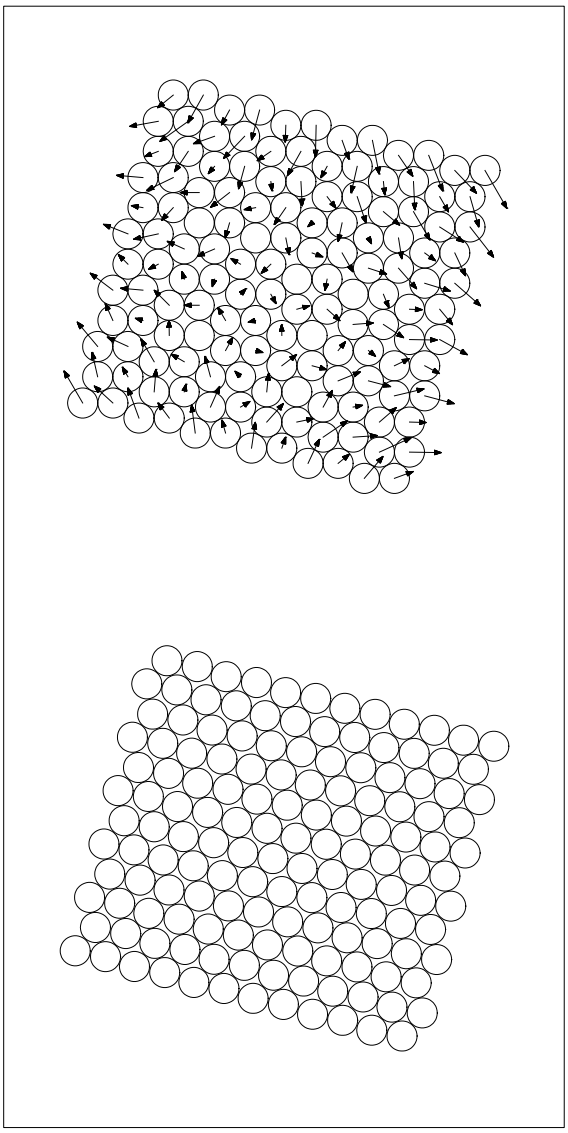

Figure 4.7: *Example of a lattice deformation.* The above periodic packing (packing 3 in Ref. [142]) is collectively jammed, but not strictly jammed. It can be continuously sheared toward the triangular lattice by deforming the lattice in a volume non-decreasing manner, as shown here.

which we can also write in indicial form suitable for computational use as [recall that $s \equiv (p,q)$ determined how the vectorization of the upper/lower triangle of $\boldsymbol{\varepsilon}$ was done to obtain $\widehat{\boldsymbol{\varepsilon}}$]

$$\left(\mathbf{A}_{ij}^{(\boldsymbol{\epsilon})}\right)_s = \frac{1}{2}\left[\left(\boldsymbol{\Lambda}\mathbf{n}_{ij}\right)_p \left(\mathbf{u}_{ij}\right)_q + \left(\boldsymbol{\Lambda}\mathbf{n}_{ij}\right)_q \left(\mathbf{u}_{ij}\right)_p\right], \tag{4.20}$$

or in matrix form,

$$\mathbf{A}_{ij}^{(\boldsymbol{\epsilon})} = \mathrm{Symm}\left\{\mathbf{u}_{ij}\left(\boldsymbol{\Lambda}\mathbf{n}_{ij}\right)^T\right\}. \tag{4.21}$$

We have not yet derived second-order order derivatives with respect to the lattice degrees of freedom, as these are significantly more complicated and are not used in this work. It is however worth carrying out the explicit calculation.

4.4.5.3 Macroscopic Stress

The condition of mechanical equilibrium at zero external load

$$\mathbf{A}\mathbf{f} = \mathbf{0},$$

reduces to the Nd microscopic force balances for each particle

$$\sum_{j \in \mathcal{N}(i)} f_{ij} \mathbf{A}_{ij}^{(\mathbf{R})} = \mathbf{0},$$

as well as the $d\left(d+1\right)/2$ conditions that there be no residual non-isotropic macroscopic stresses

$$\widehat{\boldsymbol{\sigma}} = \frac{1}{|\boldsymbol{\Lambda}|} \sum_{ij} f_{ij} \mathbf{A}_{ij}^{(\boldsymbol{\epsilon})} = p\hat{\mathbf{I}}, \tag{4.22}$$

where the pressure $p \geq 0$ is the Lagrange multiplier ("force") for the strain constraint (4.19). Here $\widehat{\boldsymbol{\sigma}}$ is the vectorized version of the upper or lower triangle of the symmetrized *macroscopic stress* (tensor) $\boldsymbol{\sigma}$, and we normalized with the reciprocal unit cell volume $|\boldsymbol{\Lambda}|$ in order to get the correct units of stress. This is expected since stress is the strain gradient of the energy *density*, and not of energy. The stress tensor is given by

$$\boldsymbol{\sigma} = \frac{1}{|\boldsymbol{\Lambda}|} \sum_{ij} f_{ij} \left\{ \text{Symm} \left[\mathbf{u}_{ij} \left(\boldsymbol{\Lambda} \mathbf{n}_{ij} \right)^T \right] \right\}, \tag{4.23}$$

which more clearly displays the tensor character through the use of the diadic product $\mathbf{u}_{ij} \left(\boldsymbol{\Lambda} \mathbf{n}_{ij} \right)^T$.

4.4.5.4 Linearization of Strain Condition

In Section 4.3.1 we demonstrated that the linearization of the impenetrability condition was in a certain sense rigorous for hard spheres: Any motion that satisfies the linearizing constraints satisfies also the true non-linear constraints. This is an important property of sphere packings that is used heavily in this Chapter, and it is therefore important to demonstrate that it continues to be true even if we allow the periodic lattice to change. In practice this means that any (first-order) solution obtained by the linear programming algorithm can be appropriately scaled to obtain a truly feasible displacement in which some particles strictly lose contact and the volume of the unit cell strictly decreases.

Motivated by Ref. [148], we consider the nonlinear corrections to (4.17) by expanding (4.16) to second order, to get

$$\begin{aligned} |\boldsymbol{\Lambda} + \Delta\boldsymbol{\Lambda}| &= |\boldsymbol{\Lambda}| \, |\mathbf{I} + \boldsymbol{\varepsilon}| = \\ |\boldsymbol{\Lambda}| \prod_{i=1}^{d} (1 + \lambda_i) &= |\boldsymbol{\Lambda}| \left(\sum_i \lambda_i + \sum_{i>j} \lambda_i \lambda_j + \text{higher order terms} \right), \end{aligned}$$

where λ_i are the eigenvalues of the strain, which are all real because of its symmetry and also have a nonpositive sum due to (4.17). If $\text{Tr}\left(\boldsymbol{\varepsilon}\right) = \sum \lambda_i < 0$, then the first order term will dominate for sufficiently small deformations and the nonlinear constraint (4.16) will be satisfied. Furthermore, if $\text{Tr}\left(\boldsymbol{\varepsilon}\right) = 0$, then we have that

$$\sum_{i>j} \lambda_i \lambda_j = -\frac{1}{2} \sum_i \lambda_i^2 < 0$$

for any nontrivial deformation, which shows that the second-order term is of the correct sign and the volume of the unit cell strictly decreases for sufficiently small deformations (i.e., it strictly decreases at the second-order level).

4.4.6 Strict Jamming and Macroscopic Rigidity

The interpretation of the conditions for strict jamming in terms of a macroscopic stress and strain tensor is physically important, and in this section we elaborate on some technical details related to this interpretation. In order to justify the interpretation of $\boldsymbol{\varepsilon} = (\Delta\boldsymbol{\Lambda})\,\boldsymbol{\Lambda}^{-1}$ as a *macroscopic* strain and of $\boldsymbol{\sigma} = \sum_{ij} f_{ij}\left[\mathbf{u}_{ij}^T\,(\boldsymbol{\Lambda}\mathbf{n}_{ij})\right]$ as a (un-symmetrized) *macroscopic* stress we look at the infinite periodic packing obtained by replicating the generating packing along all dimensions (i.e., unwrapping the torus), and consider whether these tensors are invariant to the choice of the unit cell.

In particular, a unit cell may be chosen to be larger and contain several multiples $\mathbf{N}_c$ of the unit cell used when testing for jamming. This new enlarged packing may not be jammed even if the smaller packing was jammed, due to the existence of unjamming motions involving large numbers of particles. A reasonable first definition of jamming appropriate for infinite repetitive packings is that every finite subpacking of the infinite packing be jammed. This is the definition of *finite stability* in Ref. [136], and in that reference some shortcomings are discussed an a stronger *uniform stability* condition is defined. Here we consider counting arguments for jamming when the unit cell is enlarged, demonstrating that the collective jamming condition is insufficient to give finite stability for hard sphere packings, where as the strict jamming condition may be sufficient.

4.4.6.1 Invariance of the Macroscopic Strain

The unit cell of a periodic system is not uniquely defined. For example, one may take a larger unit cell as the reference cell, i.e., take the lattice to be a sublattice of the original lattice

$$\boldsymbol{\Lambda}' = \boldsymbol{\Lambda}\mathbf{N}_c,$$

where $\mathbf{N}_c$ is a diagonal matrix with positive integer entries. Now consider a lattice deformation with periodicity determined by $\boldsymbol{\Lambda}$ in the primed notation, where $\Delta\boldsymbol{\Lambda}' = (\Delta\boldsymbol{\Lambda})\,\mathbf{N}_c$. The macroscopic strain is

$$\boldsymbol{\varepsilon}' = (\Delta\boldsymbol{\Lambda})\,\boldsymbol{\Lambda}^{-1} = (\Delta\boldsymbol{\Lambda})\left(\mathbf{N}_c\mathbf{N}_c^{-1}\right)\boldsymbol{\Lambda}^{-1} = \boldsymbol{\varepsilon},$$

i.e., the strain is independent of the exact choice of the unit cell. This is a very important invariance property which makes our results more physical.

4.4.6.2 Invariance of the Macroscopic Stress

The expression for the macroscopic stress (4.22) appears to be invariant with respect to choosing a different unit cell as the reference cell, as it should on physical grounds. However, this is difficult to show in general as the $\mathbf{n}_{ij}$'s depends non-trivially on the choice of the cell.

Take a contact $\{i,j\}$ carrying force f_{ij} and having $\mathbf{n}_{ij} = \langle 1,1\rangle$ when the lattice is $\boldsymbol{\Lambda}$. Now take $[N_x \times N_y]$ unit cells put together as a new unit cell, i.e., take $\boldsymbol{\Lambda}' = \boldsymbol{\Lambda}\mathbf{N}_c$ with

$$\mathbf{N}_c = \begin{bmatrix} N_x & \\ & N_y \end{bmatrix}$$

as the new lattice. With this new choice of lattice, the original contact gets replaced with $N_x - 1$ contacts with $\mathbf{n}_{i'j'} = \langle 0,1\rangle$, $N_y - 1$ contacts with $\mathbf{n}_{i'j'} = \langle 1,0\rangle$, and 1 contact with $\mathbf{n}_{i'j'} = \langle 1,1\rangle$. All of these image contacts carry the same force f_{ij} and have the same direction vector $\mathbf{u}_{ij}$. Looking at equations (4.21) and (4.22), we see that the portion in the expression for the macroscopic stress which depends on the choice of the unit cell, $\boldsymbol{\Lambda}\mathbf{n}_{ij}$, gets replaced with

$$\frac{\boldsymbol{\Lambda}'}{N_xN_y}\sum_{\{i',j'\}}\mathbf{n}_{i'j'} = \boldsymbol{\Lambda}\left\{\frac{(N_x-1)}{N_xN_y}\mathbf{N}_c\begin{bmatrix}0\\1\end{bmatrix} + \frac{(N_y-1)}{N_xN_y}\mathbf{N}_c\begin{bmatrix}1\\0\end{bmatrix} + \frac{1}{N_xN_y}\mathbf{N}_c\begin{bmatrix}1\\1\end{bmatrix}\right\} \;=\; \boldsymbol{\Lambda}\mathbf{n}_{ij}.$$

This vector is therefore unchanged by the change of the unit cell, and therefore the macroscopic stress is invariant under this change of unit cells, as it should be. This is not quite a polished proof, but indicates that expression (4.22) is physical.

The relation between microscopic forces and the macroscopic stress in a force network has been considered in the granular materials community, and various related expressions of different generality have been derived [149, 150, 151, 152]. Expression (4.22) for the macroscopic strain in a periodic network has not been previously derived in this form to our knowledge. However, it is closely related to results based on averaging that have been obtained and used in the context of disordered contact networks, as we demonstrate next. Consider the form of the stress tensor given in Eq. (4.23). First, notice that only the bars with nonzero $\mathbf{n}_{ij}$ contribute to the macroscopic stress, i.e., only the bars that cross the "boundary" of the unit cell contribute. Compare this to the expressions found in Ref. [149] for the macroscopic stress in a disordered network (recast into a form more suitable for our presentation)

$$\boldsymbol{\sigma} = \frac{1}{V} \sum_{\{i,j\} \in V} f_{ij} l_{ij} \left(\mathbf{u}_{ij} \mathbf{u}_{ij}^T \right) = \frac{1}{V} \sum_{i \in V,\ j \notin V} f_{ij} \left(\mathbf{u}_{ij} \mathbf{r}_i^T \right), \tag{4.24}$$

the second of which also only involves microscopic forces crossing the boundary of a given reference (averaging) volume V, i.e. only the bars $\{i,j\} \in \partial V$. For a periodic system it is natural to take the unit cell as the averaging volume. Consider a contact $\{i,j\}$ with nonzero $\mathbf{n}_{ij}$. It will appear twice in the sum in Eq. (4.24), once as the "original" with direction $\mathbf{u}_{ij}$, and once as an "image" $\{i',j'\}$ with $\mathbf{u}_{i'j'} = -\mathbf{u}_{ij}$ and $\mathbf{r}_{i'} = \mathbf{r}_i - \boldsymbol{\Lambda}\mathbf{n}_{ij} + l_{ij}\mathbf{u}_{ij}$. Therefore the contribution from this bar to the averaged macroscopic stress in Eq. (4.24) is

$$\frac{1}{|\boldsymbol{\Lambda}|} f_{ij} \left[\mathbf{u}_{ij}^T \left(\boldsymbol{\Lambda} \mathbf{n}_{ij} \right) \right] - \frac{1}{|\boldsymbol{\Lambda}|} f_{ij} l_{ij} \left(\mathbf{u}_{ij} \mathbf{u}_{ij}^T \right).$$

The first term in this expression is identical to the one in Eq. (4.23). If we take a large unit cell, in the spirit of the averaging in Eq. (4.24), the second term will become negligible.

4.4.6.3 Counting Contacts for Repetitive Packings

Unlike local jamming, both collective and strict jamming are global properties of the packing, and impose constraints on the total number of contacts needed in order to ensure jamming (we are focusing here on ideal packings, for which contacts can be counted exactly). Namely, the number of impenetrability constraints (contacts) must be larger than the number of degrees of freedom by at least one. This is because $k+1$ *inequality* constraints are needed to eliminate k degrees of freedom (just as $d+1$ neighbours are needed for local jamming). Such counting arguments have been used frequently since Maxwell employed them to study bar frameworks; see, for example, Ref. [137].

We apply here such counting arguments to repetitive packings obtained by periodically repeating a generating packing $P(\widehat{\mathbf{R}})$ on a lattice $\boldsymbol{\Lambda}$ (see section 4.4). Assume first that with imposed periodic boundary conditions (4.14) for a given choice of unit cell, which contains N spheres and M interparticle contacts, the packing is collectively jammed. Furthermore, assume that the packing is *isostatic*, i.e., it has the minimal number of contacts needed for jamming. It is expected that such isostaticity is a generic property of disordered *polydisperse* packings (see, for example, Ref. [143] for a discussion of randomly perturbed triangular lattice disk packings), since additional contacts imply some degeneracy in the packing, as we discuss in detail in Section 5.4.

In Section 4.3.1.1 we gave the counting constraints for periodic boundary conditions, Eq. (4.8) for collective jamming and Eq. (4.9) for strict jamming. What does jamming for this particular choice of unit cell (i.e., torus) imply about the infinite repetitive (cover) packing? This is an important question because periodic boundary conditions are used to emulate infinite (or very large) systems, rather than to study packings on a flat torus. This question is difficult and has not yet been fully answered. However, some important conclusions are simple to deduce. For example, any unjamming motion for a particular choice of unit cell is also an (repetitive) unjamming motion of the infinite packing, and also for all integer-multiples of unit cell, $\widetilde{\boldsymbol{\Lambda}} = \boldsymbol{\Lambda}\mathbf{N}_c$, where $N_c = \det(\mathbf{N}_c) > 1$.

However, the converse is not true: In other words, the fact that a packing is jammed with one choice of unit cell does not imply that it will be so for other choices. In fact, for collectively jamming it is *guaranteed* that the packing will *not* be jammed for larger unit cells. Specifically, Eq. (4.8) becomes

$$\widetilde{M} = N_c M = dN_c N - (d-1)N_c < d\widetilde{N} - (d-1),$$

which shows that the packing cannot be jammed with the new choice of unit cell. For strict jamming, however, this is not the case, and it is possible that the packing is strictly jammed for all choices of unit cell, as seen by writing constraint (4.9) for the new boundary conditions:

$$\widetilde{M} = N_c M = \begin{cases} 2N_c N + N_c > 2\widetilde{N} + 1 \text{ for } d = 2 \\ 3N_c N + 3N_c > 3\widetilde{N} + 3 \text{ for } d = 3 \end{cases} .$$

This shows that the packing has redundant contacts with the new boundary conditions. This does not prove that the packing is jammed with the new choice of unit cell, but it *may* be.

This argument demonstrates that strictly (but not collectively) jammed isostatic periodic packings are candidates for "jammed" infinite packings. In this work we have focused on finite systems, and generalizations to infinite packings are discussed in Ref. [136]. In particular, *uniformly stable* infinite packings can be considered the analog of strictly jammed finite packings.

4.5 Beyond the ASD: Transition States and Paths

In this section we more formally explore the issue of unjamming motions in the presence of nonzero interparticle gaps, i.e., for non-ideal packings. Assume that a configuration $\mathbf{R}_J$ represents a jammed ideal packing of N particles with packing fraction ϕ_J, where there are M interparticle contacts. Next, decrease the density slightly by reducing the scaling the particle sizes by a factor $\mu = 1-\delta$, so that the packing fraction is lowered to $\phi = \phi_J (1-\delta)^d$. The definition of jamming (rigidity) we used for ideal packings is that in the limit $\delta \to 0$ the jamming basin $\mathcal{J}_{\mathbf{\Delta R}} \to \{\mathbf{R}_J\}$. Here, *we define a jammed packing as one in which growth of the particles must eventually lead to the unique jamming point under consideration, regardless of the (thermal) motions of the particles.* That is, the jamming basins correspond to pockets of the feasible (available) region of configuration space that are not connected to other basins. In this sense, every jammed packing has a maximal allowed dilution $\delta_{\max} > 0$ beyond which transitions to other basins become possible. Our goal in this section is to identify the *transition states* and *transition paths* between basins explicitly, given a jammed ideal packing. Work is underway to implement and apply the algorithm described in this section for small packings of binary hard disks.

4.5.1 Vertices of $\mathcal{J}_{\mathbf{\Delta R}}$

The jamming basin itself is a region in configuration space bounded by the impenetrability conditions, which are quadratic inequalities that prohibit the configuration point from entering a hypercylinder. Formally, that is, the feasible region is the complement of the union of hypercylinders, or the set of solutions to a collection of $N(N-1)/2$ inequalities (one for each pair of particles $\{i,j\}$). At a given δ, only a subset of $M + \Delta M$ of these inequalities will be *active*, reducing to the M true contacts in the jamming limit. For sufficiently small δ, the jamming basin approaches the *jamming polytope* (a closed polyhedron in arbitrary dimension) $\mathcal{P}_{\mathbf{\Delta R}} \subseteq \mathcal{J}_{\mathbf{\Delta R}}$ [52, 145], as discussed in Section 4.3.1.1.

Each face of the polytope $\mathcal{P}_{\mathbf{\Delta R}}$ corresponds to one of the contacts being active, i.e., the two particles involved touching exactly. The configurations where as many contacts are active as possible, that is, any motion of the particles must break some of the active contacts, are the *vertices* of $\mathcal{P}_{\mathbf{\Delta R}}$. When the curvature of the constraints is taken into account, the faces of the polytope become curved (outward, i.e., concave), however, the combinatorial structure still remains like that of a polytope. In particular the vertices become stretched and more pointed as δ increases. At the critical $\delta = \delta_{\max}$, it is exactly one of these vertices that opens up. That is, the transition state is a vertex of $\mathcal{J}_{\mathbf{\Delta R}}$, and the transition path is a motion whose linear component (velocity) $\mathbf{V} = \dot{\mathbf{R}}$ is perpendicular to all of the normal vectors of the constraints active at the vertex. It is now clear how to identify transition states and paths: Track a vertex of $\mathcal{J}_{\mathbf{\Delta R}}$ as δ grows away from 0, and identify the point where the vertex becomes open. Do this for all vertices, and this will identify all the transition states and paths. Note that new vertices may appear for $\delta > 0$ due to the presence of near-contacts in the packing that are not active at $\delta = 0$.

4.5.1.1 Tracking the Vertices

We make the following non-degeneracy assumption: At every vertex of $\mathcal{J}_{\mathbf{\Delta R}}$ there are exactly $N_f = (N-1)d$ active contacts. This is not true for certain regular arrangements such as crystal packings, however, it should be true for "disordered" packings. It is not necessary that the jammed packing under consideration be isostatic, rather, it is required that the polytope $\mathcal{P}_{\mathbf{\Delta R}}$ be a *simple polytope*. If the packing does happen to be isostatic, then $\mathcal{P}_{\mathbf{\Delta R}}$ is also simplicial, i.e., it is a simplex, as illustrated in Fig. 4.1.

Under this assumption, it is possible to change δ by shrinking all the particles uniformly, while maintaining in contact all of the contacts in the contact network at the vertex in question. This is the same idea used in the Zinchenko algorithm [28], where the spheres are grown in order to produce a jammed packing instead of being shrunk away from a jammed packing. The motion/velocities of the particles i and i necessary to maintain the contact $\{i, j\}$ closed are determined from the conditions

$$\mathbf{v}_{ij}^T \mathbf{u}_{ij} = -\dot{\delta}\tilde{D}_{ij},$$

where $\tilde{D}_{ij} = (D_i + D_j)/2$. In matrix form, this gives the linear system of equations

$$\mathbf{B}^T\mathbf{V} = -\dot{\delta}\tilde{\mathbf{D}},$$

where $\mathbf{B}$ is the submatrix of $\mathbf{A}$ formed by the columns corresponding to the contacts (constraints) that are closed (active) at the vertex in question. Our non-degeneracy assumption means that the matrix $\mathbf{B}$ is invertible, and the velocities are simply

$$\mathbf{V} = -\dot{\delta}\mathbf{B}^{-T}\tilde{\mathbf{D}}.$$

Tracking a vertex fails when the matrix $\mathbf{B}$ becomes non-invertible, and this is exactly the point when the vertex becomes a transition state. Since $\mathbf{B}$ is not invertible, it has a null vector $\mathbf{U}$, and this null vector is in fact exactly the transition path, i.e., the particle velocities that lead to escape from the jamming basin through the vertex in question. At this point one simply stops tracking this vertex, and records the transition state and path that were just identified. The same transition state and path should be identified for one of the other jammed states as well, since after passing through the transition state the configuration point moves into another basin.

Tracking the vertex also fails when during the tracking process a new contact forms, i.e., one of the near contacts becomes active. At this point, the vertex splits into a number of new vertices, each of which needs to be tracked separately, so that all vertices are found and thus all transition states are identified. At the point of formation of the extra contact, the packing is isostatic, and there exist a set of interparticle forces $\mathbf{f}$ satisfying force balance. Let the column of the rigidity matrix corresponding to the newly formed contact be $\mathbf{A}_{ij}$ (see also the useful expressions in Section 6.2.2.1). If we put a force f_{ij} on the newly formed contact, then the forces on the other contacts are determined from

$$\mathbf{Bf} = -f_{ij}\mathbf{A}_{ij} \Rightarrow \mathbf{f} = -f_{ij}\mathbf{B}^{-1}\mathbf{A}_{ij}.$$

If we make the forces unique by normalizing them, $\tilde{\mathbf{D}}^T\mathbf{f} + f_{ij}\tilde{D}_{ij} = 1$, we get

$$f_{ij} = \frac{1}{\tilde{D}_{ij} - \tilde{\mathbf{D}}^T\mathbf{B}^{-1}\mathbf{A}_{ij}}.$$

At the same time, the total contact velocity for contact ij is negative (since the contact closed)

$$v_{ij} = \dot{\delta}\tilde{D}_{ij} + \mathbf{v}_{ij}^T\mathbf{u}_{ij} = \dot{\delta}\tilde{D}_{ij} + \mathbf{V}^T\mathbf{A}_{ij} = \frac{\dot{\delta}}{f_{ij}} < 0,$$

that is $f_{ij} < 0$ (since $\dot{\delta} > 0$). The newly formed contact must be kept if one wants to continue decreasing δ, and the argument above shows that the contacts with negative forces may not be broken, while any of the forces with positive forces may be broken in order to have $\dot{\delta} > 0$ and maintain nonoverlap. This was derived in an alternative way in Ref. [28], where the opposite situation $\dot{\delta} < 0$ was considered since the particles grow in size in the Zinchenko algorithm, until a jammed point where all interparticle forces are positive is reached.

In essence, the algorithm described above is the inverse of the Zinchenko packing algorithm [28]. By following the vertices away from the jamming point one is following a possible path in the Zinchenko algorithm from the liquid phase to the jamming point under consideration. Our algorithm is slightly more involved since in the Zinchenko algorithm one makes an arbitrary choice of which contact to break when a new contact forms. Here we are trying to track all vertices and thus all possibilities need to be taken into account. The algorithm is however very well-suited for implementation using recursion, since the set of vertices being tracked with possible splittings into new vertices forms a tree. In Ref. [28], it is mentioned that every now and then the algorithm ran unto "singular points" which caused very large transverse displacements of the particles. It is interesting that he notes that these singular points were observed in the density interval $0.52 - 0.58$, which is often quoted as the location of the kinetic glass transition for hard spheres.

4.5.1.2 Numerical Details

When tracking a vertex, an ODE needs to be solved, and the rigidity sub-matrix $\mathbf{B}$ inverted. Also, the points where $\mathbf{B}$ becomes singular, $|\mathbf{B}| = 0$, or when new contacts form, need to be identified so that the ODE solver can be stopped or restarted with a new contact network. Technically, this kind of problem is known as event location in ODEs [91, 131] and there are algorithms that incorporate such event-location into ODE solvers. This extra complication may not actually be necessary, and simpler schemes (requiring smaller steps in the ODE solver) may be sufficient in practice.

An additional important numerical detail is the handling of the inversion of $\mathbf{B}$, especially due to the possibility of singularities. For large packings, sparse-matrix technology and iterative solvers need to be invoked, and some of this is discussed by Zinchenko (although a lot more developments have occurred since 1990 in this field of applied mathematics). But for small problems dense linear algebra will suffice. Additionally, the inverse of $\mathbf{B}$ does not need to be formed, instead, only the adjoint

$$\mathrm{adj}\,(\mathbf{B}) = |\mathbf{B}|\,\mathbf{B}^{-1},$$

which is well-behaved even for a singular matrix. In fact, under non-degeneracy assumptions, near a singular point the adjoint becomes of rank-1, $\mathrm{adj}\,(\mathbf{B}) \sim \mathbf{U}\mathbf{U}^T$, and in fact $\mathbf{U} \sim \mathbf{V}_{\|}$ is the transition path. Therefore, the recommended ODE to solve is

$$\begin{aligned} \dot{\mathbf{R}} &= -\mathrm{adj}\left(\mathbf{B}^T\right)\tilde{\mathbf{D}} \\ \dot{\delta} &= |\mathbf{B}|\,, \end{aligned} \tag{4.25}$$

where it is assumed that the signs are chosen so that $|\mathbf{B}| > 0$ when the vertex is first formed. Note that $\mathbf{B}$ changes as the configuration $\mathbf{R}$ changes. As it can be seen from Eq. (4.25), the transition point, when $|\mathbf{B}| = 0$, corresponds to a point where $\dot{\delta}$ switches sign, that is, it goes from increasing to decreasing. One can in fact continue solving the ODE past the transition location comfortably, without encountering any singularities, and come to another jammed packing, namely, the other end of the transition path. During this process, however, some contacts may disappear when new ones are formed, and the algorithm becomes exactly like the Zinchenko algorithm.

4.5.2 The Salsburg Approach

In Ref. [52], a simplified approach to finding out when the jamming polytope becomes open is taken. Namely, only contacts that were exactly closed at close packing are considered, and also only unjamming motions that pass through the origin are considered. That is, to test whether the curved polytope is closed, Salsburg considers starting at the origin and pushing the particles in a given direction $\widehat{\Delta\mathbf{R}}$, until a wall of the curved polytope is encountered. If no such wall is encountered, that is, if none of the (original) contacts resists the motion of the particles, than this is considered an unjamming motion. This of course neglects the possibility that new contacts form and resist the motion, and also does not really identify the transition states since it only considers motions away from the origin (the true escape path will actually be curved itself). We give the mathematics of this approach here because it may be useful for large isostatic packings in identifying the "weak" vertices, that is, directions in which the polytope is very elongated. This is alike identifying soft modes of the Hessian for soft-particle systems without taking into account anharmonic effects.

Substituting $h_{ij} = \delta D$ in Eq. (4.6) we get the impenetrability constraint

$$-\frac{\Delta\mathbf{r}_{ij}^T\mathbf{u}_{ij}}{D} - \frac{\Delta\mathbf{r}_{ij}^T\Delta\mathbf{r}_{ij}}{2D^2} = \delta_{ij}(\Delta\mathbf{R}) \leq \delta\left(1 - \frac{\delta}{2}\right)$$

The jamming gap at which a given displacement $\Delta\mathbf{R}$ becomes feasible for this particular constraint is thus $\delta = 1 - \sqrt{1 - \delta_{ij}^2(\Delta\mathbf{R})} \approx \delta_{ij}(\Delta\mathbf{R})$ and therefore one can define a potential function $\tilde{\delta}(\Delta\mathbf{R}) = \max_{\{i,j\}} \delta_{ij}(\Delta\mathbf{R})$ the contours of which reveal the feasible regions at different jamming gaps (approaching the polytope of course as $\delta \to 0$). Now consider $\Delta\mathbf{R} = s\widehat{\Delta\mathbf{R}}$, where $\widehat{\Delta\mathbf{R}}$ is a given unit vector, $\left\|\widehat{\Delta\mathbf{R}}\right\| = 1$, along which one tries to move away from the *origin*. Then we have a $\delta_{ij}(s)$ that is an inverted parabola, with a maximum of height

$$\sqrt{\delta_{ij}(\widehat{\Delta\mathbf{R}})} = -\frac{\widehat{\Delta\mathbf{r}}_{ij}^T\mathbf{u}_{ij}}{\sqrt{2}\left\|\widehat{\Delta\mathbf{r}}_{ij}\right\|},$$

which is the largest jamming gap that a particular contact can support for this chosen direction of particle motion. Note that here the scalings due to s cancel in the numerator and the denominator. Therefore, the largest jamming gap at which the curved polytope is closed is given by

$$\sqrt{2\delta_{\max}} = \min_{\Delta\mathbf{R}} \max_{\{i,j\}} \frac{-\Delta\mathbf{r}_{ij}^T\mathbf{u}_{ij}}{\|\Delta\mathbf{r}_{ij}\|} = \min_{\Delta\mathbf{R}} \max_{\{i,j\}} \cos\theta_{ij},$$

where θ_{ij} is the angle between the relative displacement and the contact vector for neighbor particles i and j. The condition $\delta \leq \delta_{\max}$ is necessary for the curved polytope to be closed. Considering the fact that all polytope faces are an equal distance 1 away from the origin, for a point on the surface of the polytope, $\Delta\mathbf{r}_{ij}^T\mathbf{u}_{ij} = -\delta D$, so that $\cos\theta_{ij} \sim \|\Delta\mathbf{r}_{ij}\|^{-1}$. The minimum above is achieved at one of the vertices of the polytope, since otherwise one could move the ray and make a more oblique angle with the faces that contain the point of intersection of the ray with the polytope surface. Therefore, one could look at each of the vertices of the polytope and calculate $\sqrt{2\delta_{\text{vertex}}} = \max_{\{i,j\}} \cos\theta_{ij}$ over all the active contacts (at that vertex), and identify the vertex with the smallest δ_{vertex} as the weakest one. For an isostatic packing the polytope is a simplex and the vertices can easily be calculated even for large packings, see Section 6.2.2.1. It would be interesting to consider the behavior of the randomized LP algorithm for such a simplex polytope, and determine the probability that the weakest vertex, or at least a good approximation of the smallest δ_{vertex}, is found by the LP for a random load $\mathbf{B}$.

4.6 Algorithmic Details

In this section we outline in detail two algorithms to test for jamming in hard-sphere packings. The first one is applicable to ideal packings, while the second one deals with non-ideal packings. Although the core concept used in both is the randomized linear programming algorithms presented in Sections 4.3 and 4.4.3, the two differ in their goals and the way they process the results of the linear programming step: The first one attempts to give a binary classification of packings into jammed and not jammed, while the second tries to explore the extent of $\mathcal{J}_{\Delta\mathbf{R}}$ by trying to continuously displace the particles as much as possible, as discussed in Section 4.3.3.

4.6.1 Algorithm: Ideal Packings

We summarize the proposed algorithm to test for collective or strict jamming in ideal packings, applicable also to packings with very small interparticle gaps (Algorithm 13). This algorithm removes spheres which are not locally jammed. Once a rattling sphere is removed, this removes some contacts from the packing, which can make other spheres not locally jammed. Therefore an implementation in which neighbors of rattlers are recycled on a stack of spheres to be checked is needed. This algorithm also classifies packings which have an ideal jammed subpacking as jammed, even if they have some rattling particles or rattling clusters of particles.

Algorithm 13: Randomized Linear Programming Algorithm

1. If testing for collective jamming, fix the strain $\boldsymbol{\varepsilon} = 0$.
2. Choose a suitable gap tolerance δ, $\delta D \sim \Delta r_{\text{large}}$, in Eq. (4.5) and add all potential contacts $\{i, j\}$ between neighboring particles to the contact network.
3. If there are no spheres in the packing, declare the packing as not jammed and terminate.
4. Test for local jamming (rattlers):
   (a) Make a stack of all the spheres.
   (b) Remove the top (pop) sphere i from the stack and solve the LP (4.15) with $\mathbf{b}_i = \mathbf{A}_i\mathbf{e}$ using the simplex algorithm. If $\|\Delta\mathbf{r}_i\| \geq \Delta r_{\text{large}}$, remove the sphere from the packing, push all its neighbors not on the stack back on the stack, and remove all its contacts from the contact network.
   (c) Go back to step 4a if the stack is not empty.
   (d) Repeat step 3.
5. Choose a random load $\mathbf{B}$.
6. Solve the LP (4.10) along with constraint (4.17) (if testing for strict jamming).
7. Remove all spheres i displaced by the load from the packing, $\|\Delta\mathbf{r}_i\| \geq \Delta r_{\text{large}}$.
8. Repeat steps 3-4, reverse the direction of $\mathbf{B}$, $\mathbf{B} \leftarrow -\mathbf{B}$, and repeat steps 6-7.
9. If no spheres were displaced by either load, declare the (sub)packing jammed and terminate. Otherwise go back to step 3.

4.6.2 Algorithm: Nonideal Packings

When dealing with nonideal packings, one has to abandon the binary "jammed" versus "not jammed" binary classification. Instead we focus on trying to judge the *extent* of $\mathcal{J}_{\Delta\mathbf{R}}$ by trying to displace the spheres away from their current position by as much as possible. We first give the algorithm to do this in Algorithm 14 and then we discuss specific steps and the choices one can make in each step. Some illustrative results are given in Section 4.7.

We discuss the different steps of this algorithm separately in the following subsections. Note that the proposed algorithm is not as efficient as possible, mostly because not all linear programs need to be solved to full accuracy. Linear optimizers, and in particular interior-point algorithms, spend most of their effort in the final stages of the optimization, looking for the exact optimal vertex (or face) of the feasible polytope. Therefore, early termination is most desirable, and in future work we will develop specialized implementations that will replace step 5 with several Newton steps of a feasible interior-point algorithm. In a sense, the above algorithm resembles a Sequential Linear Programming (SLP) algorithm for finding equilibrium configurations of packings under applied loads. It remains to be explored whether including information about the curvature of the nonlinear impenetrability constraints, as is done in most modern nonlinear optimization algorithms, will be useful in light of the increased complexity of the linear algebra involved. For packings of nonspherical particles, such as packings of ellipsoids, including second-order information is necessary in order to find feasible directions of displacements. Numerous optimizations related to reuse of information in the iterative process and linear solvers as well as parallelization will also be investigated.

We stress that one cannot directly use off-the-shelf nonlinear optimization software to explore $\mathcal{J}_{\Delta\mathbf{R}}$, since feasibility must be strictly maintained throughout the process. Furthermore, efficiency also demands a specialized implementation. This is why we present in this work algorithms based on linear programming, which allows one to use any of the numerous linear programming libraries available today without the complexity of dealing with nonlinear programming algorithms.

4.6.2.1 Choosing the Gap Tolerance

First, we discuss the choice of the gap tolerance δ. The larger this tolerance, the more possible particle contacts we will add to the set of constraints, and thus the more computational effort we need. Further-

Algorithm 14: Sequential Random Loading Algorithm

1. If testing for collective jamming, fix the strain $\boldsymbol{\varepsilon} = 0$.
2. Choose a suitable gap tolerance δ, $\delta D \sim \Delta r_{\text{large}}$, in Eq. (4.5) and add all potential contacts $\{i, j\}$ between neighboring particles to the contact network.
3. Test for rattlers:
   For all spheres i, solve the LP (4.15) using the simplex algorithm and two randomly chosen loads $\mathbf{b}_i$ of opposite direction. If $\|\Delta\mathbf{r}_i\| \geq \Delta r_{\text{large}}$, mark the sphere as a rattler.
4. Choose a random load $\mathbf{B}$ and set $\mathbf{b}_i = 0$ for rattling particles.
5. Solve the LP (4.10) along with constraint (4.17) to obtain a linearized unjamming motion $\Delta\mathbf{R}$.
6. Find the largest scaling $\tau > 0$ for the displacements so that no spheres overlap for displacements from 0 to $\tau\Delta\mathbf{R}$ and also require that the volume of the unit cell does not increase, $\det[\mathbf{I} + \tau\boldsymbol{\varepsilon}] \leq 1$. Displace the spheres to a new configuration, $\mathbf{R} \leftarrow \mathbf{R} + \tau\Delta\mathbf{R}$, $\boldsymbol{\Lambda} \leftarrow (\mathbf{I} + \tau\boldsymbol{\varepsilon})\,\boldsymbol{\Lambda}$. Note that this changes the rigidity matrix $\mathbf{A}$ of the packing and requires updating the contact network.
7. If any particle was displaced by a significant amount, $\tau\|\Delta\mathbf{r}_i\| > \beta D$, go back to step 5. Also keep statistics of $\tau\|\Delta\mathbf{r}_i\|$ over all spheres, such as average $\overline{\|\Delta\mathbf{r}_i\|}$ and maximum value $\max\|\Delta\mathbf{r}_i\|$.
8. Optionally repeat step 3 and set $\mathbf{b}_i = 0$ for (new) rattling particles.
9. Reverse the direction of $\mathbf{B}$ and repeat steps 5-7.
10. If the average or maximal particle displacement exceed thresholds, declare packing as not "jammed" and terminate. Otherwise go back to step 4 until convinced packing is "jammed".

more, we are including more redundant and/or stricter-than-necessary linearized impenetrability constraints. Choosing a very small tolerance makes it hard to treat systems with moderately large interparticle gaps (say of the order of $\delta = 0.1D$), since crucial constraints which become relevant as soon as the magnitude of the displacements becomes comparable to δD are omitted. We have found values of $\delta \approx 0.1D - 1.0D$ reasonable, depending on the dimensionality and type of packing. The general rule is that the contacts of each sphere with *all* spheres in (only) its first coordination shell should be included, and of course physical intuition and close examination of the results are very helpful.

4.6.2.2 Testing for Rattlers

Unlike the case of ideal packings, where we permanently remove rattlers from the packing, here we simply avoid placing a load on the rattlers, but still consider them as part of the packing, as they may provide important constraints as the spheres displace. It is desirable not to place a load on rattlers because for some smaller gap tolerances δ, the contact network may not provide sufficient constraints to locally trap all particles. The particles that are not locally trapped will displace by very large distances under any nonzero load, leading to a very small scaling τ and very slow progress of the algorithm. Unfortunately, some linear programming solvers may return a large displacement for a rattler even if no load is applied on it, as most solvers initialize the variables independently of the user. In practice, step 3 of the algorithm only helps in cases when there is a small number of clear rattlers, as it enables one to use a smaller gap tolerance δ and thus reduce the size of the linear programs, and in such cases the first test for rattlers will already identify the troublesome particles. In other cases, one simply must use a sufficiently large δ.

4.6.2.3 Scaling the Displacements

We emphasized in Section 4.3.1 that any solution of the linearized impenetrability constraints is also a solution to the full nonlinear impenetrability constraints. However, there are several reasons why it is important to choose an appropriate scaling for the displacements τ. First, we do not include all pairs of particles in the constraints, and therefore any $\Delta\mathbf{R}$ for which some particle displaces by more than δD is not necessarily a feasible displacement, and may need to be scaled down appropriately. Furthermore, the linearized constraints

are significantly stricter than the nonlinear ones for larger displacements, and therefore it is often possible to scale up the displacements by a significant factor without violating feasibility.

Since our aim is to displace the particles as much as possible from their initial configuration, we choose the largest scaling factor possible. To find this scaling, one thinks of $\mathbf{\Delta R}$ as a vector of particle velocities and finds the time of the first interparticle collision τ. This can be found with exactly the same procedure as used to build collision schedules in the Lubachevsky-Stillinger packing algorithm [12]. For highest efficiency, the computational domain is partitioned into cells and only collisions between particles in neighboring cells are considered, along with transfers of the particles between the cells. The same partitioning is used when building the contact network of the packing after displacing the particles, though depending on the value of the gap tolerance δ more than just the neighboring cells might need to be searched. One should also ensure that the volume of the unit cell does not increase during the deformation of the lattice when testing for strict jamming.

4.6.2.4 Termination Criteria

We do not give detailed criteria on when to terminate the iterated linear programming in step 7, since these should really be adapted to a nonlinear feasible interior-point algorithm to be used in place of the linear optimizer. Typically $\beta \approx 0.01 - 0.1$. When none of the particles can be displaced further despite repeating step 5, the dual variables obtained by the LP solver will become (close to) the true interparticle forces that resist the load. A primal-dual nonlinear solver would also terminate at such a point and return the appropriate Lagrange multipliers. However, outside the ASD these forces are no longer unique [143], nor is it guaranteed that a packing that can support a random load $\mathbf{B}$ and $-\mathbf{B}$ can support all loads. Therefore, we need to use several random loads. We do not have estimates or bounds on how many loads need to be used, however, experience has shown that only a few $(3 - 5)$ loads are sufficient to find large displacements if such displacements exist.

4.6.2.5 Interpreting the Results

Processing the results of the above algorithm is somewhat of an art. However, by observing the statistics of the multiple particle displacements, and especially by visualizing the path traversed by the particles during the loading, one can get a sense of the character of $\mathcal{J}_{\mathbf{\Delta R}}$. Particularly useful is observing the average magnitude of the particle displacements, and we use this metric in reporting some results for disordered computer-generated packings in Section 4.7.2. It may also be useful to observe the distribution of displacement magnitudes among the particles.

In general, it is best to first try the algorithm of this section, and then use a visualization tool (like our VRML animations) or a histogram of the magnitudes of the particle displacements to judge whether there appears to be a jammed subpacking (to within a tight tolerance), or whether all particles seem to be able to displace significantly. If the former is the case, then using the algorithm of Section 4.6.1 one can identify such a jammed subpacking if it exists within the tolerances used. This is illustrated in Fig. 4.8, where we show the results of applying the algorithm to test for *collective* jamming in ideal packings to an equimolar bidisperse disk packing of 250 disks ($\varphi = 0.846$) . A gap tolerance of $\delta = 0.01$ was used, and all disks that displaced by more than $\Delta r_{\text{large}} = 10^{-3}$ were removed to leave a jammed subpacking of 232 disks, for which the average displacement during the test was $\overline{\|\Delta \mathbf{r}_i\|}/D \approx 7 \cdot 10^{-7}$ and the maximal was $\max_i \|\Delta \mathbf{r}_i\| \approx 2 \cdot 10^{-5}$, indicating a high numerical accuracy in the packing algorithm (about $20,000$ collisions per particle were used). If the rattling particles or rattling clusters of particles were not removed, the displacements observed would have been higher, as, for example, in Table 4.2. On the other hand, if overly strict tolerances were chosen in the algorithm of section 4.6.1 (for example, $\Delta r_{\text{large}} = 10^{-4}$), then no jammed subpacking would have been found. With reasonably tight tolerances, *there is no strictly jammed subpacking* of this packing. Note that it may be possible to remove some of the disks from the collectively jammed subpacking and still maintain the jamming property.

4.6.3 Numerical Implementation

We have developed a practical numerical implementation of the randomized LP algorithms using a primal-dual interior-point linear optimizer. We have applied the algorithms to test for the different jamming

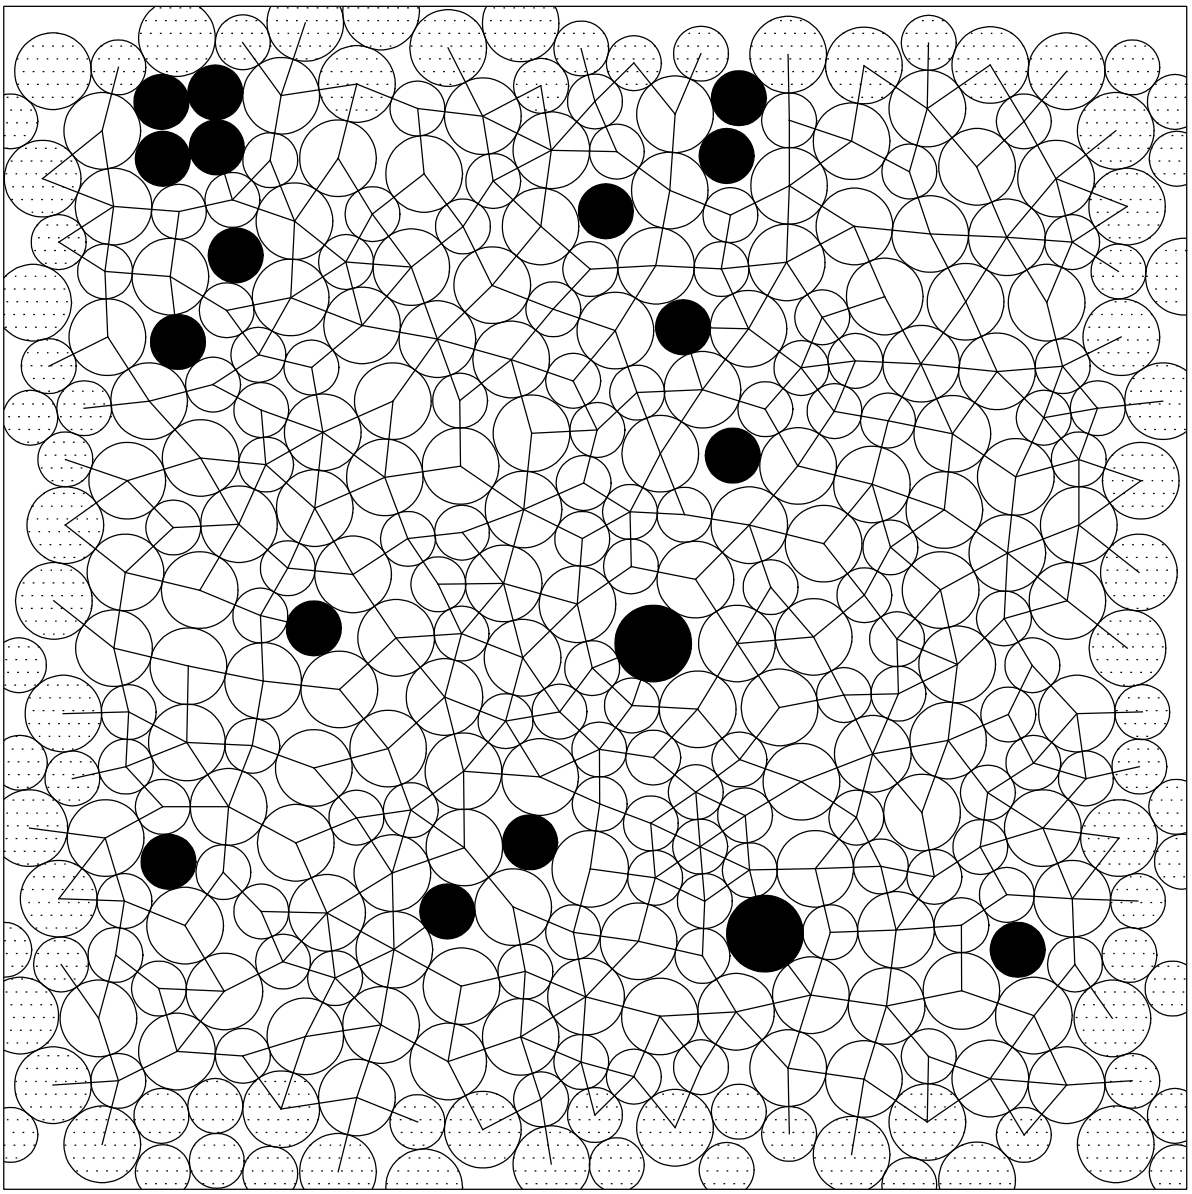

Figure 4.8: *Results from the algorithm of Section* 4.6.1 (Algorithm 13). The algorithm identified a collectively jammed subpacking, after removing the rattlers (black disks). The dotted disks represent periodic images.

categories in practice and verified their utility and efficiency. Illustrations of results obtained using our implementations are given throughout, and results from the application of the algorithm of Section 4.6.2 to large computer-generated monodisperse and bidisperse disk and sphere packings are given in Ref. [72].

We have implemented an efficient numerical solution of (4.10) using the primal-dual interior-point algorithm in the LOQO optimization library [153]. Both Fortran 95 codes which directly invoke the LOQO library, and Algebraic Modeling Programming Language (AMPL) models have been developed, along with VRML visualization tools. The AMPL models are particularly simple to use and modify, and are available on our website [135]. We wish to emphasize that by using primal-dual interior point algorithms we automatically get both forces and displacements using the same implementation. For example, both LOQO and PCx (see Ref. [153]) return both primal and dual solutions to the user. Primal-dual interior-point algorithms are very well suited for problems of this type. Nonetheless, for three-dimensional problems the available high-quality implementations of interior-point algorithms (such as Ref. [153]) are based on direct linear solvers are too memory-demanding and inefficient. Tuned implementations based on conjugate-gradient iterative solvers are needed.

Testing for strict jamming typically takes more time, by as much as 25%, since additional denser rows/columns are included in the rigidity matrix, and this is more pronounced in three dimensions where more of the contacts are on the boundary. The exact way the strain and the associated constraints are handled makes a difference in this case. We emphasize that for three-dimensional packings the sparse factorization linear solver in LOQO is not the best choice, so much smaller running times are possible with specialized implementations. The running time of the linear solver depends non-trivially on both the number of spheres *and* the number of contacts in the contact network. The number of contacts is very sensitive to the choice of the gap tolerance δ, which we usually decreased as the packing size increased (and thus the average displacements decreased). Therefore, the running times shown shortly should not be taken as a measure of the scaling of the LP solver computational effort with the number of spheres, but rather as typical runtimes for some representative packing sizes.

4.7 Results

In this section we apply the LP algorithms to representative ordered (crystalline) and disordered (amorphous) packings of hard disks and spheres, and summarize our observations.

4.7.1 Periodic Lattice Packings

Table 1 in Ref. [15] gives a classification of some common simple lattice packings into jamming categories for hard-wall boundary conditions. Table 4.1 modifies this for periodic boundary conditions. The results in principle will depend on the choice of unit cell, so the terminology "lattice XXX is YYY jammed" is used loosely here. We illustrate some unjamming motions for lattice disk packings in Figs. 4.9 and 4.10.

We point out for the curious that the triangular lattice is not the only strictly jammed ordered disk packing; two other examples are shown in Fig. 4.11. We show below that one can remove at most one quarter of the disks from a triangular lattice packing and still maintain strict jamming. Using the Lubachevsky-Stillinger packing algorithm for small packings, we recently found a new family of strictly jamming packings obtained by reinforcing with triangular regions a particular tiling of the plane with three congruent pentagons. An example is shown in Fig. 4.11.

4.7.1.1 Low-Density Strictly Jammed Triangular Packings

We give a simple argument here proving that the reinforced Kagomé lattice (see Fig. 4.11) is the lowest density strictly jammed subpacking of the triangular lattice disk packing. We start with a triangular packing and remove disks to form vacancies, preserving the strict jamming property. Observe that one cannot have any compact divacancies, i.e., neighbouring vacancies, since these unjam the packing. Therefore, each disk removal removes *exactly* 6 contacts from the contact network. At a minimum, we need at least 2 contacts per particle in a (collectively or) strictly jammed large packing, since the number of impenetrability constraints must be at least the number of degrees of freedom. Therefore, if we start with $N \gg 1$ disks, and remove n

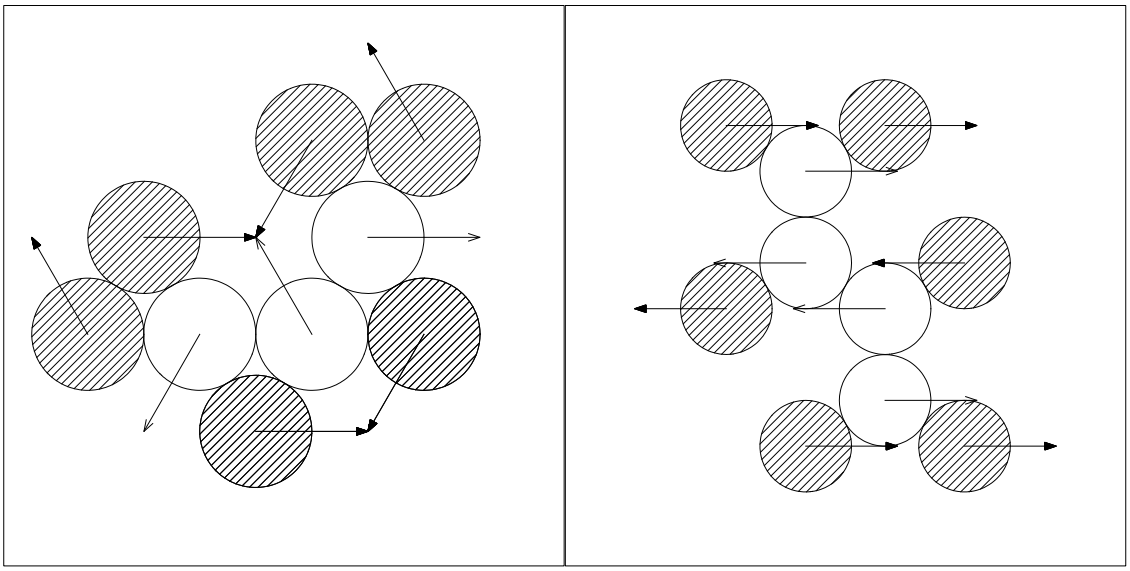

Figure 4.9: *Simple collective mechanisms in the Kagomé and honeycomb lattices*, respectively. These lattices are not collectively jammed with periodic boundary conditions, as the sample unjamming motions for the Kagomé (left) and for the honeycomb (right) packings shown here illustrate. The shaded disks represent periodic images.

Lattice	φ	Z	N_s	L	C	S
Honeycomb	0.605	3	4	Y	N	N
Kagomé	0.680	6	3	Y	N	N
Square	0.785	4	2	Y	N	N
Triangular	0.907	6	1	Y	Y	Y
Diamond	0.340	4	4	Y	N	N
SC	0.524	6	2	Y	N	N
BCC	0.680	8	2	Y	N	N
FCC	0.741	12	1	Y	Y	Y
HCP	0.741	12	2	Y	Y	Y

Table 4.1: *Classification of some simple lattices into jamming categories for periodic boundary conditions.* We give the packing (i.e., covering) fraction φ (to three decimal places), the coordination number Z, and the number of disks/spheres N_s per unit cell, as well as an assessment of whether the lattice is locally (L), collectively (C) or strictly (S) jammed (Y is jammed, N is not jammed). We chose the smallest unit cells for which an unjamming motion exists (illustrated on our webpage [135]), if there is one.

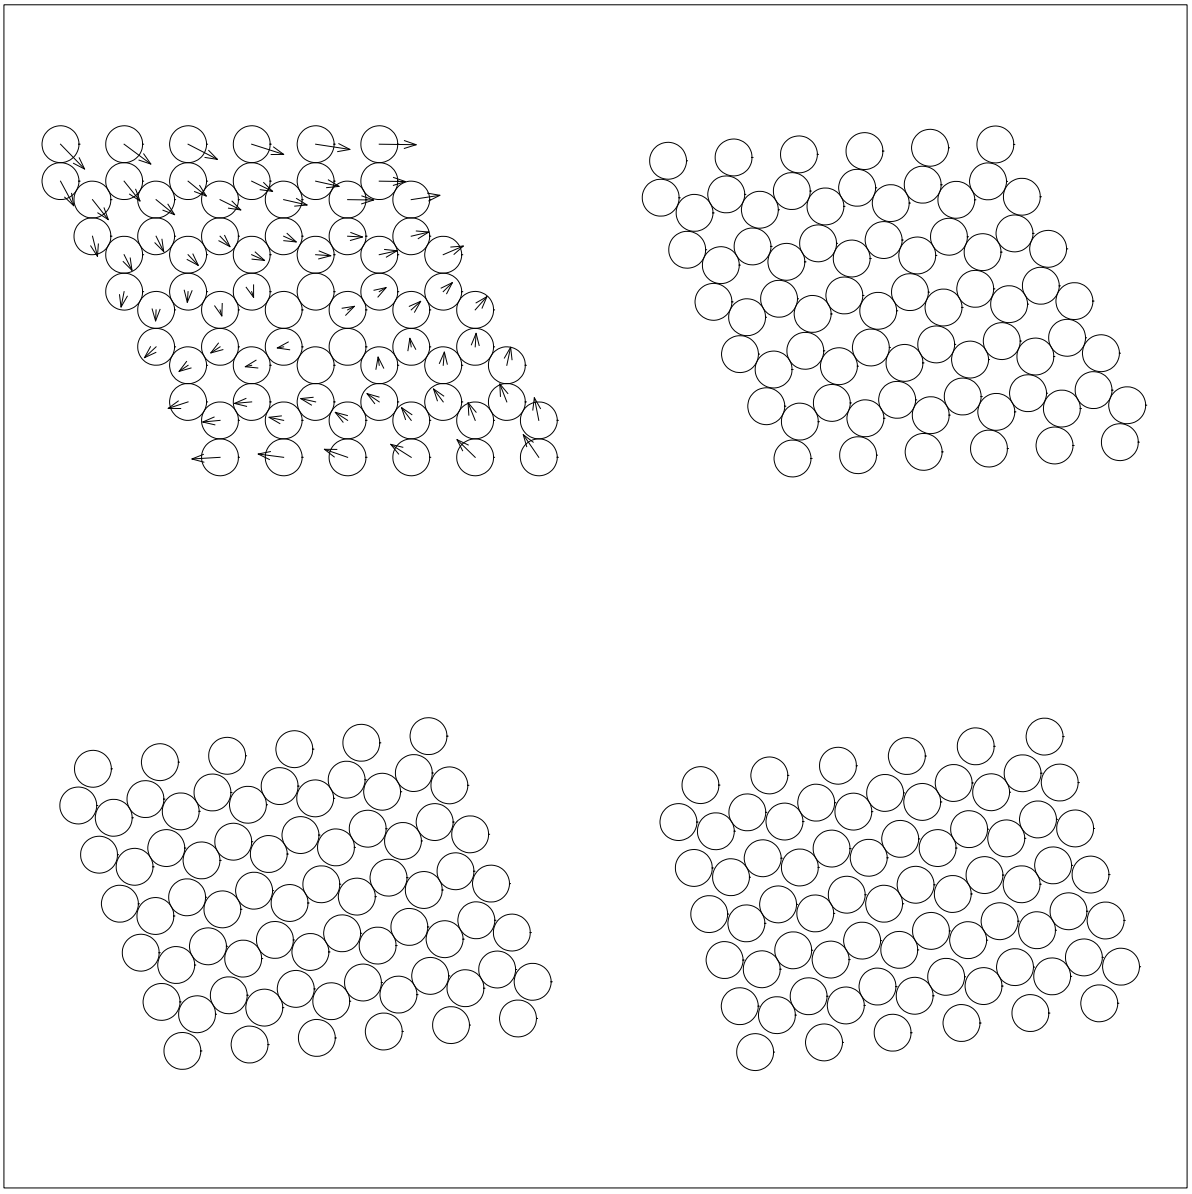

Figure 4.10: *Shearing the honeycomb lattice.* The honeycomb lattice is not strictly (or collectively) jammed, and an example of a lattice deformation, replicated on several unit cells to illustrate the shear character of the strain $\boldsymbol{\varepsilon} = (\boldsymbol{\Delta\Lambda})\boldsymbol{\Lambda}^{-1}$. Note that only three (original) spheres are involved in the actual calculation of this unjamming motion, the rest are image spheres.

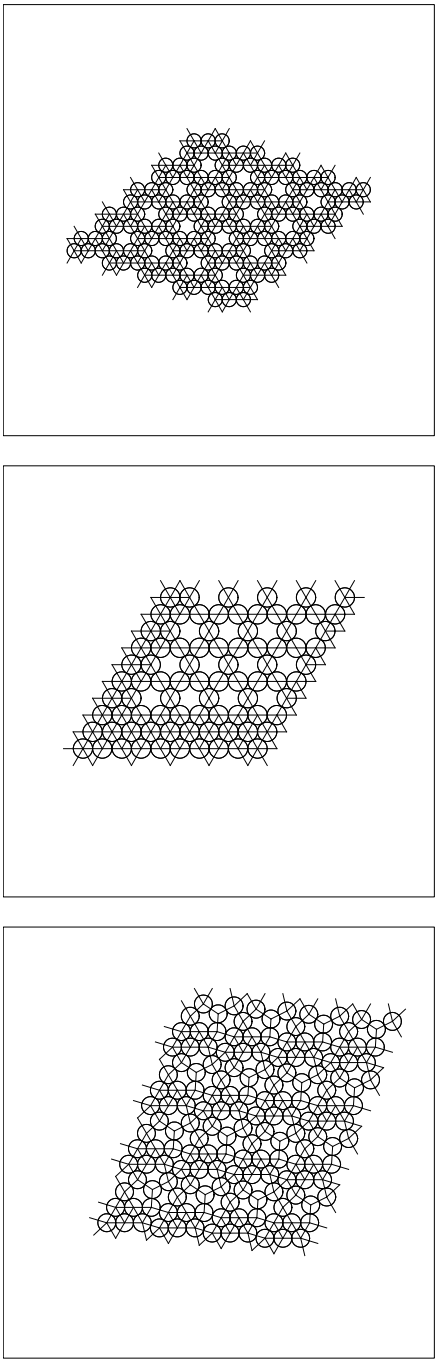

Figure 4.11: *Examples of strictly jammed lattices* in two dimensions (from Ref. [71]). The 6/7th lattice (packing number 2 in Ref. [142] and the last packing in Ref. [59]), top, is obtained by removing every 7th disk from the triangular lattice. The reinforced Kagomé lattice, middle, is obtained by adding an extra "row" and "column" of disks to the Kagomé lattice and thus has the same density in the thermodynamic limit, namely, it has every 4th disk removed from the triangular packing (see also Ref. [59]). The pentagonal packing shown at the bottom is obtained from a particular tiling of the plane with three rotated congruent pentagons, and is just one member of a whole family of strictly jammed packings.

disks, we must have that:

$$3N - 6n \geq 2(N - n) \Rightarrow n \leq \frac{N}{4} \tag{4.26}$$

that is, we can at most remove one quarter of the disks. Since the reinforced Kagome lattice is strictly jammed (we do not prove this here) and asymptotically realizes this bound, the bound (4.26) is tight.

4.7.2 Results for Disordered Hard Sphere and Disk Packings

The linear programming algorithms we developed were specifically designed for disordered packings. We have used them to analyze jamming in disordered packings produced by a variety of packing algorithms, namely the Lubachevsky-Stillinger packing algorithm [12], an energy minimization algorithm as presented in Ref. [31], as well as the Zinchenko packing algorithm [28].

4.7.2.1 Packing Algorithms

We produced most packings using the Lubachevsky-Stillinger compression algorithm [12] with periodic boundary conditions. We also obtained sample monodisperse sphere and bidisperse disk packings from the authors of Ref. [31]. These packings are not of perfectly hard spheres, but rather soft spheres interacting via repulsive potentials when there is overlap between the cores of diameter D. They use energy minimization for harmonic and Hertzian potentials, descending to an energy minimum using the conjugate gradient algorithm from a random initial configuration (i.e., a rapid quench from $T = \infty$ to $T = 0$). The packings we analyzed were just above the "jamming threshold" density ϕ_c, meaning that there was only very small (less than $10^{-5}D$) overlap between the outer cores. We therefore simply scaled the sizes of the particles by a factor very close to unity to obtain overlap-free hard-sphere packings. Since the jamming threshold densities found in Ref. [31] were very close to the final densities produced by the Lubachevsky-Stillinger algorithm (with reasonably large compression rates), we expected these packings to behave very similarly, and have confirmed this with computational tests. Therefore, here we focus on and present the results for the Lubachevsky-Stillinger packings. Finally, we also had available disordered three-dimensional packings produced with the contact network building Zinchenko packing algorithm [28], and confirmed that they behaved like the packings produced by the other algorithms. Unfortunately, we do not know of a two-dimensional implementation of this algorithm.

All of these algorithms seem to produce collectively jammed packings in both two and three dimensions, excluding rattlers and allowing for appropriate numerical tolerances. This can be proved rigorously for the Zinchenko packing algorithm, and under certain additional assumptions also for the energy minimization algorithm. In principle, only locally jammed configurations are possible final states for the Lubachevsky-Stillinger algorithm since they give infinite collision rates, however, we believe that local configurations are unstable attractors for this algorithm. On the other hand, none of these algorithms produces truly strictly jammed packings *a priori*. Indeed, the packings that that we tested were never truly strictly jammed. This is not surprising because none of them incorporates deformations of the periodic lattice, but rather, they all use a fixed (typically cubical) unit cell. It is not hard to incorporate boundary deformations into these algorithms, and we are presently working on such extensions. In particular, the Lubachevsky-Stillinger algorithm can easily incorporate a deforming lattice in the spirit of Parinello-Rahman molecular dynamics [107]. We have in fact implemented such an extended Lubachevsky-Stillinger algorithm (described in Section 3.2.4.2) and used it to produce *a priori* strictly jammed packings. In packing algorithms based on energy minimization, as in Ref. [31], one need only include the strain as part of the degrees of freedom in order to allow relaxation of the lattice and produce strictly jammed packings. The same is true of the Zinchenko packing algorithm.

Both monodisperse and bidisperse packings were studied. The main reason for including bidisperse packings in this preliminary study is that monodisperse disk packings crystallize easily, forming large ordered almost-triangular domains (grains) with high packing fraction $\varphi \approx 0.88$. This is because in two dimensions the locally densest configuration coincides with the globally densest triangular lattice, unlike in three dimensions, where the locally optimal (tetrahedral) configuration cannot tile three-dimensional space [6]. It is only by introducing polydispersity that one can produce disk packings with no apparent (or little) short-range order (i.e., amorphous), as can be determined by, for example, bond-orientational order metrics,[6] and in

particular, the *local* Q_6 order metric. We used an equimolar mixture of disks with diameter ratio of 1.4 as done in Ref. [31]. For amorphous monodisperse three-dimensional packings the typical packing fraction is around $\varphi \approx 0.64$, and such a packing is shown in Fig. 4.12. For the aforementioned amorphous binary disk packings $\varphi \approx 0.84$, and such a packing is illustrated in Fig. 4.8.

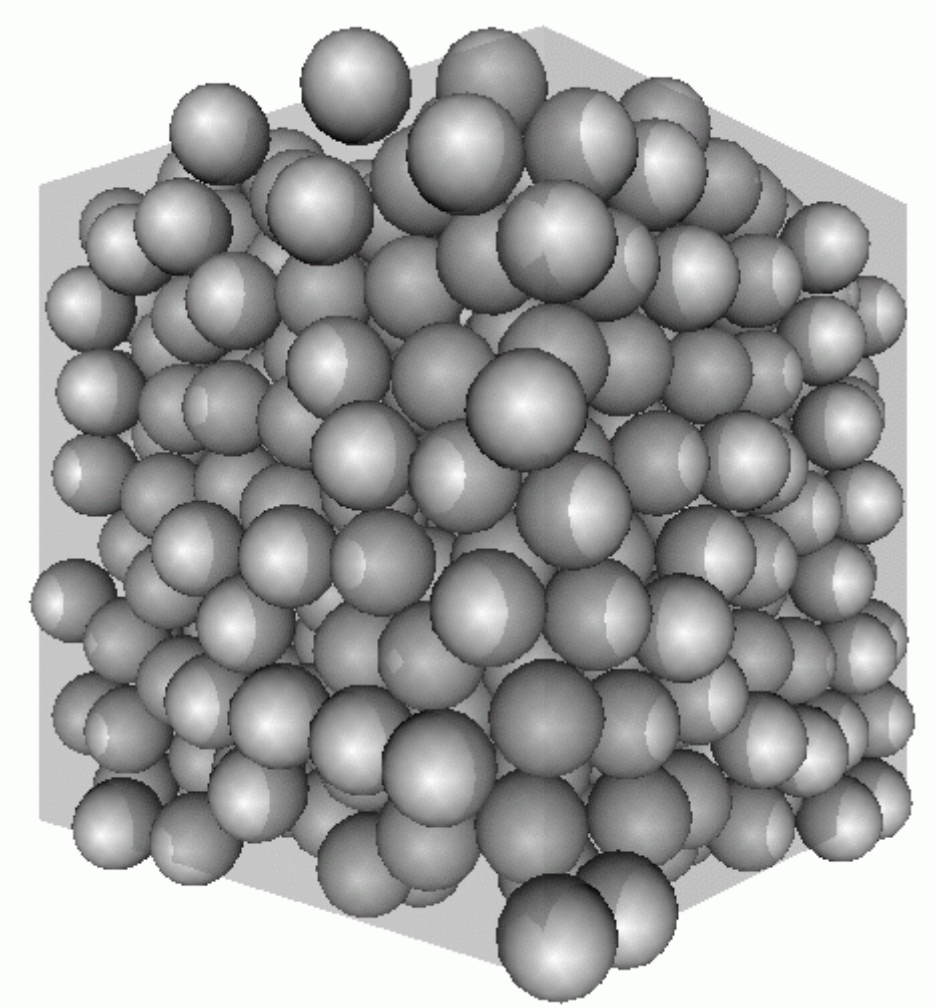

Figure 4.12: *Virtually strictly jammed sphere packing.* This random packing of 500 spheres with density $\varphi = 0.64$ was produced using the (original) Lubachevsky-Stillinger algorithm and it is collectively jammed and practically strictly jammed. The (cubical) unit cell is also shown.

In a truly disordered (generic) packing, it is expected that the average number of interparticle contacts per particle (coordination number) will be $Z = 2d$, i.e., that the packing will be isostatic. Thus, it is expected that $Z = 4$ in two dimensions. However, collectively jammed monodisperse disks packings are rather dense ($\varphi \approx 0.86 - 0.88$) and crystalline and they have $Z > 4$ (This should be compared to $Z = 6$ for the triangular crystal). Disordered bidisperse disk packings do have $Z \approx 4$, and similarly in three dimensions monodisperse packings have $Z \approx 6$, consistent with an assumption of generic character. However, the exact number one gets depends rather sensitively on the criterion for assigning contacts and on whether rattling particles are excluded or not. It is important to note that a large packing must have $Z \geq 2d$ in order to be collectively or strictly jammed, and $Z \geq d + 1$ to be locally jammed. We also point out that our algorithm to test for jamming does not depend sensitively on the criterion for selecting contacts.

4.7.2.2 Methodology

Although most of the packings we analyzed had small interparticle gaps and can also be studied within the framework of ideal packings and classified as jammed or not jammed, we instead consider them nonideal and explicitly deal with the interparticle gaps. We wish to stress that the results below are *not averages* over many packings with the same number of spheres/disks, but rather they are results for *particular* packings produced by the Lubachevsky-Stillinger algorithm. These packings seemed to be typical of the types of packings produced by the algorithm under a relatively wide range of expansion rates and packing sizes. We therefore believe that the numbers presented here serve well as a semi-qualitative illustration of the behavior of random disk and sphere packings commonly used in many computational studies. The primary reason we do not give averaged results this is that detailed average results should be given only once it is determined what quantitative metric of jamming is physically appropriate (which is likely to be different for different types of packings and different applications), and results should also be correlated with more characteristics of the packings (i.e., not just the covering fraction) and to various relevant parameters of the algorithm used to generate the packing [56].

As a quantitative measure of jamming in these packings, we report the average particle displacement $\overline{\|\Delta \mathbf{r}_i\|}$ achieved during random loading. This choice is not ideal, and attaching a physical picture to the numbers is difficult. For example, rattlers often contribute most to the average displacement for packings which might be "more jammed" if the rattlers are removed. Moreover, although an entry in Table 4.4 below might say that the average displacement for a monodisperse disk packing was only 10% of the particle size, the character of the particle motion might be such that very significant rearrangements happen in the packing because grain boundaries move (see Fig. 4.14), and this has to be seen to be appreciated. We will share our VRML visualizations with interested readers, and many examples are provided on our webpage [135].

Another statistic we report is the time (in seconds) spent by the AMPL implementation (with some Fortran) of the testing algorithm on a typical personal computer[14]. Since most of the computational time is spent in LOQO, similar running times are typical of the Fortran codes as well. For each packing, we applied three different random loads (with opposite orientations), and for each load we successively solved three linear programs (so a total of 18 linear programs for each packing). The running times below should not be taken as a measure of the scaling of the LP solver computational effort with the number of spheres, but rather as typical runtimes for some representative packing sizes. This is because the computational effort depends nontrivially on many of the parameters in the algorithm, and on the exact implementation.

4.7.2.3 Results

Qualitatively different results were observed for the amorphous monodisperse sphere packings and binary disk packings, and the polycrystalline monodisperse disk packings.

For the amorphous packings, we give results for bidisperse packings in two dimensions in Table 4.2, and for monodisperse packings in three dimensions in Table 4.3, with similar trends. In general, these packings *were* collectively jammed, in the sense that only small (average) displacements of the particles are possible. The small feasible displacements are mostly due to rattlers and/or early termination of the packing algorithm and we believe that any true final Lubachevsky-Stillinger packing with infinite collision rate will in fact have an *ideal* collectively jammed subpacking (similarly for the other packing algorithms). The packings *were not* strictly jammed for small system sizes, however, the magnitude of the feasible displacements decreased as the packings became larger, and therefore *large amorphous packings were virtually collectively and strictly jammed*. This can be understood by thinking of the distinction between collective and strict jamming as a boundary effect: As the packings become larger the boundary effects diminish. Therefore, even though none of the packing algorithms is meant to produce strictly jammed packings *a priori*, they do so for large amorphous packings.

Importantly, very different results were observed for monodisperse disk packings, which are invariantly nearly triangular (i.e., crystalline). We wish to point out that crystallization into a triangular lattice poses a convergence obstacle for the Lubachevsky-Stillinger algorithm since near triangular regions have very high collision rates even when the disks diameters are not at their maximal value. Therefore it was only for monodisperse disk packings that some of the final packings were not collectively jammed (large particle

[14] More precisely, calculations were performed on an 1666MHz AMD Athlon PC running Linux.

N	ϕ	t (s) coll.	$\overline{\|\Delta\mathbf{r}_i\|}/D_i$ coll.	$\overline{\|\Delta\mathbf{r}_i\|}/D_i$ strict
50	0.845	2.1	0.010	0.060
100	0.842	6.4	0.0034	0.011
250	0.846	21	0.0037	0.0053
500	0.847	72	0.0016	0.0067
750	0.849	88	0.0022	0.012
1000	0.849	130	0.0016	0.018
1500	0.848	247	0.0016	0.020
2500	0.849	248	0.0039	0.010

Table 4.2: Results of Algorithm 14 for equimolar binary disk packings of diameter ratio 1.4. The first column shows the total number of particles N, the second the packing fraction, the third the running time for the AMPL model that tests for collective jamming, and the last two columns show the average particle displacement during collective (i.e., with a fixed lattice) and strict jamming (i.e., with a deforming lattice) testing. Notice that the displacements are significantly larger for the strict jamming test, especially for small packings. The analogous table for three dimensions, Table 4.3, shows similar behavior but significantly larger computational times due to the inefficiency of the direct linear solver in LOQO for three-dimensional contact networks.

N	ϕ	t (s) coll.	t (s) strict	$\overline{\|\Delta\mathbf{r}_i\|}/D$ coll.	$\overline{\|\Delta\mathbf{r}_i\|}/D$ strict
50	0.628	23	29	0.0012	0.12
100	0.644	53	76	0.00043	0.15
250	0.636	164	210	0.0021	0.031
500	0.641	480	597	0.0037	0.014
750	0.641	900	1017	0.0015	0.0035
1000	0.642	1822	1866	0.011	0.013

Table 4.3: Results for monodisperse sphere packings. The columns are as in Table 4.2, and here we show the running times for both the testing for collective and strict jamming.

rearrangements were possible near grain boundaries). Most packings were however collectively jammed just as for amorphous packings and we present results for these in Table 4.4. By using certain tricks in the Lubachevsky-Stillinger algorithm, such as collections of frozen particles or very large expansion rates, one can obtain apparently "jammed" amorphous monodisperse disk packings near a packing fraction $\varphi \approx 0.83$. However, due to numerical instabilities or the presence of an artificial boundary of fixed disks, these packings *were not* collectively jammed, as illustrated in Fig. 4.13. One of the important observations is that *none of the large Lubachevsky-Stillinger monodisperse disk packings were strictly jammed*. In fact, typical grain boundaries are very fragile under shear, and so even for the large packings significant rearrangements of the grain boundaries are feasible, as illustrated in Fig. 4.14.

N	ϕ	t (s) coll.	$\overline{\|\Delta\mathbf{r}_i\|}/D$ coll.	$\overline{\|\Delta\mathbf{r}_i\|}/D$ strict
50	0.832	2.9	0.0022	0.39
100	0.863	8.9	$5.4{\cdot}10^{-8}$	0.18
250	0.886	21	0.0014	0.86
500	0.891	78	$6.7{\cdot}10^{-5}$	0.16
750	0.887	103	0.0040	0.26
1000	0.882	153	0.0017	0.23

Table 4.4: Results for monodisperse disk packings. The columns are as in Table 4.2. Notice the very large displacements during the test for strict jamming, even for large packings, as well as the high packing densities for larger packings.

It is also important to verify that any packing algorithm claimed to produce jammed packings can indeed

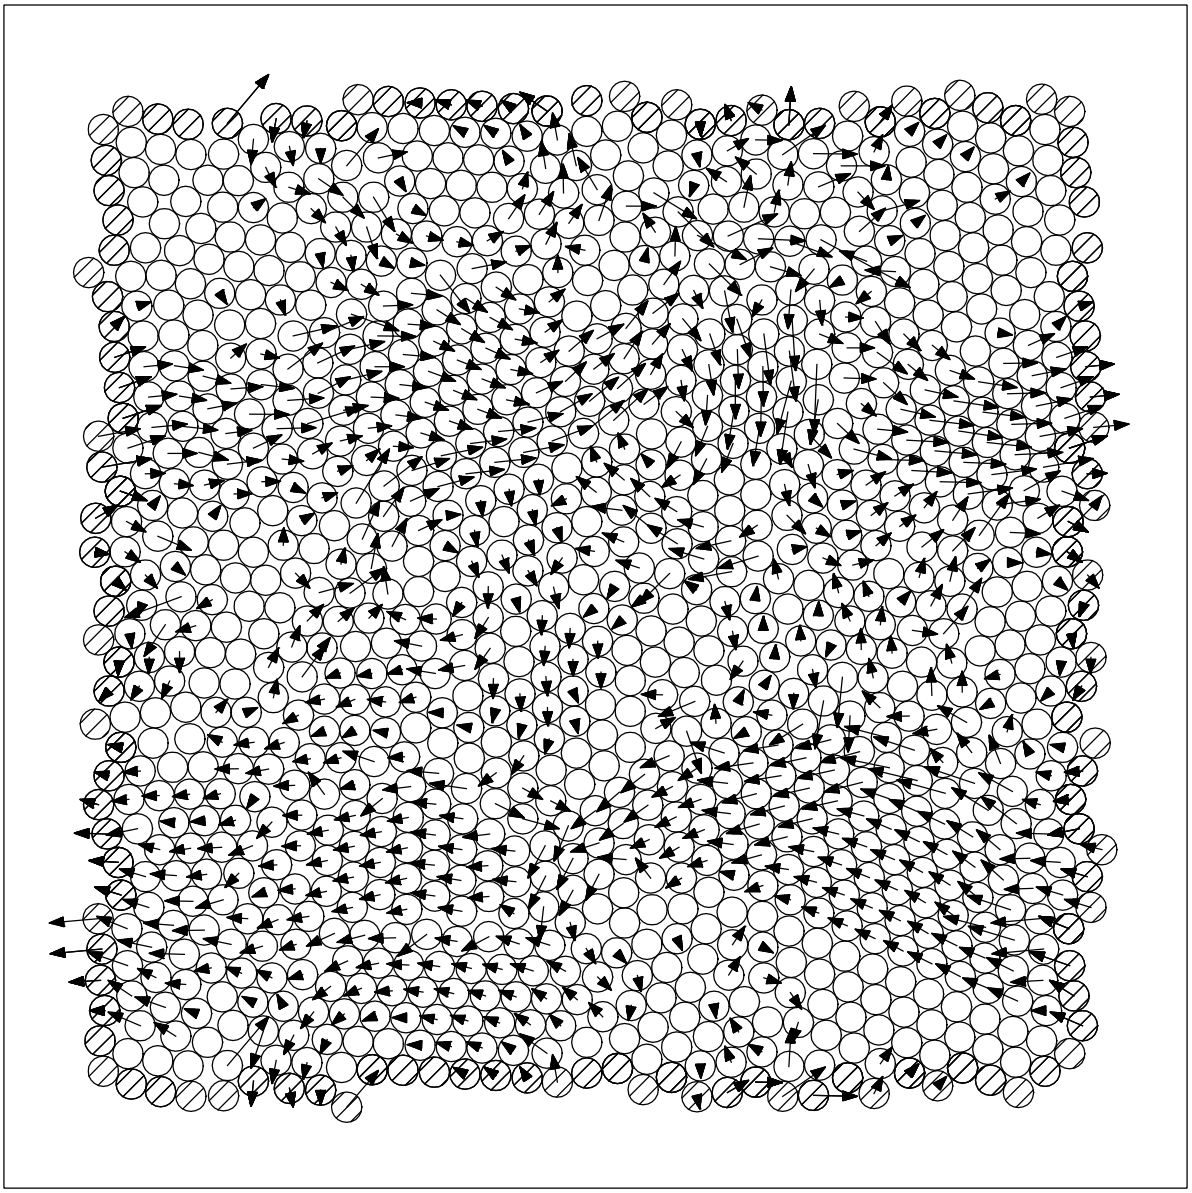

Figure 4.13: *Locally jammed disk packing.* A random packing ($\varphi = 0.82$) of 1000 disks that is *not collectively jammed*, and a representative periodic unjamming motion. More insightful animations can be found at the webpage [135].

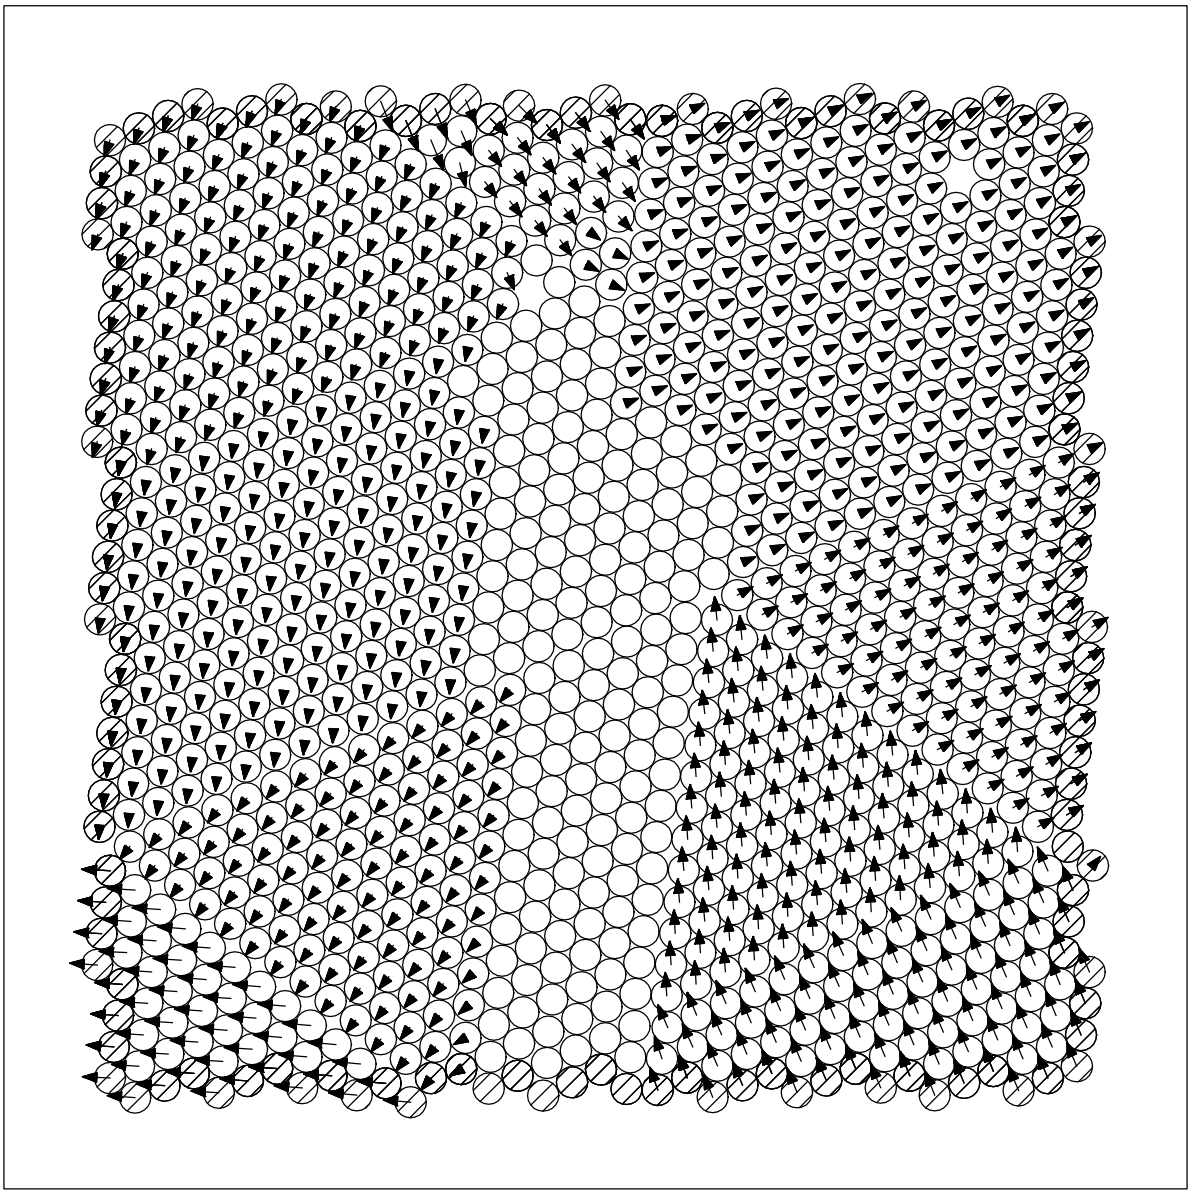

Figure 4.14: *Collectively jammed disk packing.* A *dense* ($\varphi = 0.89$) random packing of 1000 disks that is *collectively jammed but not strictly jammed*, and a representative unjamming motion. One can see the grains gliding over each grain boundary due to the shear, bringing this packing closer to a triangular lattice.

produce jammed ideal packings, in the sense that all tolerances in the test for jamming can be tightened progressively as the numerical accuracy is increased and the convergence criteria in the packing algorithm are tightened. We demonstrate this for collective jamming in monodisperse sphere packings in Table 4.5. The corresponding results for strict jamming, given in Table 4.6, illustrate that the (traditional) Lubachevsky-Stillinger packings do not have a strictly jammed ideal subpacking, but are practically strictly jammed for large system sizes. This is unlike monodisperse disk packings, which are far from being strictly jammed, as illustrated in Table 4.7.

N / $N_{\rm coll}(10^3)$	1	5	10	25
50	0.041	0.015	.0018	$4.9{\cdot}10^{-10}$
100	0.036	0.016	0.0011	0.00014
250	0.050	0.023	0.0015	0.00036
500	0.047	0.024	0.0028	0.0014
750	0.046	0.019	0.0030	0.0011
1000	0.052	0.020	0.0025	0.00067

Table 4.5: The average particle displacement $\overline{\|\Delta\mathbf{r}_i\|}/D$ during the test for collective jamming is shown for a series of sphere packings produced by the (original) Lubachevsky-Stillinger algorithm. From top to bottom the packing size N increases, and from left to right the number of collisions per particle $N_{\rm coll}$(in thousands) increases (and thus the density also slowly increases). No special handling of rattlers was employed. It is easily observed that the packings uniformly become "more jammed" as the packing algorithm is run longer (though rattlers may continue to give a finite contribution to the observed displacements). Similar behavior is expected of any algorithm which in the limit of infinite numerical precision produces packings with a collectively jammed subpacking.

N / $N_{coll}(10^3)$	1	5	10	25
50	0.083	0.057	0.059	0.051
100	0.066	0.042	0.023	0.026
250	0.052	0.027	0.010	0.0097
500	0.056	0.024	0.012	0.010
750	0.048	0.027	0.014	0.014
1000	0.060	0.025	0.0040	0.0021

Table 4.6: This table is the analog of Table 4.5 but for strict jamming. In this case it is seen that the average displacements do not converge uniformly toward zero, an indication that the packings do not have a strictly jammed ideal subpacking (similar results are observed for amorphous binary disk packings). However, the average displacements are quite small for large packings (this is even more pronounced for the binary disk packings).

$N_{coll}(10^3)$	1	5	10	25
Collective	0.12	0.007	0.00050	$1.7{\cdot}10^{-5}$
Strict	0.45	0.24	0.12	0.12

Table 4.7: Just as an illustration, this short table shows results for a two dimensional disk packing with 250 disks, corresponding to the results presented for amorphous sphere packings in Tables 4.5 and 4.6. It is seen that although the packing has an ideal collectively jammed subpacking, it is clearly far from being strictly jammed, as typical for monodisperse disk packings produced by the Lubachevsky-Stillinger packing algorithm with a fixed lattice.

We have implemented an extension of the Lubachevsky-Stillinger packing algorithm in which the lattice deforms during the molecular dynamics run [11], as dictated by the collisional (contact) stress induced by the particle collective. Details of the algorithm and the packings it produces will be given elsewhere. For relatively

small numbers of particles, this algorithm typically produces truly strictly jammed packings, and for these packings $\overline{\|\Delta \mathbf{r}_i\|}$ is similar for both collective and strict jamming. The algorithm produces similar amorphous packings (in packing fraction and disorder) to the original Lubachevsky-Stillinger algorithm, however, for monodisperse disks it frequently terminates with completely crystal packings, and also produces complete triangular lattices with special types of defects, such as monovacancies and peculiar "dislocation cores". One such strictly jammed disk packing is shown in Fig. 4.15. Investigation of these strictly jammed disk packings as well as extensions of other packing algorithms to allow for deforming boundaries should be carried out in the future. We note in passing that we also used the extended Lubachevsky-Stillinger algorithm to try the shrink-and-bump heuristic [71] to test for strict jamming by also allowing the lattice to deform while the particles bump around. This seemed to detect disordered packings which are not strictly jammed, however, the test is significantly slower than the linear programming algorithm and is also very heuristic and much less reliable.

4.8 Conclusions

In this chapter we proposed, implemented, and tested a practical algorithm for verifying jamming categories in finite sphere packings based on linear programming. We demonstrated its simplicity and utility, and presented some representative results for ordered lattices and random packings. Interestingly, the large computer-generated monodisperse random packings that we tested were virtually strictly jammed in three dimensions, but *not* in two dimensions.

Our results have important implications for the classification of random disk and sphere packings and suggest a number of interesting avenues of inquiry for future investigations. Random disk packings are less well-understood than sphere packings. The tendency of disk packings to "crystallize" (to form ordered, locally dense domains) at sufficiently high densities is well established. For example, Quickenden and Tan experimentally estimated the packing fraction of the "random close packed" (RCP) state to be $\varphi \approx 0.83$ and found that the packing fraction could be further increased until the maximum value of $\varphi = 0.906$ is achieved for the triangular lattice packing [154]. By contrast, random sphere packings at φ in the range $0.63 - 0.66$ cannot be further densified.

Our recent understanding of the ill-defined nature of random close packing and of jamming categories raises serious questions about previous two-dimensional studies, particularly the stability of such packings. Our present study suggests that disordered random disk packings are not collectively jammed at $\varphi \approx 0.83$; at best they may be locally jammed. This brings into question the previous widespread belief that the two-dimensional analog of the RCP sphere-packing state has density about $\varphi \approx 0.82 - 0.83$ [41]. Collectively jammed disk packings seem to have significantly higher densities $\varphi \approx 0.88$ and consist of large triangular grains, but even at such high densities they are not strictly jammed. An interesting question is whether the grain size becomes small compared to the system size for large collectively jammed disk packings, or whether the appearance of grain boundaries is in fact a finite-size boundary effect. It may be that the preponderance of collectively and strictly jammed large disk packings are very crystalline, with a distribution of the local bond-orientational parameter Q_6 (see Ref. [6]) highly peaked around some relatively large value. Furthermore, it is important to ascertain if the strong distinction between only collectively and strictly jammed disk packings persists in the limit of very large packings. Careful investigations of very large collectively and strictly jammed disk packings produced with a variety of packing algorithms are still required to answer these questions.

The old concept of the RCP state incorrectly did not account for the jamming category of the packing. Previous attempts to estimate the packing fraction of the "random loose" state [139] are even more problematic, given that this term is even less well-defined than the RCP state [155, 156]. Furthermore, as our investigations of disk packings show, the "stability" of packings cannot be judged based solely on local criteria, as suggested in Ref. [141] for sphere packings, and using such local criteria in estimating mean coordination numbers or densities of packings [140, 139] is at best an exercise in modeling locally jammed packings. The best way to categorize random disk packings is to determine the maximally random jammed (MRJ) state [58] for each of the three jamming categories (local, collective and strict). Such investigations will be carried out in the future, and we have some preliminary results and promising avenues of approach.

The identification of the MRJ state for strictly jammed *disk* packings is an intriguing open problem. On

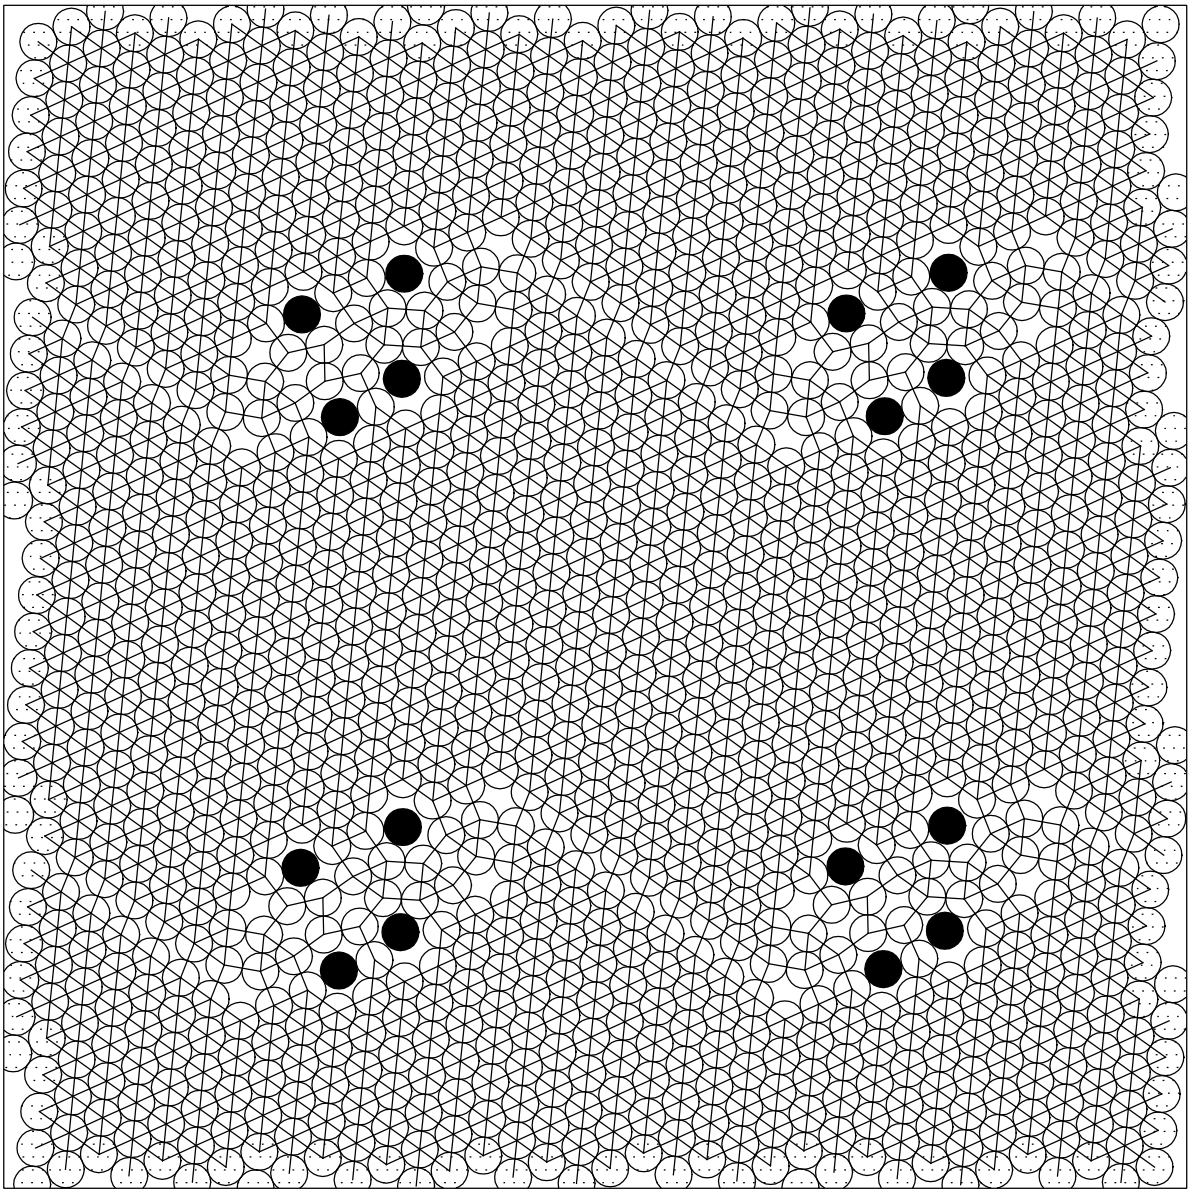

Figure 4.15: *Strictly jammed disk packing.* We show here $2x2$ unit cells of a *dense* ($\varphi = 0.88$) random packing of 250 disks that is *strictly jammed*, modulo four rattling particles, shown in black. This packing was produced with the extended Lubachevsky-Stillinger algorithm which allows for deformations of the lattice during the compression. We also display the contact network of the packing. The striking feature of this and similar strictly jammed disk packings we have produced is the appearance of peculiar "dislocation cores" and the appearance of large perfectly triangular regions.

the one hand, we have shown that random packings exist with densities in the vicinity of the maximum possible value ($\varphi = \frac{\pi}{2\sqrt{3}}$) that are not strictly jammed, and on the other hand, there is a conjectured achievable lower bound $\varphi \geq \frac{\sqrt{3}\pi}{8}$ corresponding to the "reinforced" Kagomé lattice (see Fig. 4.11). It may therefore be that the search for the MRJ state for strictly jammed disk packings should focus on randomly diluted triangular packings. For random sphere packings, an initial study undertaken in Ref. [56], using the LP algorithm described in this work, found that maximally disordered random packings around $\varphi \approx 0.63$ were strictly jammed, suggesting a close relation between the conventionally accepted RCP state and the MRJ state for strictly jammed packings. Much less obvious is what the MRJ state for collectively jammed sphere packings is. Finally, a completely unexplored question concerns the identification of the MRJ state for locally jammed disk and sphere packings.

Chapter 5

Hypostatic Jammed Packings of Hard Ellipsoids

In this chapter we apply the collision-driven algorithm presented in Chapter 3 to generate, for the first time, disordered jammed packings of hard ellipsoids in three [78] and two dimensions. We show that ellipsoids can randomly pack more densely; up to $\varphi = 0.68 - 0.71$ for spheroids with an aspect ratio close to that of M&M'S Candies®, and even approach $\varphi \approx 0.74$ for non-spheroidal ellipsoids. In several industrial processes such as sintering and ceramic formation interest exists in increasing the density and number of contacts of powder particles to be fused. It is worth noting that using ellipsoidal instead of spherical particles, we may increase the density of a randomly poured and compacted powder to almost the limit of a perfectly crystalline system. Both the density of the random packings ϕ_J and the average contact number $\bar{Z}$ rise sharply from their sphere values as asphericity is introduced, showing that the sphere point is a singular point and that orientational degrees of freedom dramatically impact the properties of jammed packings. The contact numbers reach values close to isostatic only for very aspherical particles, while the packings of nearly spherical ellipsoids are rather hypostatic, in sharp contrast with random sphere packings, which have been shown to be isostatic.

In order to understand the observed behavior, we consider the mathematics of jamming in packings of smooth strictly convex nonspherical hard particles [74]. We explain why the isostatic conjecture, stating that for large disordered jammed packings the average contact number is twice the number of degrees of freedom per particle, does not apply to nonspherical particles. We develop first- and second-order conditions for jamming, and demonstrate that ellipsoid packings can be jammed even though they are hypostatic. We apply an algorithm using these conditions to the computer-generated hypostatic ellipsoid and ellipse packings and demonstrate that our algorithm does produce jammed packings, even close to the sphere point. We also consider packings that are nearly jammed and draw connections to packings of deformable (but stiff) particles. Finally, we consider the jamming conditions for nearly spherical particles and explain quantitatively the behavior we observe in the vicinity of the sphere point.

5.1 Introduction

The structure of liquids, crystals and glasses, and transitions between these phases is intimately related to volume fractions of ordered and disordered (random) hard-sphere packings [70]. In dimensions higher then three, packing problems [8] are of current interest for insulating stored data from noise [9], and in two and three dimensions in relation to flow and jamming of granular materials [157, 158, 159] and glasses [4]. Most of the computational studies of jammed packings of hard particles have focused on hard spheres, primarily because it is much simpler to study hard spheres both computationally and theoretically. However, this leaves open the question of whether hard spheres are an appropriate model for systems where the particles are not truly spherical.

Through computer simulations, we find that asphericity, as measured by the deviation of the *aspect ratio* α from unity, dramatically affects the properties of jammed packings. In particular, we find that the *packing*

fraction (density) at jamming ϕ_J and the average *coordination* (*contact*) *number* per particle $\bar{Z}$ increase sharply from the typical sphere values $\phi_J \approx 0.64$ and $\bar{Z} = 6$ when moving away from the sphere point $\alpha = 1$. If one views ϕ_J and $\bar{Z}$ as functions of the particle shape, they have a cusp (non-analytic) minimum at the sphere point. The rapid increases are unrelated to any observable increase in order in these systems that develop neither crystalline (periodic) nor liquid crystalline (nematic or orientational) order.

There have been conjectures [160, 161, 162, 163, 143] that large disordered jammed packings of hard frictionless particles are *isostatic*, meaning that the number of contacts is equal to the number of degrees of freedom. Most of previous discussions of isostaticity have been in the context of spheres, and for sphere packings this *isostatic conjecture* has been verified computationally [75, 164]. For a general particle shape, the obvious generalization of the conjecture would produce the expectation $\bar{Z} = 2d_f$, where d_f is the number of degrees of freedom per particle ($d_f = 2$ for disks, $d_f = 3$ for ellipses, $d_f = 3$ for spheres, $d_f = 5$ for spheroids, and $d_f = 6$ for general ellipsoids). Since d_f increases discontinuously with the introduction of rotational degrees of freedom as one makes the particles non-spherical, the isostatic prediction would be that $\bar{Z}$ would have a jump at $\alpha = 1$. We do not observe such a discontinuity, rather, we observe that ellipsoid packings are *hypostatic*, $\bar{Z} < 2d_f$, near the sphere point, and only become close to isostatic for large aspect ratios (but still remain hypostatic).

We generalize our previous theoretical and computational investigations of jamming in sphere packings [75, 71] to packings of nonspherical particles, and in particular, packings of hard ellipsoids. We generalize the mathematical theory of rigidity of tensegrity frameworks [73, 165] to packings of nonspherical particles, and demonstrate rigorously that the ellipsoid packings we studied in Ref. [78] are jammed even very close to the sphere point. Armed with this theoretical understanding of jamming, we also obtain a quantitative understanding of the cusp-like behavior of ϕ_J and $\bar{Z}$ around the sphere point. We will repeat some of the discussion about jamming from Chapter 4, partly in order to make this chapter more self-contained, and partly in order to remove all of the assumptions we made that are true only for spheres. The reader is however encouraged to read at least Section 4.2 before reading this chapter.

In Section 5.2 we generate disordered jammed packings of hard ellipses and ellipsoids and study their properties as a function of the particle aspect ratio. In Section 5.4 we define jamming and we investigate the reasons for the failure of the isostatic conjecture for nonspherical particles. In Section 5.5 we develop the first- and second-order conditions for jamming in a system of nonspherical particles, and then design and use a practical algorithm to test these conditions in Section 5.5.3. In Section 5.7 we consider the thermodynamical behavior of hypostatic packings that are close to, but not quite at, the jamming point. In Section 5.8 we discuss the connections between jammed packings of hard particles and strict energy minima for systems of deformable particles. In Section 5.9 we focus on packings of nearly spherical ellipsoids, and finally, offer conclusions in Section 5.10.

5.2 MRJ Packings of Hard Ellipsoids

We use the generalized Lubachevsky-Stillinger (LS) sphere-packing algorithm [12], described in Chapter 3, to generate disordered jammed packings of hard ellipsoids. A typical configuration of 1000 oblate ellipsoids (the aspect ratio $\alpha = b/a = 1.9$) is shown in Fig. 5.1, with density of about $\varphi \approx 0.70$. The nematic order parameter is $S \approx 0.02 - 0.05$, which is consistent with complete absence of orientational (nematic) ordering for this size of the system [43]. Based on our experience with spheres [56], we believe that our algorithm (with rapid particle expansion) produces final states that represent the MRJ state well. The algorithm closely reproduces the packing fraction measured experimentally [78] for M&M'S Candies® , as detailed in Chapter 7. In simulations a contact is typically defined by a cutoff on the gap between the particles. Fortunately, over a wide range ($10^{-9} - 10^{-4}$) of contact tolerances, $\bar{Z}$ is reasonably constant, as discussed in detail in Chapter 9 for spheres (the discussion applies to hard ellipsoids as well).

In Fig. 5.2 we show the jamming density ϕ_J and contact number $\bar{Z}$ of jammed monodisperse packings of hard ellipsoids in three dimensions as a function of the particle shape. The ellipsoid semiaxes have ratios $a : b : c = 1 : \alpha^\beta : \alpha$ where $\alpha > 1$ is the *aspect ratio* (for general particle shapes, α is the ratio of the radius of the smallest circumscribed to the largest inscribed sphere), and $0 \le \beta \le 1$ is the "oblateness", or *skewness* ($\beta = 0$ corresponds to prolate and $\beta = 1$ to an oblate spheroid). It is seen that the density rises as a linear function of $\alpha - 1$ from its sphere value $\phi_J \approx 0.64$, reaching densities as high as $\phi_J \approx 0.74$ for the self-dual

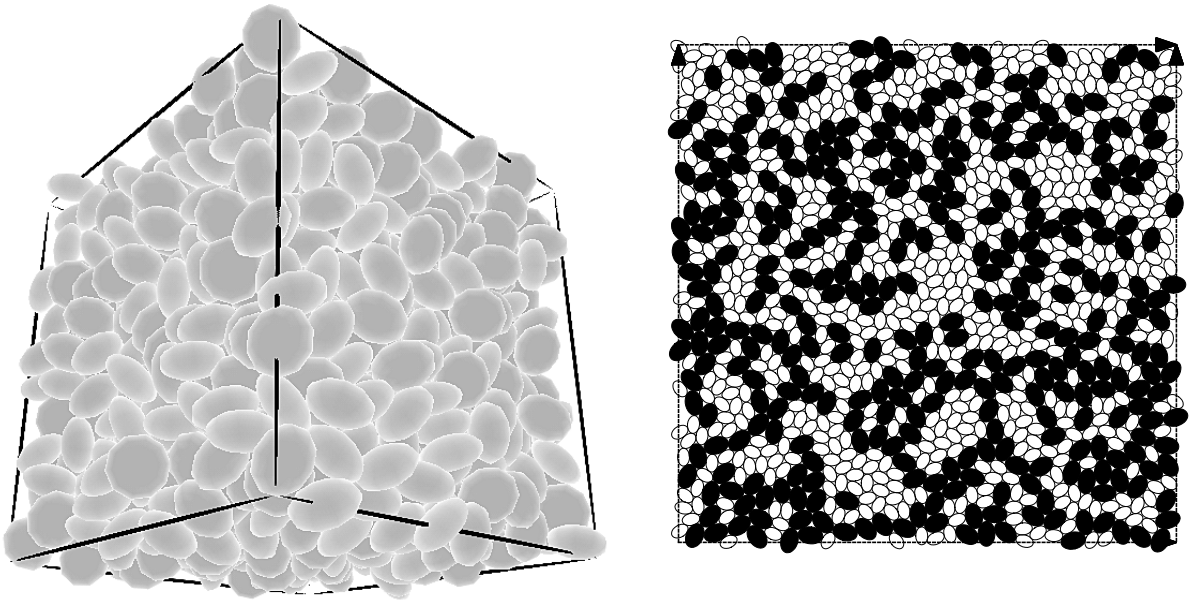

Figure 5.1: (*Left*) Computer-generated packing of 1000 oblate ellipsoids with $\alpha = 1.9$, close to the shape of M&M'S Candies®. (*Right*) Computer-generated packing of 1000 ellipses of aspect ratio $\alpha = 1.5$, being among the densest generated ellipse packings and almost as dense as the triangular disk crystal packing ($\phi \approx 0.906$).

ellipsoids with $\beta = 1/2$. The jamming density eventually decreases again for higher aspect ratios, however, we do not investigate that region in this work. The contact number also shows a rapid rise with $\alpha - 1$, and then plateaus at values somewhat below isostatic, $\bar{Z} \approx 10$ for spheroids, and $\bar{Z} \approx 12$ for nonspheroids (for $\beta = 1/4$ and $\beta = 3/4$, as well as $\beta = 1/2$). In Section 5.9 we will need to revert to two dimensions (ellipses) in order to make some analytical calculations possible. We therefore also generated jammed packings of ellipses, and show the results in Fig. 5.3. Since monodisperse packings of disks always crystallize and do not form disordered jammed packings, we used a binary packing of particles with one third of the particles being 1.4 times larger than the remaining two thirds. The ellipse packings show exactly the same qualitative behavior as ellipsoids, and an example packing is shown in Fig. 5.1.

Previous simulations for random sequential addition (RSA) [14], as well as gravitational deposition [38, 34], produce a similarly shaped curve, with a maximum at nearly the same aspect ratios $\alpha \approx 1.5$ (for both prolate and oblate spheroids), but with substantially lower volume fractions (such as $\varphi \approx 0.48$ for RSA). It is interesting to note that for both spheroids and general ellipsoids $\bar{Z}$ reaches a constant value close to the isostatic prediction $\bar{Z} = 2d_f$ (but still less), at approximately the aspect ratio for which the density has a maximum. This supports the claim the density decrease for large α comes from exclusion volume effects at constant coordination number [27]. In Section 5.9 we explain quantitatively why the density and contact numbers rise sharply near the sphere point. Since the density increases for aspect ratios near unity and then decreases for large aspect ratios, it is clear it must have a maximum. We do not yet understand why different dimensions, different particle shapes [27], and packings as different as RSA and MRJ, show a maximum at approximately the same aspect ratio $1.5 < \alpha < 2.0$.

Several interesting features can be noted in Fig. 5.2 concerning the variations of the density with changes in the particle shape, and especially with variations in the skewness β. Firstly, we observe the same packing density for ellipsoids with skewness β and their dual ellipsoids with skewness $1 - \beta$. In particular, to within statistical and algorithmic variations, prolate and oblate ellipsoids have identical packing densities within the ranges of aspect ratio we have studied (up to $\alpha \approx 10$, though for large aspect ratios our results are not as accurate and we do not show them here). The self-dual ellipsoids with $\beta = 1/2$ show the highest packing densities, and are in a sense most aspherical given a certain aspect ratio. Correspondence between prolate

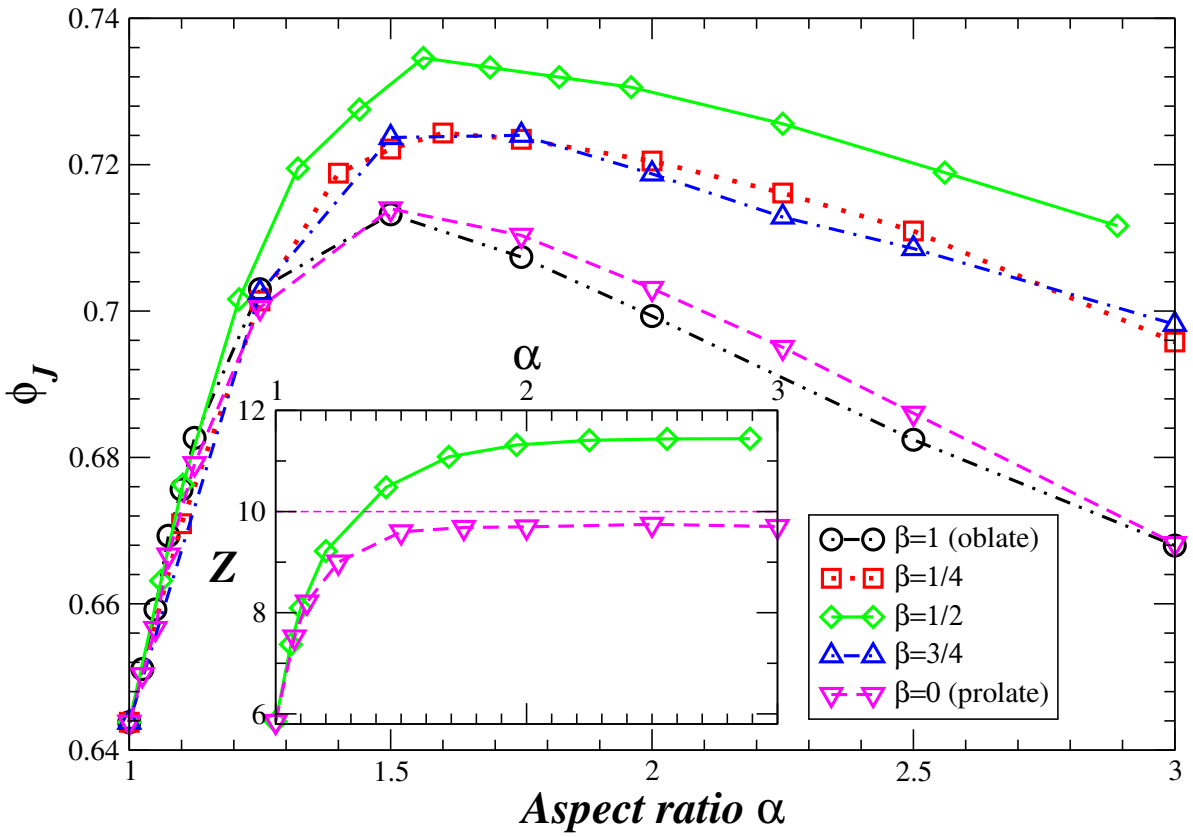

Figure 5.2: Jamming density and average contact number (inset) for packings of $N = 10000$ ellipsoids with ratios between the semiaxes of $1 : \alpha^{\beta} : \alpha$ [c.f. Fig. (2) in Ref. [78]]. The isostatic contact numbers of 10 and 12 are shown as a reference.

and oblate ellipsoids has been observed in equilibrium simulations of hard-ellipsoid fluids [166, 43], and it has been noted that the second virial coefficient, or more specifically, the scaled excluded volume

$$v_{ex} = \frac{V_{ex}}{8V},$$

is equal for a prolate and an oblate ellipsoid for the same aspect ratio. Here V_{ex} measures the orientation-averaged volume excluded to another particle with identical shape [43], and V is the particle volume, while the normalization factor of eight is used so that $v_{ex} = 1$ for spheres. The excluded volume plays a crucial role in many theories of the liquid state, and is invariant under the change of β with $1-\beta$ [167, 168]. We plot in Fig. 5.4 the data from Fig. 5.2 but with the aspect ratio replaced by v_{ex}, with the hope that comparisons among different particle shapes may find better agreement if one uses v_{ex} instead of α. Note however that even at the same scaled excluded volume and contact number, the packing density is different for different β.

Another interesting question is the exact form of the density and contact number increase around the sphere point. The inset shown in Fig. 5.2 suggests that the density rises linearly in $\alpha - 1$, and indeed we will predict exactly this behavior using analytical calculations in Section 5.9. It is computationally difficult to study $\bar{Z}$ very close to the sphere point due to the difficulty of distinguishing true contacts from near contacts, as discussed in Chapter 9, however, (semi-)analytical calculations in Section 5.9 suggest that near the sphere point $\bar{Z} \approx 2d + Z_{\alpha}\sqrt{\alpha - 1}$, and numerical results are consistent with such behavior. We note that random sequential addition (RSA) [14] also seems to produce a comparable linear increase in the saturation density as asphericity is introduced. Our RSA data has large statistical errors because efficient simulation of the RSA process is rather difficult for nonspherical particles however, especially near the sphere point. This is

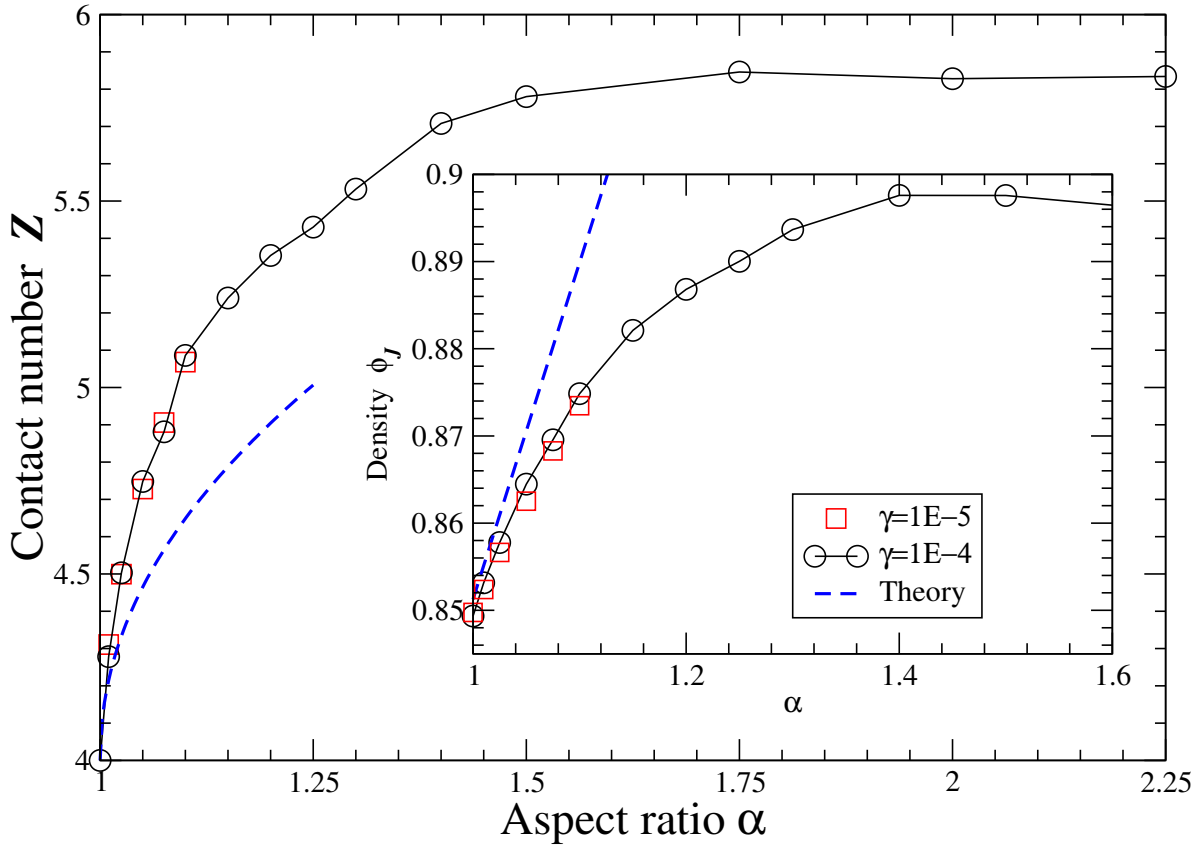

Figure 5.3: Average contact number and jamming density (inset) for bi-disperse packings of $N = 1000$ ellipses with ratios between the semiaxes of $1 : \alpha$, as produced by the MD algorithm using two different expansion rates γ (affecting the results only slightly). The isostatic contact number is 6. The results of the leading-order (in $\alpha - 1$) theory presented in Section 5.9 are shown for comparison.

because finding the correct orientation for insertion of a particle inside a tight void requires a lot more trials than just finding the right position of the centroid.

Finally, it is interesting to observe that the number of rattlers, i.e., particles without any touching neighbors, decays rapidly as asphericity is introduced. Where as sphere packings can have as many as 5% rattling particles in two dimensions or 2.5% in three dimensions, this number decreases rapidly for ellipsoids and for aspect ratios larger than about 1.25 we usually do not observe any rattling particles even in very large packings. The reasons for the occurrence of rattlers and their impact on the thermodynamic or mechanical properties of packings are not well understood. Experimental packings can be made to not contain any rattlers, without significantly affecting the packing properties. For example, this can be done by packing marbles inside an elastic balloon and compressing the balloon slowly until rattling is no longer heard when the packing is shaken. When we discuss contacts or near contacts, and quote contact numbers such as $\bar{Z}$, we will remove rattlers from consideration due to the additional complexity they would introduce otherwise.

5.3 Impenetrability and Interparticle Forces

In Chapter 4 we discussed jamming in hard-sphere systems in great detail, including linear programming algorithms to test for jamming in the different jamming categories of local, collective and strict jamming. In this section we consider extending the theory and algorithms to nonspherical particles. We do not repeat the discussion of boundary conditions and related details from Chapter 4 since those details apply for both

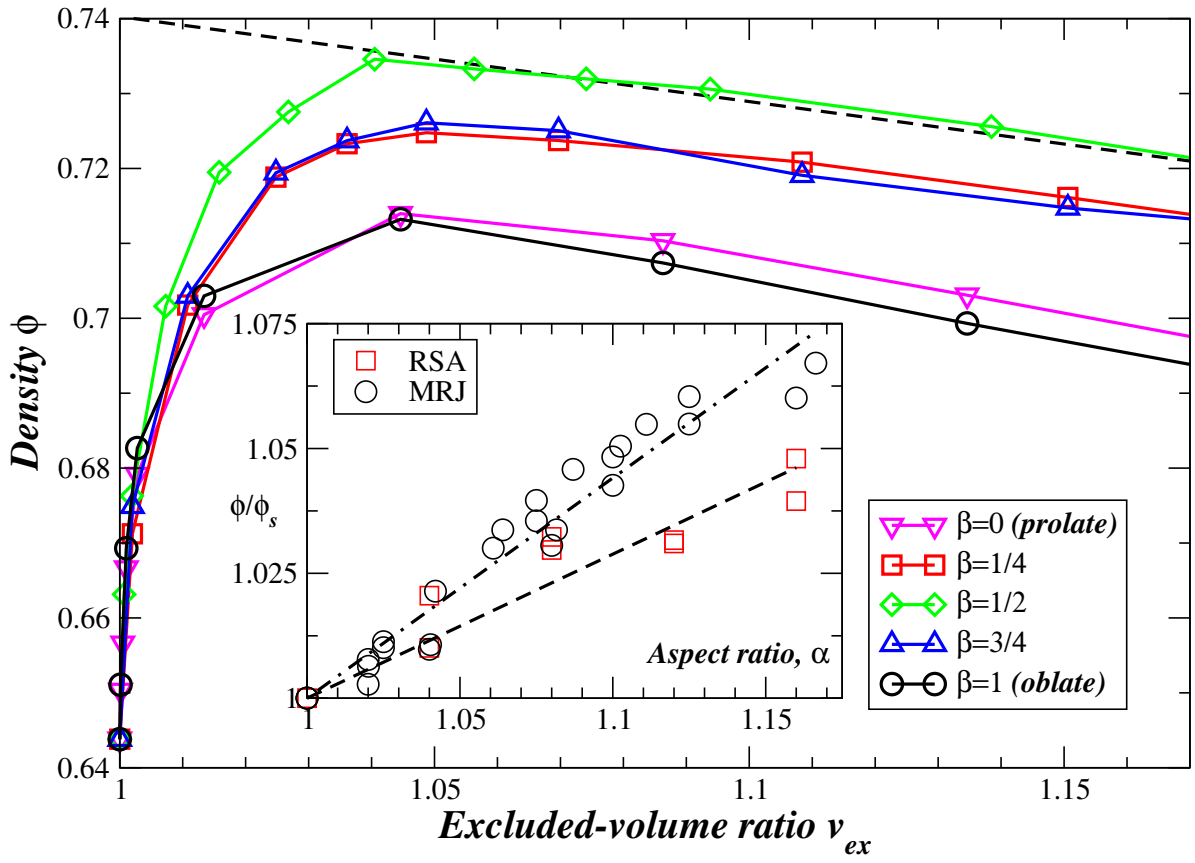

Figure 5.4: The density data as in Fig. 5.2, but with aspect ratio replaced by the scaled excluded volume. An interesting, but perhaps naive, extrapolation of the density for $\beta = 1/2$ from large aspect ratios to the sphere point produces a jamming density of about 0.74, very close to the density of the hard sphere crystal. The inset focuses on the region close to the sphere point and shows the jamming density for both MRJ packings and for saturated packings produced by random sequential addition (RSA). We do not try to distinguish the data for different β due to the presence of large statistical and systematic errors very close to the sphere point.

spherical and non-spherical particles. Instead, we focus on the differences between the impenetrability constraints for spherical and non-spherical particles.

5.3.1 The Rigidity Matrix

The addition of rotational degrees of freedom changed the rigidity matrix of the packing. In a very general sense, the rigidity matrix is the gradient of the overlap potential $\boldsymbol{\zeta} = (\zeta_{ij})$ with respect to the configuration

$$\mathbf{A} = \boldsymbol{\nabla}_{\mathbf{Q}} \boldsymbol{\zeta}.$$

This $[N_f \times M]$ matrix connects, to first order, the change in the interparticle gaps to the particle displacements, $\Delta\boldsymbol{\zeta} = \mathbf{A}^T \Delta\mathbf{Q}$. The rigidity matrix is sparse and has two blocks of d_f non-zero entries in the column corresponding to the particle contact $\{i, j\}$, namely, $\boldsymbol{\nabla}_i \zeta_{ij}$ in the block row corresponding to particle i and $\boldsymbol{\nabla}_j \zeta_{ij}$ in the block row corresponding to particle j (unless one of these particles is frozen). Represented schematically:

$$\mathbf{A} = \begin{array}{c} \\ i \rightarrow \\ \\ j \rightarrow \\ \\ \end{array} \begin{array}{c} \{i,j\} \\ \downarrow \\ \left[\begin{array}{c} \vdots \\ \boldsymbol{\nabla}_i \zeta_{ij} \\ \vdots \\ \boldsymbol{\nabla}_j \zeta_{ij} \\ \vdots \end{array}\right] \end{array} .$$

It may sometimes be more convenient to work with surface-to-surface interparticle gaps, $\Delta\mathbf{h} = \mathbf{A}_E^T \Delta\mathbf{Q}$ (the subscript E stands for Euclidean), especially if second-order terms are not considered, and this is what we did in Chapter 4 for spheres. For spheres this change of convention does nothing more than rescale the columns of the rigidity matrix by the sum of sphere radii, and it does not affect any of the results presented in this work. All of the discussion about boundary conditions and its impact on the rigidity matrix carries over from spheres to ellipsoids, and we do not repeat the tedious details here but rather focus on the impenetrability constraints between pairs of particles.

5.3.2 Interparticle Forces

Hard particles in contact can exert a compressive (positive) *contact force* $\mathbf{f} = f\mathbf{n}$ directed along the normal vector (for frictionless particles). The total force and torque exerted on a given particle i by the contacts with its neighbors $\mathcal{N}(i)$ is

$$\Delta\mathbf{b}_i = -\sum_{j\in\mathcal{N}(i)} f_{ij} \begin{bmatrix} \mathbf{n}_{ij} \\ \left(\mathbf{r}_{iC}^{ij} \times \mathbf{n}_{ij}\right) \end{bmatrix} = \sum f_{ij} \left(\boldsymbol{\nabla}_i h_{ij}\right),$$

or, considering all particles together

$$\Delta\mathbf{B} = \mathbf{A}_E \mathbf{f}.$$

The fact that the matrix (linear operator) connecting force imbalances to contact forces is the transpose of the rigidity matrix is well-known and can also be derived by considering the work done by the contact forces to displace the particles

$$\begin{aligned} W &= \Delta\mathbf{B}^T \Delta\mathbf{Q} = \left(\tilde{\mathbf{A}}\mathbf{f}\right)^T \Delta\mathbf{Q} = \mathbf{f}^T \left(\tilde{\mathbf{A}}^T \Delta\mathbf{Q}\right) = \\ &= \mathbf{f}^T \Delta\mathbf{h} = \mathbf{f}^T \left(\mathbf{A}_E^T \Delta\mathbf{Q}\right), \end{aligned}$$

showing that $\tilde{\mathbf{A}} = \mathbf{A}_E^T$. In this work we will use forces $\mathbf{f}$ that are a rescaled version of the physical forces $\mathbf{f}_E$, $f_{ij} = \left(\mathbf{r}_{ij}^T \mathbf{n}_{ij}\right) f_{ij}^E/2$, so that $\mathbf{A}\mathbf{f} = \mathbf{A}_E \mathbf{f}_E$. This scaling is more natural for our choice of overlap potential, and does not affect any of the results.

In static packings, the contact forces must be balanced, i.e., the force/torque equilibrium condition

$$\mathbf{A}\mathbf{f} = \mathbf{0} \text{ and } \mathbf{f} \geq \mathbf{0}$$

must be satisfied. The number of equations is equal to the total number of degrees of freedom $N_f = Nd_f$, and the total number of equations is equal to the number of contacts M. The actual magnitudes of the forces are determined by external loads (for example the applied pressure for a system of deformable particles), history of the packing preparation, etc.; however, the relation between the forces at different contacts is determined by the packing geometry, or more specifically, by $\mathbf{A}$. Typically forces are rescaled to a mean value of unity, $\mathbf{e}^T\mathbf{f} = M$, and it is has been observed that the distribution of rescaled contact forces has some universal features, for example, there is an exponential tail of contacts carrying a large force, and also a large number of contacts supporting nearly zero force [75, 169]. We will see later that these *force chains*, or *internal stresses*, are an essential ingredient of jamming for hard particles.

5.4 The Isostatic Conjecture

In the granular materials literature special attention is often paid to so-called isostatic packings, and it has been postulated several times that realistic packings of hard particles are isostatic. There are several different definitions of isostaticity, and most of the discussions in the literature are specifically applied to hard spheres or to frictional particles. In this section we summarize several relevant definitions of and arguments for isostaticity and generalize them to nonspherical particles. Our arguments will arrive to the following definition of isostaticity, which we believe is the correct generalization to systems of nonspherical particles. A packing is *isostatic* if the number of constraints (contacts) is equal to the number of degrees of freedom

$$N_c = N_f \text{ or } N_c = N_f + 1,$$

where we can explicitly add +1 if we do not count the density as a degree of freedom. Packings with less contacts than isostatic are called *hypostatic*, and packings with more contacts than isostatic are *hyperstatic*. The *isostatic conjecture* states that *large disordered packings of hard particles are isostatic*. Defining what precisely is meant by a disordered packing is difficult in itself [170, 58], however, intuitively, in a disordered packing there is only the minimal degree of correlations between particles, as necessitated by the constraints of impenetrability and jamming. Therefore, it is expected that in a certain sense disordered packings are *generic* [137], and that "special" configurations with geometric degeneracies will not appear.

Note that for large systems the majority of the degrees of freedom come from the particles themselves, $N_f \approx Nd_f$, and the majority of constraints come from contacts shared between two particles, $N_c \approx M = N\bar{Z}/2$, giving the isostatic property

$$\bar{Z} = 2d_f. \tag{5.1}$$

Equation (5.1) has been verified to very high accuracy for jammed hard-sphere packings [75], however, disordered packings of hard ellipsoids are always hypostatic and thus contradict the isostatic conjecture [78]. In this Section we attempt to deconstruct previous discussions of the isostatic conjecture and jamming in hard-particle packings, and we hope that through our discussions it will become clear why previous "proofs" of the isostatic conjecture do not apply to nonspherical particles, or to put it the other way around, what makes disordered sphere packings isostatic.

5.4.1 Jamming, Rigidity and Stability

An essential initial step is defining more precisely what is meant by a stable, rigid, or jammed packing. All of these terms have been used in the literature, and in fact we equate each of them with a particular perspective on jamming:

Kinematic A packing is *jammed* if none of the particles can be displaced in a non-trivial way without introducing overlap between some particles.

Static A packing is *rigid* if it can resolve any externally applied forces through interparticle ones, without changing the packing configuration.

Perturbation A packing is *stable* if the structure of the packing changes smoothly for small (local or global) perturbations of the packing.

We will consider each of these approaches separately. It will shortly become clear that all of them are closely related, and under certain mild conditions they are actually equivalent. This is the case for spheres, and we focused on the kinematic perspective in Section 4.2; however, one can use any of the different views of jamming as well.

We will use the term jamming as an umbrella term, and later give our preferred definition of jamming, which is based on the kinematic perspective. We note that it is important to precisely specify the boundary conditions applied regardless of the view used in considering jamming; different boundary conditions lead to different jamming categories, specifically *local*, *collective* or *strict* jamming [15, 71], as discussed in Section 4.2. Here, we will sometimes use local jamming in simple examples but mostly focus on collective jamming; all particles are allowed to move collectively and we will employ periodic boundary conditions with fixed

lattice vectors. In order to eliminate trivial uniform translations of the systems, we can freeze the centroid of one of the particles, to obtain a total of

$$N_f = Nd_f - d$$

degrees of freedom. The exact boundary conditions affect the counting of constraints and degrees of freedom, however, the correction is not extensive in N and therefore is negligible for large system when considering quantities on a per particle basis (such as, for example, $\bar{Z}$).

An important point to note is that the above definitions of jamming treat all degrees of freedom identically. In particular, translational motion (forces) is treated on the same footing as rotational motion (torques). This is not necessarily the most appropriate definition, as is easily seen by considering the case of spheres, which can rotate in place freely even though they are (translationally) jammed. This distinction between translations and rotations will become more explicit in Section 5.7 when considering packings that are nearly, but not quite, jammed.

5.4.1.1 Kinematic View

The kinematic perspective considers a packing jammed if it is not possible to continuously displace the particles in a non-trivial way without introducing overlap. That is, the impenetrability conditions preclude any motion of the particles. Here trivial motions are those that do not change the distances between any two particles, such as global translations when periodic boundary conditions are used. We can assume that such trivial motions have been eliminated via some artificial constraint, such as fixing the centroid of one particle externally when using periodic boundary conditions.

Mathematically, for any continuous motion $\Delta\mathbf{Q}(t)$ there exists a $T > 0$ such that at least one of the impenetrability constraints between a touching pairs of particles

$$\boldsymbol{\zeta}\left[\mathbf{Q}_J + \Delta\mathbf{Q}(t)\right] \geq \mathbf{0} \tag{5.2}$$

is violated for all $0 < t < T$. A motion $\Delta\mathbf{Q}(t)$ such that for all $0 < t < T$ none of the constraints are violated is an *unjamming motion*. One can in fact restrict attention to analytic paths $\Delta\mathbf{Q}(t)$, and also show that a jammed packing is in a sense isolated in configuration space, since the only way to get to a different packing is via a discontinuous displacement $\|\Delta\mathbf{Q}\| > 0$ [73].

A similar definition of jamming was used by Alexander in Ref. [160]. He considers a packing to be *geometrically rigid* if it cannot be "deformed continuously by rotating and translating the constituent grains without deforming any of them and without *breaking* the contacts between any two grains". This definition implies that a packing in which particles can be moved so as to break contacts (for example, imagine a pebble resting on other pebbles in gravity, and moving it upward away from the floor) is jammed. Later in the manuscript Alexander talks about adding constraints to block motions that break contacts. We in fact have in a certain sense a choice in the matter, determining whether we work with inequality or equality constraints. We choose to work with inequality constraints, since this is the natural choice for frictionless hard particles; there is no cohesion between the particles maintaining contacts. In effect, when counting degrees of freedom, we count the possible collective rescaling of the particle shapes (i.e., the density ϕ) necessary to maintain contacts as a single degree of freedom.

5.4.1.2 Static View

The static perspective considers a packing *rigid* if it can resolve any applied forces through interparticle ones, without changing the packing configuration. This is sometimes referred to as *static rigidity*, to be contrasted with *kinematic rigidity* as defined in Section 5.4.1.1. For hard particles, there is no scale for the forces, and so the actual magnitude of the forces does not matter, only the relative magnitudes and the directions. The particles do not deform, but can exert an arbitrary positive contact force.

Mathematically, we consider the existence of a solution to the force-equilibrium equations

$$\mathbf{Af} = -\mathbf{B}, \text{ where } \mathbf{f} \geq 0, \tag{5.3}$$

for all resolvable external loads $\mathbf{B}$. The space of resolvable loads is determined by the boundary conditions: certain forces such as pulling on the walls of a container cannot be resolved by packings and need to be

excluded. This is similar to the definition used by Witten in Ref. [162]: A packing is *mechanically stable* "if there is a nonzero measure set of external forces which can be balanced by interbead ones." The problem with this definition of rigidity and in particular Eq. (5.3) is that it does not take into account the fact that the geometry of the packing, i.e., the rigidity matrix $\mathbf{A}$, changes when an external load is applied on the packing. That is, physically, forces arise only through deformation, and this deformation, however small, may need to be taken into account. Forces are in essence Lagrange multipliers associated with the impenetrability constraints in Eq. (5.2); the very existence of such Lagrange multipliers may require a change in the packing configuration.

Various counting arguments related to force equilibrium constraints, starting with the seminal work of Maxwell, have appeared in the engineering literature on mechanical structures (see references in Ref. [137]). In this literature, a structure is called *statically determinate* if the forces necessary to support a given external load can be uniquely determined from the equations of force balance. An isostatic jammed packing is necessarily statically determinate up to a multiplier of the self-stress of the packing. A hyperstatic packing is statically underdetermined, since there are multiple ways to resolve almost any applied load. In this case constitutive (elastic) laws need to be invoked to determine the forces. A hypostatic packing on the other hand is statically overdetermined, and as such is considered unstable in the literature on mechanical structures.

5.4.1.3 Perturbation View

The perturbation perspective considers a packing to be stable if the structure of the packing changes smoothly for small perturbations of the packing. In particular, the structure of the packing includes the positions of the particles *and* the contact force network. Perturbations to be considered should include changes in the grain internal geometry (deformation), strain, and stress (external forces due to shaking, vibration, or a macroscopic load). In great generality we can restrict our perturbations to small perturbations of the distances between contacting particles *combined* with small perturbations of the applied forces.

Mathematically, we consider the sensitivity of the configuration and force chains to *all* perturbations of the interparticle gaps $\Delta\boldsymbol{\zeta}$ and applied forces $\Delta\mathbf{B}$ away from zero, i.e., we look for solutions of the *coupled* system of equations of preserving contacts and maintaining force equilibrium:

$$\begin{aligned} \left[\mathbf{A}\left(\mathbf{Q}+\Delta\mathbf{Q}\right)\right]\left(\mathbf{f}+\Delta\mathbf{f}\right) &= -\varepsilon\Delta\mathbf{B} \\ \boldsymbol{\zeta}\left(\mathbf{Q}+\Delta\mathbf{Q}\right)-\Delta\zeta_{\mu} &= -\varepsilon\Delta\boldsymbol{\zeta} \\ \mathbf{e}^{T}\Delta\mathbf{f} &= \mathbf{0}, \end{aligned} \tag{5.4}$$

where $\varepsilon > 0$ is a small number and we have assumed $\mathbf{f} > \mathbf{0}$. Similarly to the external forces, the space of resolvable gap perturbations is determined by the boundary conditions: global expansions will lead to gaps that cannot all be closed unless the particles grow by a certain scaling factor $\mu = 1 + \Delta\mu$. It is therefore convenient to include $\Delta\zeta_{\mu} \approx 2\Delta\mu$ as an additional variable. An added constraint is that the normalization $\mathbf{e}^{T}\mathbf{f} = M$ be maintained. It is important to note that we explicitly account for the dependence of the rigidity matrix on the configuration in the force-balance equation. Notice that when we combine perturbations of the geometry and forces together, the total number of variables is $M + N_f$, and the total number of constraints is also $M + N_f$ (here we include the global particle rescaling $\Delta\zeta_{\mu}$ as a degree of freedom). Therefore there are no underdetermined (linear) systems as found in counting arguments that consider geometry and forces separately, as is typically done in the literature.

5.4.2 Isostaticity

In this section we will attempt to deconstruct previous arguments in justification of an isostatic conjecture, mostly in the context of sphere packings, and try to identify the problems when the same arguments are applied to nonspherical particles.

The isostatic conjecture (property) is usually justified in two steps. First, an inequality $\bar{Z} \leq 2d_f$ is demonstrated, then, the converse inequality $\bar{Z} \geq 2d_f$ is invoked to demonstrate the equality $\bar{Z} = 2d_f$. We will demonstrate that it is the second of these steps that fails for non-spherical particles. However, first we recall some typical justifications for the inequality $\bar{Z} \leq 2d_f$. Isostaticity is also extensively discussed by Roux [143], and although there are close connections between our discussion here and the one in Ref. [143], we will not discuss the similarities or differences due to space limitations.

5.4.2.1 Why $\bar{Z} \leq 2d_f$ applies

A packing with $\bar{Z} > 2d_f$ is overconstrained, and in a certain sense geometrically degenerate and thus not "random." It can be argued that such a packing is not stable against small perturbations of the packing geometry, since all contacts cannot be maintained closed without deforming some of the particles. For example, Witten [162] considers hard-sphere packings with a small polydispersity, so that particles have slightly different sizes, to conclude that "the creation of a contact network with coordination number higher than $2d$ occurs with probability zero in an ensemble of spheres with a continuous distribution of diameters." Moukarzel [161, 171] considers how the actual stiffness modulus of deformable particles affects the interparticle forces and concludes that making the particles very stiff will eventually lead to negative forces and thus breaking of contacts, until the remaining contact network has $\bar{Z} \leq 2d_f$: "The contact network of a granular packing becomes isostatic when the stiffness is so large that the typical self-stress...would be much larger than the typical load-induced stress...granular packings will only fail to be isostatic if the applied compressive forces are strong enough to close interparticle gaps establishing redundant contacts." A similar argument is made by Sir Edwards in Ref. [163] for frictional grains: "if $z > 4$ then there is a solution with no force on $z - 4$ contacts, and there is no reason why other solutions would have validity."

These arguments apply also to nonspherical particles, however, it is important to point out that they specifically only apply to truly hard-particle packings or to packings of deformable particles in the limit of zero applied pressure (see Section 5.8.1). In real physical systems, particles will have a finite stiffness and the applied forces will be non-negligible, and such packings will have more contacts than the idealized hard-particle construction.

5.4.2.2 Why $\bar{Z} \geq 2d_f$ does not apply

The converse inequality, stating that a minimum of $2d_f$ contacts is necessary for jamming (rigidity), does not apply to nonspherical particles. We can demonstrate this vividly with a simple example of an ellipse jammed between three other stationary (fixed) ellipses, as shown in Fig. 5.5. Jamming a disk requires at least three touching disks; the additional rotational degree of freedom of the ellipse would seem to indicate that four touching ellipses would be needed in order to jam an ellipse. However, this is not true: if the normal contact vectors intersect at a single point, three ellipses can trap another ellipse, as shown in Fig. 5.5. We will shortly develop tools that can be used to demonstrate rigorously that this example is indeed jammed.

The above example shows that the claim of Ref. [160] that "One requires 4 $(= 3 + 1)$ contacts to fix the DOF [degrees of freedom]...of an ellipse in the plane" is wrong. Similarly, it shows that the argument in Ref. [161], namely, that the minimum number of contacts needed for a packing of N spheres in d dimensions to be rigid is dN, cannot be generalized to nonspherical particles by simply replacing d with d_f. Claims that the number of constraints must be larger than the number of degrees of freedom have been made numerous times within the kinematic perspective on jamming, for example, in Ref. [163]. Our careful analysis of the conditions for jamming in the next section will elucidate why this is correct for spheres but not necessarily correct for nonspherical particles, and under what conditions a hypostatic packing can be jammed.

Another argument, made within the static perspective on jamming [c.f. Eq. (5.3)], is given by Witten [162]: "The number of equilibrium equations Nd should not exceed the number of force variables N_c; otherwise these forces would be overdetermined." The example in Fig. 5.5 demonstrates why this argument cannot be applied to nonspherical grains. The system of force and torque balance equations is not really overdetermined since the torque balance equation is identically zero due to the peculiar geometry of the packing. Since the normal vectors at the points of contact intersect at a point, a torque around that point cannot be resolved by any set of normal forces between the particles. Yet the packing is jammed, and if built in the laboratory it will resist the torque by slight deformations of the particles, so that the normal vectors no longer intersect in one point and the contact forces can resist the applied torque. The connection between the geometry of the contact network, i.e., $\mathbf{A}$, and the packing configuration $\mathbf{Q}$, must be taken into account when considering the response of hypostatic packings to external perturbations. We discuss this connection in detail in the next section.

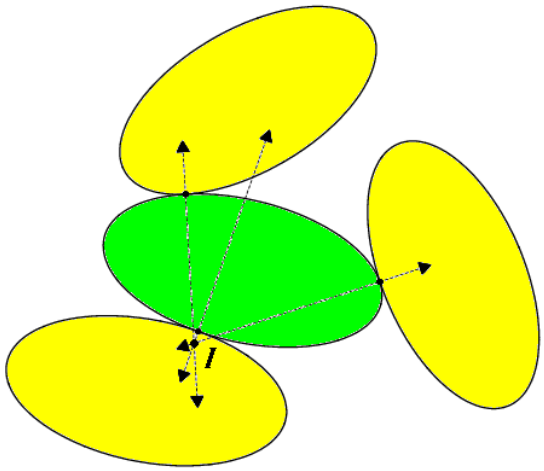

Figure 5.5: A mobile ellipse (green) jammed between three fixed ellipses (yellow). All ellipses are of the same size and have an aspect ratio $\alpha = 2$. This packing was produced by a Lubachevsky-Stillinger type algorithm, where the three particles were kept fixed by giving them infinite mass and no initial velocities. The normal vectors at the points of contact intersect at a common point I, as is necessary to achieve torque balance. This leads to a geometrically-degenerate configuration which would usually be dismissed as a probability-zero configuration. The number of force balance constraints here is 2, and the number of torque constraints is 1, giving a total of 3 constraints.

5.5 Conditions for Jamming

In this Section we develop first and second order conditions for jamming, using a kinematic approach. Statics (forces) will emerge through the use of duality theory. The discussion here is an adaptation of the theory of first-order, pre-stress, and second-order rigidity developed for tensegrities in Ref. [73]. The discussion in this section is technical and in Section 5.8 the rigorous hard-particle results are explained more simply by considering the conditions for local (stable) energy minima in soft-particle systems.

We consider an analytic motion of the particles

$$\Delta\mathbf{Q}\,(t) = \dot{\mathbf{Q}}t + \ddot{\mathbf{Q}}\frac{t^2}{2} + O(t^3),$$

where $\dot{\mathbf{Q}}$ are the velocities, and $\ddot{\mathbf{Q}}$ are the accelerations. Expanding the distances between touching particles to second-order, and taking into account that $\boldsymbol{\zeta}\,(\mathbf{Q}_J) = \mathbf{0}$, gives

$$\boldsymbol{\zeta}(t) \approx \mathbf{A}^T\dot{\mathbf{Q}}t + \left[\dot{\mathbf{Q}}^T\mathcal{H}\dot{\mathbf{Q}} + \mathbf{A}^T\ddot{\mathbf{Q}}\right]\frac{t^2}{2} = \dot{\boldsymbol{\zeta}}t + \ddot{\boldsymbol{\zeta}}\frac{t^2}{2}, \tag{5.5}$$

where the Hessian $\mathcal{H} = \boldsymbol{\nabla}^2_{\mathbf{Q}}\boldsymbol{\zeta} = \boldsymbol{\nabla}_{\mathbf{Q}}\mathbf{A}$ can be thought of as a higher-rank symmetric matrix.

5.5.1 First-Order Terms

Velocities $\dot{\mathbf{Q}} \neq \mathbf{0}$ for which $\dot{\boldsymbol{\zeta}} = \mathbf{A}^T\dot{\mathbf{Q}} \geq \mathbf{0}$ represent a *first-order flex* (using the terminology of Ref. [73]). If we can find an unjamming motion $\dot{\mathbf{Q}}$ such that $\dot{\boldsymbol{\zeta}} > \mathbf{0}$ (note the strict inequality), then the packing is *first-order flexible*, and there exists a $T > 0$ such that none of the impenetrability conditions [c.f. Eq. (5.2)] are violated for $0 \leq t < T$. We call such a $\dot{\mathbf{Q}}$ a *strict first-order flex*. If on the other hand for at least one constraint $\dot{\zeta} < 0$ for every $\dot{\mathbf{Q}}$, then the packing is jammed, since every non-trivial movement of the particles violates some impenetrability condition for all $0 < t < T$ for some $T > 0$. We call such a packing *first-order*

jammed. Finally, a $\dot{\mathbf{Q}}$ such that $\dot{\boldsymbol{\zeta}} = \mathbf{0}$ is a *null first-order flex*, often referred to as *zero* or *floppy mode* in the physics literature.

A packing is first-order jammed if and only if there are no (non-trivial) first order flexes. A packing is first-order flexible if there exists a strict first-order flex. Some packings are neither first-order jammed nor first-order flexible; One must consider higher-order terms to access whether such packings are jammed, and if they are not, to identify an unjamming motion. We will consider the second-order terms later; in this section we develop conditions and algorithms to verify first-order jamming and identify first-order flexes if they exist. The algorithms are closely based on work in Ref. [71].

5.5.1.1 Strict Self-Stresses

Let us first focus on a single contact $\{i, j\}$, and ask whether one can find a first order flex that is strict on that contact, i.e.,

$$\dot{\zeta}_{ij} = \left(\mathbf{A}^T\dot{\mathbf{Q}}\right)_{ij} = \left(\mathbf{A}^T\dot{\mathbf{Q}}\right)^T \mathbf{e}_{ij} = (\mathbf{A}\mathbf{e}_{ij})^T \dot{\mathbf{Q}} > 0,$$

where $\mathbf{e}_{ij}$ denotes a vector that has all zero entries other than the unit entry corresponding to contact $\{i, j\}$. If it exists, such a flex can be found by solving the linear program (LP)

$$\begin{array}{lll} \max\limits_{\dot{\mathbf{Q}}} & (\mathbf{A}\mathbf{e}_{ij})^T \dot{\mathbf{Q}} & \\ \mathbf{A}^T\dot{\mathbf{Q}} & \geq & \mathbf{0}. \end{array} \tag{5.6}$$

If this LP has optimal objective value of zero, then there is no first-order flex that is strict on the contact in question. Otherwise, the LP is unbounded, with an infinite optimal objective value. The dual LP of (5.6) is a feasibility problem

$$\begin{aligned} \mathbf{A}\left(\tilde{\mathbf{f}} + \mathbf{e}_{ij}\right) &= \mathbf{0} \\ \tilde{\mathbf{f}} &\geq \mathbf{0}, \end{aligned} \tag{5.7}$$

where the contact forces $\tilde{\mathbf{f}}$ are the Lagrange multipliers corresponding to the impenetrability constraints $\mathbf{A}^T\dot{\mathbf{Q}} \geq \mathbf{0}$. If the dual LP (5.7) is feasible, then the primal LP (5.6) is bounded. If we identify $\mathbf{f} = \tilde{\mathbf{f}} + \mathbf{e}_{ij} \geq \mathbf{0}$, $\tilde{f}_{ij} \geq 1$, we are naturally led to consider the existence of non-trivial solutions to the force-equilibrium equations

$$\mathbf{A}\mathbf{f} = \mathbf{0} \text{ and } \mathbf{f} \geq \mathbf{0}. \tag{5.8}$$

A set of non-negative contact forces $\mathbf{f} \neq \mathbf{0}$ that are in equilibrium as given by Eq. (5.8) is called a *self-stress*. In Ref. [73] these are called *proper self-stresses*, as opposed self-stresses which are not required to be non-negative. Self-stresses can be scaled by an arbitrary positive factor, so we will often add a normalization constraint that the average force be unity, $\mathbf{e}^T\mathbf{f} = M$. A self-stress that is strictly positive on a given contact is *strict* on that contact. A self-stress $\mathbf{f} > 0$ is a *strict-self stress*. The existence of a (strict) self-stress can be tested by solving the linear program

$$\begin{aligned} \max_{\mathbf{f},\varepsilon} \quad & \varepsilon \\ \mathbf{A}\mathbf{f} &= \mathbf{0} \\ \mathbf{f} &\geq \varepsilon\mathbf{e} \\ \mathbf{e}^T\mathbf{f} &= M \end{aligned} \tag{5.9}$$

and seeing whether the optimal value is negative (no self-stress exists), positive (a strict self-stress exists), or zero (a self-stress exists). What we showed above using linear duality is that if there is a self stress that is strict on a given contact, there is no flex strict on that contact. In particular, this means that *packings that have a strict self-stress can only have null first-order flexes.*

We can also show that there is a first-order flex that is strict on all contacts that do not carry a force in any self-stress (i.e., no self-stress is strict on them). To this end, we look for a first-order flex that is strict

on a given subset of the contacts, as denoted by the positions of the unit entries in the vector $\tilde{\mathbf{e}}$

$$\begin{aligned} \max_{\dot{\mathbf{Q}},\varepsilon} \quad & \varepsilon \\ \mathbf{A}^T\dot{\mathbf{Q}} \quad & \geq \quad \varepsilon\tilde{\mathbf{e}}. \end{aligned} \tag{5.10}$$

The dual program is the feasibility problem

$$\begin{aligned} \mathbf{A}\mathbf{f} \quad &= \quad \mathbf{0} \\ \tilde{\mathbf{e}}^T\mathbf{f} \quad &= \quad 1 \\ \mathbf{f} \quad &\geq \quad \mathbf{0}, \end{aligned} \tag{5.11}$$

which is infeasible if there is no self-stress that is positive on at least on the contacts under consideration, since $\tilde{\mathbf{e}}^T\mathbf{f} \equiv 0$. Therefore the primal problem (5.9) is unbounded, that is, one can find a self-stress that is strict (since $\varepsilon \to \infty$) on the given set of contacts. This shows that *packings that do not have a self-stress are first-order flexible*. In other words, the existence of force chains in a packing is a *necessary* criterion for jamming.

In summary, if a packing has no self-stress, it is not jammed, and one can easily find a strict first-order flex by solving a linear program [71]. The analysis is simplified if the packing has a strict self-stress, since in that case all first-order flexes are null, i.e., they are solutions of a linear system of equalities $\mathbf{A}^T\dot{\mathbf{Q}} = \mathbf{0}$. This is the case of practical importance to jammed packings, so we will focus on it henceforth.

5.5.1.2 Floppy Modes

The linear system $\mathbf{A}^T\dot{\mathbf{Q}} = \mathbf{0}$ has $N_{\text{floppy}} = N_f - r$ solutions, where $r = M - N_{\text{stresses}}$ is the rank of the rigidity matrix, and N_{stresses} is the number of (not necessarily proper) self-stresses (more precisely, the dimensionality of the solution space of $\mathbf{A}\mathbf{f} = 0$). We know that $N_{\text{stresses}} \geq 1$ for a jammed packing. If the packing is not hypostatic, or more precisely, if the number of contacts is sufficiently large

$$M = N_f + N_{\text{stresses}} \geq N_f + 1,$$

then there are no non-trivial null first-order flexes (floppy modes), $N_{\text{floppy}} = 0$. Therefore, *a packing that has a strict self-stress and a rigidity matrix of full-rank is (first-order) jammed*. We will later show that this sufficient condition for jamming is also necessary for sphere packings, that is, *jammed sphere packings are never hypostatic*.

However, we will see that jammed ellipsoid packings may be hypostatic, $M < N_f + 1$. Such a hypostatic packing always has floppy modes,

$$N_{\text{floppy}} = N_f + N_{\text{stresses}} - M \geq N_f + 1 - M.$$

Every floppy mode can be expressed as a linear combination of a set of N_{floppy} basis vectors, i.e.,

$$\dot{\mathbf{Q}} = \mathbf{V}\mathbf{x} \text{ for some } \mathbf{x}, \tag{5.12}$$

where the matrix $\mathbf{V}$ is a basis for the null-space of $\mathbf{A}^T$. To determine whether any of the null first-order flexes can be extended into a true unjamming motion, we need to consider second-order terms, which we do next.

5.5.2 Second-Order Terms

Consider a given null first-order flex $\mathbf{A}^T\dot{\mathbf{Q}} = \mathbf{0}$. We want to look for accelerations $\ddot{\mathbf{Q}}$ that make the second-order term in the expansion (5.5) non-negative, i.e.,

$$\mathbf{A}^T\ddot{\mathbf{Q}} \geq -\dot{\mathbf{Q}}^T\mathcal{H}\dot{\mathbf{Q}}. \tag{5.13}$$

If we cannot find such a $\ddot{\mathbf{Q}}$ for any first-order flex, then the packing is *second-order jammed*. If we find a $\ddot{\mathbf{Q}}$ such that all inequalities in (5.13) are strict, than we call the unjamming motion $\left(\dot{\mathbf{Q}}, \ddot{\mathbf{Q}}\right)$ a *strict second-order*

flex, and the packing is *second-order flexible*, since there exists a $T > 0$ such that none of the impenetrability conditions [c.f. Eq. (5.2)] are violated for $0 \leq t < T$. If for all first-order flexes $\dot{\mathbf{Q}}$ at least one of the inequalities in (5.13) has to be an equality, then we need to consider even third- or higher-order terms, however, we will see that for sphere and ellipsoid packings this is not necessary.

5.5.2.1 The Stress Matrix

In order to find a strict second-order flex, we need to solve the LP

$$\begin{aligned} \max_{\ddot{\mathbf{Q}},\varepsilon} \quad & \varepsilon \\ \mathbf{A}^T\ddot{\mathbf{Q}} \quad & \geq \quad \varepsilon\mathbf{e} - \dot{\mathbf{Q}}^T\mathcal{H}\dot{\mathbf{Q}}, \end{aligned} \tag{5.14}$$

the dual of which is

$$\begin{aligned} \min_{\mathbf{f}} \quad & \left(\dot{\mathbf{Q}}^T\mathcal{H}\dot{\mathbf{Q}}\right)^T\mathbf{f} \\ \mathbf{A}\mathbf{f} \quad & = \quad \mathbf{0} \\ \mathbf{e}^T\mathbf{f} \quad & = \quad 1 \\ \mathbf{f} \quad & \geq \quad \mathbf{0}, \end{aligned} \tag{5.15}$$

where the common optimal objective function is

$$\varepsilon^* = \left(\dot{\mathbf{Q}}^T\mathcal{H}\dot{\mathbf{Q}}\right)^T\mathbf{f} = \dot{\mathbf{Q}}^T\left(\mathcal{H}\mathbf{f}\right)\dot{\mathbf{Q}} = \dot{\mathbf{Q}}^T\mathbf{H}\dot{\mathbf{Q}},$$

where $\mathbf{H} = \mathcal{H}\mathbf{f}$ is a form of reduced Hessian that incorporates information about the contact force and the curvature of the touching particles. The $[N_f \times N_f]$ matrix $\mathbf{H}$ plays an essential role in the theory of jamming for hypostatic ellipsoid packing and we will refer to it as the *stress matrix* following Ref. [73].

The stress-matrix has a special block structure, where all of the blocks are of size $[d_f \times d_f]$, and both the block-rows and the block-columns correspond to particles. The block entry corresponding to the pair of particles (i, j) is nonzero if and only if there is a contact between them. Written explicitly, the stress matrix is a force-weighted sum of contributions from all the contacts

$$\mathbf{H} = \sum_{\{i,j\}} f_{ij}\mathbf{H}_{ij},$$

where the contribution from a given contact $\{i, j\}$ is

$$\mathbf{H}_{ij} = \begin{array}{c} \\ i \rightarrow \\ \vdots \\ j \rightarrow \end{array} \begin{array}{c} \begin{array}{ccc} i & \cdots & j \\ \downarrow & \cdots & \downarrow \end{array} \\ \begin{bmatrix} \boldsymbol{\nabla}^2_{ii}\zeta_{ij} & \cdots & \boldsymbol{\nabla}^2_{ji}\zeta_{ij} \\ \vdots & \ddots & \vdots \\ \boldsymbol{\nabla}^2_{ij}\zeta_{ij} & \cdots & \boldsymbol{\nabla}^2_{jj}\zeta_{ij} \end{bmatrix} \end{array}. \tag{5.16}$$

If $\dot{\mathbf{Q}}^T\mathbf{H}\dot{\mathbf{Q}} < 0$ then $\varepsilon^* < 0$ and therefore the first-order flex $\dot{\mathbf{Q}}$ cannot be extended into a second-order flex. We say that the stress matrix *blocks* the flex $\dot{\mathbf{Q}}$. If on the other hand $\dot{\mathbf{Q}}^T\mathbf{H}\dot{\mathbf{Q}} > 0$, then $\varepsilon^* > 0$ and by solving the LP (5.14) we can find an unjamming motion, i.e., the packing is second-order flexible. Therefore, finding an unjamming motion at the second-order level essentially consists of looking for a null first-order flex (floppy mode) $\dot{\mathbf{Q}}$, $\mathbf{A}^T\dot{\mathbf{Q}} = \mathbf{0}$, that is also a positive curvature vector for the stress matrix.

Recalling that every floppy mode can be expressed as $\dot{\mathbf{Q}} = \mathbf{V}\mathbf{x}$ [c.f. Eq. (5.12)], we see that

$$\dot{\mathbf{Q}}^T\mathbf{H}\dot{\mathbf{Q}} = \mathbf{x}^T\left(\mathbf{V}^T\mathbf{H}\mathbf{V}\right)\mathbf{x} = \mathbf{x}^T\mathbf{H}_V\mathbf{x}. \tag{5.17}$$

If the matrix $\mathbf{H}_V$ is negative-definite, than the packing is second-order jammed. In Ref. [73] such packings are called *pre-stress stable*, since the self-stress $\mathbf{f}$ rigidifies the packing (i.e., blocks all of the floppy modes).

If $\mathbf{H}_V$ is indefinite, than the packing is second-order flexible since any of the positive-curvature directions can be converted into a strict self-stress by solving the LP (5.14).

If a packing has more than one (proper) self-stress, than it is not clear which one to use in the stress-matrix. One can try to find a self-stress that provides for jamming (pre-stress stability) by looking for a solution to Eq. (5.9) such that $\mathbf{H}_V \preceq \mathbf{0}$ (i.e., $\mathbf{H}_V$ is negative-semidefinite). This is known as semidefinite programming (SDP), and is a powerful generalization of linear programming that has received lots of attention recently [172]. It is however possible that different self-stresses are needed to block different portions of the space of floppy modes, and this general case of a second-order jammed packing is difficult to test for algorithmically. In our study of disordered sphere and ellipsoid packings, we will see that in practice the jammed packings only have one strict self-stress. In this case, testing for jamming reduces to calculating the smallest eigenvalue of $\mathbf{H}_V$. We will discuss actual numerical algorithms designed for ellipsoid packings in subsequent sections, but first we explain what makes sphere packings special.

5.5.2.2 The Stress Matrix for Hard Spheres

For hard spheres it is easy to write down the explicit form for $\mathbf{H}_{ij}$ since the overlap function is given explicitly by Eq. (2.3) and its second-order derivatives are trivial,

$$\boldsymbol{\nabla}^2_{ii} F_{ij} = \boldsymbol{\nabla}^2_{jj} F_{ij} = -\boldsymbol{\nabla}^2_{ij} F_{ij} = -\boldsymbol{\nabla}^2_{ji} F_{ij} = \frac{2\mathbf{I}_d}{(O_i + O_j)^2},$$

where $\mathbf{I}_d$ is the $[d \times d]$ identity matrix. This implies that $\mathbf{H}_{ij}$ is a positive-definite matrix, since

$$\dot{\mathbf{R}}^T \mathbf{H}_{ij} \dot{\mathbf{R}} = (\dot{\mathbf{r}}_i - \dot{\mathbf{r}}_j)^T (\dot{\mathbf{r}}_i - \dot{\mathbf{r}}_j) \geq 0.$$

Therefore, any first-order flex in fact represents a true unjamming motion, since $\dot{\mathbf{Q}}^T \mathcal{H} \dot{\mathbf{Q}} \geq \mathbf{0}$ and we can trivially use $\ddot{\mathbf{Q}} = \mathbf{0}$ in Eqs. (5.14). In other words, a sphere packing is jammed *if and only if it is first-order jammed*, and therefore it cannot be hypostatic. To test for jamming in hard-sphere packings we need only focus on the velocities of the sphere centroids and associated linear programs in Section 5.5.1. This important conclusion was demonstrated using a simple calculation in Ref. [71].

For general particle shapes, however, $\mathbf{H}_{ij}$ may be indefinite for some contacts, and testing for jamming may require considering second-order terms. If one considers general convex particle shapes but freezes the orientations of the particles, the packing will behave like a hard-sphere packing. In particular, a jammed packing of nonspherical particles must have at least as many contacts as the corresponding isostatic packing of spheres would, that is,

$$\bar{Z} \geq 2d$$

for any large jammed packing of convex hard particles.

5.5.3 Testing for Jamming in Ideal Packings

We now summarize the theoretical conditions for jamming developed in this section in the form of a procedure for testing whether a given packing of non-spherical particles is jammed. We assume that the contact network of the packing is known and available as input. For spherical particles, as already discussed, second-order terms never need to be considered, and testing for jamming can be done by solving one or two linear programs, as discussed in detail in Ref. [71]. In the formulation below, we avoid solving linear programs unless necessary, but rather use basic linear algebra tools whenever possible.

1. Find a basis $\mathbf{F}$ for the null-space of the rigidity matrix $\mathbf{A}$, i.e., find N_{stresses} linearly independent solutions to the linear system of equations $\mathbf{Af} = \mathbf{0}$, normalized to mean of unity. This can be done, for example, by looking for zero eigenvalues and the associated eigenvectors of the matrix $\mathbf{A}^T\mathbf{A}$. If
   (a) $N_{\text{stresses}} = 0$,
   (b) $N_{\text{stresses}} = 1$ but the unique self-stress is not non-negative, or
   (c) $N_{\text{stresses}} > 1$ but the linear feasibility program (5.9) is infeasible,

then declare the packing not jammed (first-order flexible), optionally identify an unjamming motion by solving the linear feasibility program $\mathbf{A}^T\dot{\mathbf{Q}} \geq \mathbf{e}$, and terminate the procedure. Otherwise, if the identified self-stress $\mathbf{f}$ is not strict, declare the test inconclusive and terminate.

2. If $N_{\text{floppy}} = N_f + N_{\text{stresses}} - M = 0$, then declare the packing (first-order) jammed and terminate the procedure. Otherwise, find a basis $\mathbf{V}$ for the null-space of $\mathbf{A}^T$, i.e. N_{floppy} linearly-independent solutions to the linear system of equations $\mathbf{A}^T\Delta\mathbf{Q} = \mathbf{0}$. Compute the stress matrix $\mathbf{H}$ using the previously-identified strict self-stress $\mathbf{f}$, and compute its projection $\mathbf{H}_V$ on the space of null first-order flexes.

3. Compute the smallest eigenvalue $\lambda_{\min}$ and associated eigenvector $\mathbf{x}_{\min}$ of the matrix $\mathbf{H}_V$. If $\lambda_{\min} < 0$, declare the packing (second-order) jammed and terminate the procedure. If $\lambda_{\min} > 0$ and $N_{\text{stresses}} = 1$ declare the packing not jammed (second-order flexible), optionally compute an unjamming motion by solving the LP (5.14) with $\dot{\mathbf{Q}} = \mathbf{V}\mathbf{x}_{\min}$, and terminate the procedure. Otherwise, declare the test inconclusive and terminate.

We will discuss the actual numerical implementation of this algorithm later, and see that in practice we do not need to solve linear programs to test for jamming in hypostatic ellipsoid packings. Essentially, the packings we encounter in our work with disordered packings of hard ellipsoids always have a single strict self-stress and a negative-definite $\mathbf{H}_V$.

5.5.4 An Example: Rectangular Lattice of Ellipses

In this section we consider a simple example of a jammed hypostatic packing of ellipses having $\bar{Z} = 4$, the minimum necessary for jamming even for disks. Namely, the rectangular lattice of ellipses, i.e., the stretched version of the square lattice of disks, is collectively jammed, and in particular, it is second-order jammed. More specifically, freezing all but a finite subset of the particles, the remaining packing is second-order jammed. At first glance, it appears that one can rotate any of the ellipses arbitrarily without introducing overlap. However, this is only true up to first order, and at the second-order level the "flat" contacts between the ellipses, that is, the contacts whose normals are along the small ellipse semiaxes, block this rotation through the curvature of the particles at the point of contact.

Consider a finite subpacking of the infinite packing that contains N ellipses ($3N$ degrees of freedom) and has M contacts between ellipses both in the subpacking, and M' contacts between an ellipse in the subpacking and a fixed (boundary) ellipse. That is, the total number of constraints (contacts) is $N_c = M + M'$. Also note that $M' + 2M = 4N$, since there are on exactly 4 contacts per ellipse. The set of self-stresses is obvious in this (sub)packing: Each row and column of ellipse contacts carries an equal force, which can be independent of the force carried by other rows or columns. That is, the dimension of the space of strict self-stresses is $N_s = M'/2$, since every boundary contact is paired with another boundary contact. Also obvious is the set of first-order flexes, i.e., particle motions which preserve contact distances to first order. Namely, a basis vector for this set is a single ellipse rotating around its centroid. The total number of first-order flexes can be found by the Maxwell counting, $N_f = 3N - N_c + N_s = N$, and so by rotating each particle around its centroid we have a basis for the full linear space of first-order flexes. This observation simplifies the calculation of the stress matrix, or more specifically, of the matrix $\mathbf{H}_V$; we only need to consider ellipsoid rotations without considering translations.

The above counting argument applies whenever one takes a jammed sphere packing and makes the particles nonspherical but does not change the normal vectors at the point of contact. This can be done, for example, by simply taking a jammed sphere packing and swelling the particles to be nonspherical, without changing the geometry or connectivity of the contact network, as illustrated for disks in Fig. 5.6. If the particles swell enough to make all of the contacts sufficiently flat, the new packing will be jammed, since all of the first-order flexes consist of particle rotations only and are blocked by the flat curvature of the contacts. In order to see this analytically, we will consider the case of one ellipse jammed among four fixed ellipses (two horizontally, two vertically). The same approach can give the full stress matrix and demonstrate that $\mathbf{H}_V$ is negative-definite, although the proof is more involved and we do not give it here.

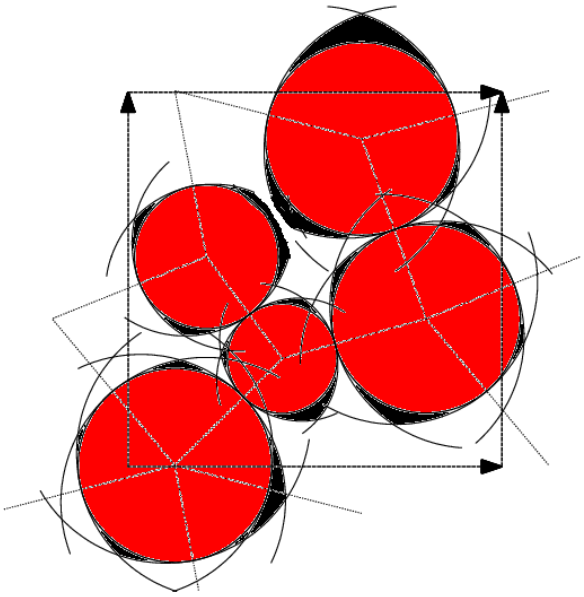

Figure 5.6: A jammed packing of hard disks (colored red) is converted into a jammed packing of nonspherical particles by simply changing the curvature of the particle surfaces at the point of contact to be sufficiently flat (the enlargements of the disks are colored black), without changing the contact network or contact forces.

5.5.4.1 Stress Matrix for a Single Ellipse

What we need to calculate is the Hessian $\partial^2\zeta/\partial\theta^2$, where ζ is the PW overlap potential and θ is the angle (orientation) of the central ellipse. For this purpose, we can locally approximate the two ellipses in contact with circles that are tangent to the ellipses at the point of contact and have a radius equal to the radius of curvature of the ellipses. Let this radius of curvature be $R = b^2/a$ for the mobile ellipse, and $R' = R$ for the fixed ellipse. Let the origin of the coordinate system be at the center of the mobile ellipse, and the initial position of the center of the tangent circle be $\mathbf{r}_0 = \langle l, 0\rangle$ for the mobile ellipse and $\mathbf{r}' = \langle l + R + R', 0\rangle$ for the fixed ellipse. As the mobile ellipse rotates by an angle $\Delta\theta$ and is scaled by a factor μ, we have $\mathbf{r} = \mu l \langle \cos\Delta\theta, \sin\Delta\theta\rangle$. The condition that the two circles be in external tangency is

$$\|\mathbf{r} - \mathbf{r}'\|^2 = (\mu R + R')^2.$$

This equation can be solved (explicitly or just implicitly) to give $\mu(\Delta\theta)$, and then $\Delta\zeta = \mu(\Delta\theta)^2 - 1$ calculated to second order to give $\partial^2\zeta/\partial\theta^2$. This calculation gives the result

$$\frac{\partial^2\zeta}{\partial\theta^2} = \frac{1-\alpha^4}{2\alpha^2},$$

where $\alpha = b/a$ is the aspect ratio.

If the contact we are considering is between curved tips, then $\alpha < 1$, and $\partial^2\zeta_{\text{curv}}/\partial\theta^2 > 0$. For the contact between the flat tips, we just replace α with α^{-1} and get that $\partial^2\zeta_{\text{flat}}/\partial\theta^2 = -\partial^2\zeta_{\text{curv}}/\partial\theta^2 < 0$. The product $\mathbf{x}^T\mathbf{H}_V\mathbf{x}$ in Eq. (5.17) simply becomes

$$\Delta_2\zeta = \left(f_{\text{flat}}\frac{\partial^2\zeta_{\text{flat}}}{\partial\theta^2} + f_{\text{curv}}\frac{\partial^2\zeta_{\text{curv}}}{\partial\theta^2}\right)\Delta\theta^2,$$

where $\mathbf{f}$ is the self-stress under consideration. Therefore, as long as $f_{\text{flat}} > f_{\text{curv}}$, the mobile ellipse is jammed, more specifically, pre-stress rigid. This result can be shown to apply to the square lattice of ellipses for an arbitrary number of ellipses. As we will see in Section 5.8, if the ellipses are not hard but rather deformable,

the packing would not support a compression along the curved contacts, but it would along the flat contacts. This is a very intuitive result: If one takes a smooth ellipsoid and presses it against a table with its most curved tip, it will buckle and the only stable configuration is one where the flat tip presses against the table. Note however that the hard-ellipse equivalent is jammed and can resist any finite external forces, including a compression along the curved contacts. The anharmonicity of the hard-sphere potential becomes essential in this example, since the packing can choose the correct internal (self) stresses (forces) needed to provide mechanical rigidity. In systems of deformable particles, the internal stresses are fixed and determined by the state of compression.

5.5.5 Outside the Kinematic Perspective

It is worthwhile to briefly consider the connections between the jamming criteria developed above using the kinematic approach to jamming, and the static and perturbation approaches.

5.5.5.1 Static View

We have already seen that forces appear naturally as Lagrange multipliers corresponding to impenetrability constraints, in the form of a strict self-stress $\mathbf{f} > 0$. In the static view, we ask whether a packing can support a given applied external force $\mathbf{B}$ by a set of non-negative interparticle forces. The key observation is that we can add an arbitrary positive multiple of a self-stress to any set of interparticle forces that support $\mathbf{B}$ in order to make them non-negative, without affecting force balance. Therefore, if the rigidity matrix $\mathbf{A}$ is of full-rank, as it has to be for jammed sphere packings, any (supportable) load $\mathbf{B}$ can be balanced with non-negative interparticle forces, and kinematic and static rigidity become equivalent [138].

For jammed hypostatic ellipsoid packings however, such as the one in Fig. 5.5, supporting some loads may require a small deformation of the packing, such as a slight rotation of the mobile ellipse in the example in Fig. 5.5. After this small deformation, the normal vectors at the points of contact will change slightly and the interparticle forces $\mathbf{f}$ can support the applied force $\mathbf{B}$. The larger the magnitude of the forces is, the smaller the deformation needed to support the load is. Therefore every jammed packing can support any applied force in a certain generalized sense. Another way to look at this is to observe that, if the interparticle forces are much larger than the applied ones, the applied load will act as a small perturbation to the packing and the static view becomes equivalent to the perturbation view (with $\Delta\boldsymbol{\zeta} = \mathbf{0}$). We consider the perturbation view next and show how the stress matrix appears in the response of the packing to perturbations.

5.5.5.2 Perturbation View

In the perturbation view we consider how the configuration *and* the contact forces respond to perturbations consisting of small changes of the contact geometry and small applied forces. Counting geometric and force constraints separately, as done in the literature, is incorrect when $\mathbf{f} > 0$: There is coupling between the particle positions and the interparticle forces as represented by the Hessian $\mathbf{H} = \mathcal{H}\mathbf{f}$.

With this in mind, we can expand Eq. (5.4) to first order in $\{\|\Delta\mathbf{Q}\|, \|\Delta\mathbf{f}\|\}$, to get the linear system of equations

$$\begin{bmatrix} \mathbf{A} & -\mathbf{H} & 0 \\ \mathbf{0} & \mathbf{A}^T & -2\mathbf{e} \\ \mathbf{e} & \mathbf{0} & 0 \end{bmatrix} \begin{bmatrix} \Delta\mathbf{f} \\ \Delta\mathbf{Q} \\ \Delta\mu \end{bmatrix} = -\varepsilon \begin{bmatrix} \Delta\mathbf{B} \\ \Delta\boldsymbol{\zeta} \\ 0 \end{bmatrix}. \tag{5.18}$$

It can be demonstrated that if the reduced Hessian $\mathbf{H}_V$ is definite, this system will have a solution for any $\Delta\mathbf{b}$ and $\Delta\boldsymbol{\zeta}$. Furthermore, if $\mathbf{H}_V$ is negative-definite the response to perturbations will be stable, in the sense that applied forces will do a positive work in order to perturb the packing.

Equation (5.18) can be used to find the jamming point starting with a packing that is nearly jammed, i.e., a packing that has nonzero interparticle gaps $\varepsilon\Delta\boldsymbol{\zeta}$ and a self-stress that has a small imbalance $\varepsilon\Delta\mathbf{B} = \mathbf{Af}$. This works well for small packings, however, for large disordered packings, the force chains are very sensitive to small changes in the geometry and the linearization of the perturbation response is not a good approximation even for packings very close to the jamming point. Additionally, we note that to first order in ε, the solution to Eq. (5.18) has

$$\Delta\mu/\varepsilon = \mathbf{f}^T\Delta\boldsymbol{\zeta}/2M = \mathbf{f}_E^T\mathbf{h}/2M \tag{5.19}$$

, which can be used to quickly estimate the jamming gap of a nearly-jammed packing from just the interparticle gaps $\Delta\boldsymbol{\zeta} = \boldsymbol{\zeta}$ and the interparticle forces, without knowing the actual jamming point [75].

5.6 Numerically Testing for Jamming in Hypostatic Ellipsoid Packings

In this section we will apply the criteria for jamming and the algorithm to test for jamming from Section 5.5.3 to our computationally-generated hypostatic packings of ellipsoids. The numerical results show that the packings are indeed second-order jammed, even very close to the sphere point.

We derived the necessary overlap potentials and its derivatives for hard ellipses and ellipsoids in Chapter 2. In particular, the stress-matrix is built from blocks as given in Eq. (5.16), where each of the four blocks $\boldsymbol{\nabla}^2_{\alpha\beta}\zeta$ (α and β denote either A or B) involves both translations and rotations,

$$\boldsymbol{\nabla}^2_{\alpha\beta}\zeta = \begin{bmatrix} \boldsymbol{\nabla}^2_{\mathbf{r}_\alpha}\zeta & \boldsymbol{\nabla}^2_{\boldsymbol{\varphi}_\alpha \mathbf{r}_\beta}\zeta \\ \boldsymbol{\nabla}^2_{\mathbf{r}_\alpha \boldsymbol{\varphi}_\beta}\zeta & \boldsymbol{\nabla}^2_{\mathbf{r}_\beta}\zeta \end{bmatrix},$$

and the necessary second-order derivatives are given in Section 2.2.4.2.

The numerical implementation of the algorithm given in Section 5.5.3 poses several challenges. The most important issue is that that algorithm was designed for ideal packings, that is, it was assumed that the true contact network of the packing is known. Packings produced by the MD algorithm, although very close to jamming (i.e., very high pressures), are not ideal. In particular, it is not trivial to identify which pairs of particles truly touch at the jamming point. Disordered packings have a multitude of near contacts that play an important role in the rigidity of the packing away from the jamming point [173], and these near contacts can participate in the backbone (force-carrying network) even very close to the jamming point. Additionally, not including a contact in the contact network can lead to the identification of spurious unjamming motions, which are actually blocked by the contact that was omitted in error.

For hard spheres, the algorithms can use linear programming to handle the inclusion of false contacts [71]. For ellipsoids, we look at the smallest eigenvalues of $\mathbf{A}^T\mathbf{A}$, i.e., the least-square solution to $\mathbf{Af} = \mathbf{0}$. The solution will be positive if we have identified the true contact network, $\mathbf{f} > 0$, but the inclusion of false contacts will lead to small negative forces on those false contacts. The problem comes about because the calculation of the self-stress by just looking at the rigidity matrix does not take into account the actual proximity to contact between the particles. One way to identify the true contact network of the packing is to perform a long molecular dynamics run at a fixed density at the highest pressure reached, and record the list of particle neighbors participating in collisions as well as average the total transfer of collisional momentum between them in order to obtain the (positive) contact forces [75].

Once the contact network is identified, we want to look for null-vectors of the rigidity matrix. This can be done using specialized algorithms that ensure accurate answers [174], however, we have found it sufficient in practice to simply calculate the few smallest eigenvalues of the semi-definite matrix $\mathbf{A}^T\mathbf{A}$. We used MATLAB's sparse linear algebra tools to perform the eigenvalue calculation (internally MATLAB uses the ARPACK library, which implements the Implicitly Restarted Arnoldi Method). We consistently found that the smallest eigenvalue is about $3-6$ orders of magnitude smaller than the second-smallest eigenvalue, indicating that there is a near linear-dependency among the columns of $\mathbf{A}$ in the form of a self-stress. The self-stress, which is simply the eigenvector corresponding to the near-zero eigenvalue, was always strictly positive; in our experience, disordered packings of ellipsoids have a *unique strict self-stress* $\mathbf{f}$. This means that there are $N_{\text{floppy}} = N_f + 1 - M$ solutions to $\mathbf{A}^T\Delta\mathbf{Q} = \mathbf{0}$, $N_f - M$ of which are exact, and one which is approximate (corresponding to the approximate self-stress). This can be seen, for example, by calculating the eigenvalues of $\mathbf{AA}^T$, since $N_f - M$ will be zero to numerical precision, one will be very small, and the remaining ones will be orders of magnitude larger.

5.6.1 Verification of Second-Order Jamming

Once a strict self-stress is known, second-order jamming or flexibility can be determined by examining the smallest eigenvalue of $\mathbf{H}_V$, which requires finding a basis for the linear space of floppy modes. However, it

is computationally demanding to find a basis for the null-space of $\mathbf{A}^T$ due to the large number of floppy modes, and since sparsity is difficult to incorporate in null-space codes. There are algorithms to find sparse basis for this null-space [174], however, we have chosen a different approach.

Namely, we calculate the smallest eigenvalues of

$$\mathbf{H}_k = k\mathbf{A}\mathbf{A}^T - \mathbf{H},$$

which as we saw in Section 5.8.2 is the Hessian of the potential energy for a system of deformable ellipsoids where the stiffness coefficients are all k. For very large k (we use $k = 10^6$), any positive eigenvalue of $\mathbf{A}\mathbf{A}^T$ is strongly amplified and not affected by $\mathbf{H}$, and therefore only the floppy modes can lead to small eigenvalues of $\mathbf{H}_k$, depending on how they are affected by $\mathbf{H}$. We have found that MATLAB's eigs function is not able to converge the smallest eigenvalues of $\mathbf{H}_k$ for large stiffnesses k, however, the convergence is quick if one asks for the eigenvalues closest to zero or even closest to -1. This typically reveals any negative eigenvalues of $\mathbf{H}_k$ and the corresponding floppy modes.

It is also possible to perform a rigorous numerical test for positive-definiteness of $\mathbf{H}_k$ using properly rounded IEEE machine arithmetic and MATLAB's (sparse) Cholesky decomposition of a numerically reconditioned $\mathbf{H}_k$ [175]. We have used the code described in Ref. [175] to show that indeed for our packings $\mathbf{H}_k \succ \mathbf{0}$ and therefore the packings are second-order jammed. For spheroids, that is, ellipsoids that have an axes of symmetry, there will be trivial floppy modes corresponding to rotations of the particles around their own centroid. These can be removed most easily by penalizing any component of the particle rotations $\Delta\boldsymbol{\varphi}$ that is parallel to the axis of symmetry. For example, one can add to every diagonal block of $\mathbf{H}_k$ corresponding to the rotation of an ellipsoid with axes of symmetry $\mathbf{u}$ a penalization term of the form $k\mathbf{u}\mathbf{u}^T$.

We have not performed a detailed investigation of a very wide range of samples since our goal here was to simply demonstrate that under appropriate conditions the packings we generate using the modified Lubachevsky-Stillinger algorithm are indeed jammed, even though they are very hypostatic near the sphere point. In this work we have given the fundamentals of the mathematics of jamming in these packings. A deeper understanding of the mechanical and dynamical properties of nearly-jammed hypostatic ellipsoid packings is a subject for future work.

5.6.2 Extensions to Non-Ideal Packings

In Chapter 4 we used mathematical programming, and in particular, linear programming, in order to find unjamming motions even in the presence of interparticle gaps, i.e., when considering non-ideal packings. The LP formulation identifies the set of contacts that support a given external load $\mathbf{B}$, i.e., the set of *active contacts*, even in the presence of gaps. If we apply the same type of randomized algorithm to test for jamming in packings of nonspherical particles, we would have to include at least the quadratic terms in the impenetrability expansion, to obtain the primal quadratic programming (QP) problem

$$\begin{aligned} & \max_{\Delta\mathbf{Q}} \mathbf{B}^T\Delta\mathbf{Q} && \text{for virtual work} \\ \text{such that} \quad & \widetilde{\Delta\boldsymbol{\zeta}} = \mathbf{A}^T\Delta\mathbf{Q} + \Delta\mathbf{Q}^T\mathcal{H}\Delta\mathbf{Q} + \Delta\boldsymbol{\zeta} \geq \mathbf{0} && \text{for impenetrability,} \end{aligned} \tag{5.20}$$

where $\Delta\boldsymbol{\zeta} \geq \mathbf{0}$ are the interparticle gaps. The first-order Karush-Kuhn-Tucker (KKT) optimality conditions for this QP problem are essentially Eq. (5.18) restricted to the set of active contacts

$$\begin{aligned} \mathbf{A}\mathbf{f} + \mathbf{H}\Delta\mathbf{Q} &= -\mathbf{B} \\ \mathbf{f} &\geq \mathbf{0} \\ \mathbf{f}^T\widetilde{\Delta\boldsymbol{\zeta}} &= 0, \end{aligned} \tag{5.21}$$

while the second-order sufficient conditions require a strict self-stress, $\mathbf{f} + \widetilde{\Delta\boldsymbol{\zeta}} > \mathbf{0}$, and also that all first-order flexes $\dot{\mathbf{Q}}$ of the active contact network be blocked by the stress matrix, $\dot{\mathbf{Q}}^T\mathbf{H}\dot{\mathbf{Q}} < 0$, as we already derived.

The problem, however, is that non-convex quadratic programs like that in Eq. (5.20), unlike linear programs, are not solvable in polynomial time, and in fact, it is rather difficult to solve them numerically for large packings using general-purpose quadratic solvers. It may be possible to design special methods that use the pairwise structure of the constraints together with explicit separation of the variables into translational and orientational degrees of freedom, and efficiently solve the QP (5.20). This is a subject for further research.

5.7 Nearly Jammed Packings

So far we have considered ideal jammed packings, where particles are exactly in contact. Computer-generated packings however always have a packing fraction ϕ slightly lower than the jamming packing fraction ϕ_J, and the particles can rattle (move continuously) to a certain degree if agitated thermally or by shaking [75]. We can imagine that we started with the ideal jammed packing and scaled the particle sizes by a factor $\mu = 1 - \delta < 1$, so that the packing fraction is lowered to $\phi = \phi_J (1 - \delta)^d$. We call δ the *jamming gap* or distance to jamming.

It can be shown that if δ is sufficiently small the rattling of the particles does not destroy the jamming property, in the sense that the configuration point $\mathbf{Q} = \mathbf{Q}_J + \Delta\mathbf{Q}$ remains trapped in a small *jamming neighborhood* or *jamming basin* $\mathcal{J}_{\Delta\mathbf{Q}} \subset \mathcal{R}^{N_f}$ around $\mathbf{Q}_J$ [165]. In the limit $\delta \to 0$ the set of accessible configurations $\mathcal{J}_{\Delta\mathbf{Q}} \to \{\mathbf{Q}_J\}$, and in fact this is the definition of jamming used by Salsburg and Wood in Ref. [52]. Rewritten to use our terminology, this definition is: "A configuration is *stable* if for some range of densities slightly smaller than ϕ_J, the configuration states *accessible from* $\mathbf{Q}_J$ lie in the neighborhood of $\mathbf{Q}_J$. More formally, if for any small $\epsilon > 0$ there exists a $\delta > 0$ such that all points $\mathbf{Q}$ accessible from $\mathbf{Q}_J$ satisfy $\|\mathbf{Q} - \mathbf{Q}_J\| < \epsilon$ provided $\phi \geq \phi_J (1 - \delta)^d$." We call this the *trapping* view of jamming, most natural one when considering the thermodynamics of nearly jammed hard-particle systems [82]. Note that the trapping definition of jamming is in fact equivalent to our kinematic definition of jamming [165], and we actually motivated the concept of jamming by considering trapping in Chapter 4.

To illustrate the influence of the constraint curvature on jamming, we show in Fig. 5.7 four different cases with two constraints in two dimensions. In all cases a self-stress exists since the normals of the two constraints are both horizontal. If both constraint surfaces are concave (have negative or outward curvature), as constraints always are for hard-spheres, two constraints cannot close a bounded region $\mathcal{J}_{\Delta\mathbf{Q}}$ around the jamming point. One needs at least three constraints and in that case $\mathcal{J}_{\Delta\mathbf{Q}}$ will be a curved triangle. If however at least one of the constraints is convex (has positive curvature), two constraints can bound a closed jamming basin. Specifically, if the sum of the radii of curvatures of the two constraints at the jamming point $R_1 + R_2$ is positive, there is no unjamming motion, however, if it is negative then there is an unjamming motion in the vertical (floppy) direction. This is equivalent to looking at the smallest eigenvalue of the stress matrix in higher dimensions.

The jamming basin $\mathcal{J}_{\Delta\mathbf{Q}}(\delta)$ for a given jamming gap δ is the local solution to the relaxed impenetrability equations

$$\boldsymbol{\zeta}(\Delta\mathbf{Q}) \geq -\zeta_\delta = 1 - \left(\frac{1}{1-\delta}\right)^2.$$

One way to determine $\mathcal{J}_{\Delta\mathbf{Q}}(\delta)$ for a wide range of δ's is to consider the function of the particle displacements

$$\tilde{\delta}(\Delta\mathbf{Q}) = \sqrt{1 + \min[\boldsymbol{\zeta}(\Delta\mathbf{Q})]} - 1, \tag{5.22}$$

that is, to calculate by how much the particles need to be shrunk to make a given particle displacement $\Delta\mathbf{Q}$ feasible (preserving non-overlapping). The contours (level-sets) of the function $\tilde{\delta}(\Delta\mathbf{Q})$ denote the boundaries of $\mathcal{J}_{\Delta\mathbf{Q}}(\delta)$, that is, $\mathcal{J}_{\Delta\mathbf{Q}}(\delta) = \left\{\Delta\mathbf{Q} \mid \tilde{\delta}(\Delta\mathbf{Q}) \leq \delta\right\}$.

5.7.1 First-Order Jammed Packings

As a simple but illustrative example, we will consider a single mobile disk jammed between three other stationary disk, as shown in Fig. 5.8, an analog of the ellipse example from Fig. 5.5. This packing is first-order jammed, and the figure also shows a color plot of the function $\tilde{\delta}(\Delta\mathbf{Q})$ along with its contours. It is seen that for small δ the jamming basin $\mathcal{J}_{\Delta\mathbf{Q}}$ is a closed curved triangle.

These observations are readily generalized to higher dimensions. For sufficiently small δ, the jamming basin approaches a convex *jamming polytope* (a closed polyhedron in arbitrary dimension) $\mathcal{P}_{\Delta\mathbf{Q}}$, as discussed in Section 4.3.1.1. For spheres all constraint surfaces are concave and therefore $\mathcal{P}_{\Delta\mathbf{Q}} \subseteq \mathcal{J}_{\Delta\mathbf{Q}}$ [52, 145]. The jamming polytope is determined from the linearized impenetrability equations

$$\mathbf{A}^T \Delta\mathbf{Q} \geq -\zeta_\delta \approx -2\delta, \tag{5.23}$$

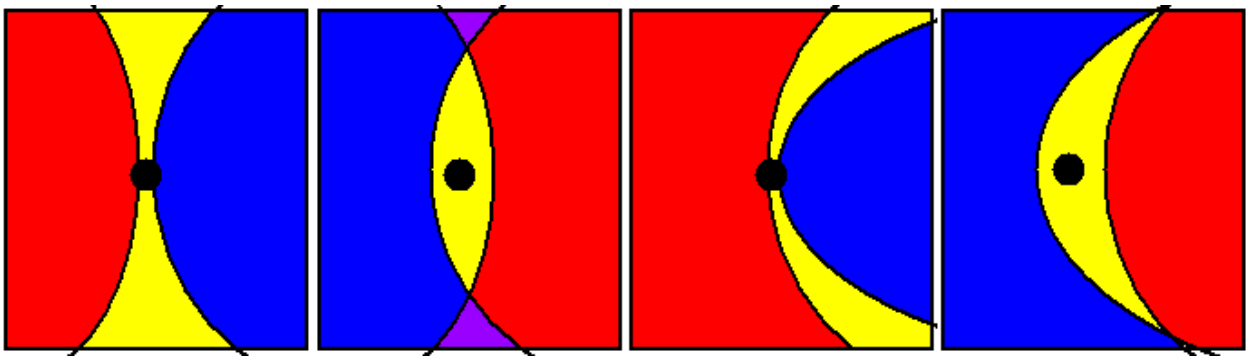

Figure 5.7: The feasible region (yellow) around a jamming point (black circle) for two curved constraints in two dimensions (black circles). The region of the plane forbidden by one of the constraints is colored red, and colored blue for the other constraint. The region forbidden by both constraints is purple. The distance from the jamming point to the constraints is approximately δ and chosen small. Four cases are shown, going from left to right: (a) Both constraints are concave, and the yellow region is not bounded. Moving along the vertical direction unjams the system (this is typical of hard spheres). (b) Both constraints are convex, and the yellow region is closed, even though it is very elongated along the vertical direction (of order $\sqrt{\delta}$). This is an example of pre-stress stability (second-order jamming). (c) One of the constraints is convex, but not enough to block the unjamming motion in the vertical direction. The motion has to curve to avoid the convex constraint, i.e., a nonzero acceleration is needed to unjam the system (second-order flexible). (d) Only one of the constraints is convex, but enough to close the yellow region (second-order jammed). If the radii of curvatures are very close in magnitude, this region can become a very elongated banana-like shape.

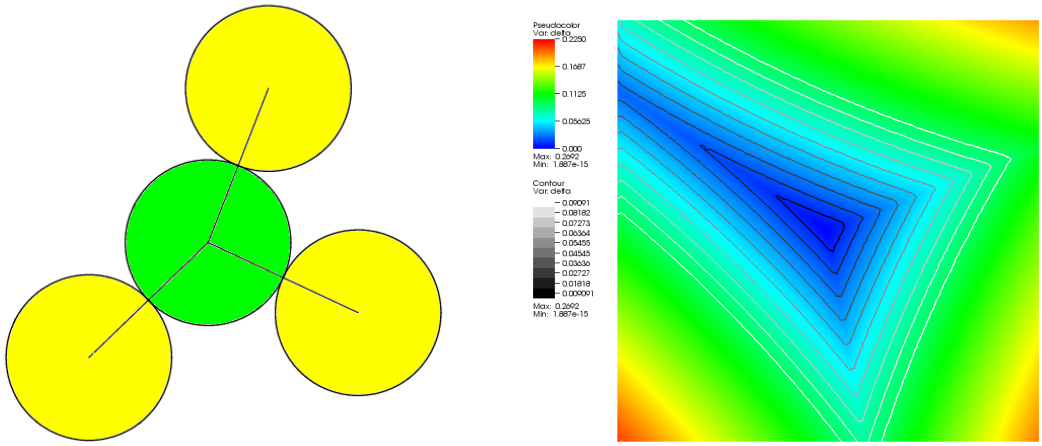

Figure 5.8: (*Left*) An example of a mobile disk (green) jammed between three fixed disks (yellow). This is analogous to the ellipse packing shown in Fig. 5.5. (*Right*) A color plot of the function $\tilde{\delta}(\Delta\mathbf{Q})$ for this disk packing along with its contours (level sets).

and we can see that its volume, which determines the (non-equilibrium) free-energy, scales like δ^{N_f}. This leads to the free-volume divergence of the pressure in the jamming limit

$$p = \frac{PV}{NkT} \approx \frac{d_f}{1 - \phi/\phi_J}, \tag{5.24}$$

which has been verified numerically for disordered isostatic hard sphere packings [75].

5.7.2 Second-Order Jammed Packings

The ellipse analog from Fig. 5.5 has three degrees of freedom, two translational and one orientational. If we fix the orientation of the (mobile) ellipse, that is, we take a planar cut through $\tilde{\delta}(\Delta\mathbf{Q})$, the situation is identical to that for the disk example above: For small δ the jamming basins $\mathcal{J}_{\Delta\mathbf{Q}}$ are closed curved triangles. However, the range of accessible orientations is rather large, on the order of $\sqrt{\delta}$, since even for a small δ the ellipse can rotate significantly. This is a consequence of the rotation of the ellipse being a floppy mode, and only being blocked by second-order effects as given by the curvature of the impenetrability constraints. In a certain sense, the packing is much more trapped in the subspace of configuration space perpendicular to the space of floppy modes than it is in the space of floppy modes. This is illustrated in Fig. 5.9.

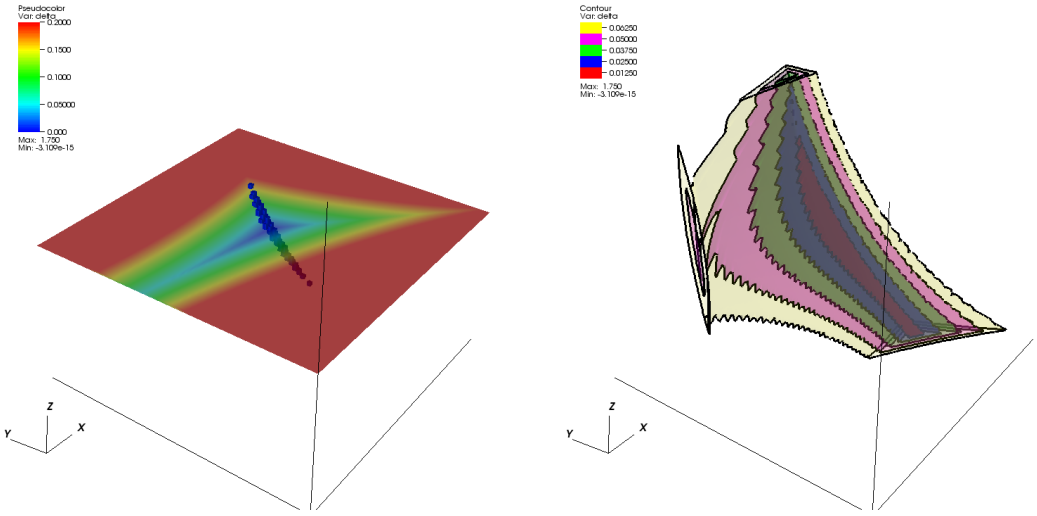

Figure 5.9: (*Left*) A plot of the function $\tilde{\delta}(\Delta\mathbf{Q})$ for the packing from Fig. 5.5. The horizontal axes correspond to the translational degrees of freedom, and the vertical to the rotational degree of freedom (the rotation angle of the major axes). The $\Delta\boldsymbol{\varphi} = \mathbf{0}$ cut is also shown (horizontal colored plane), to be compared to the right part of Fig. 5.8. We also show the jamming basin $\mathcal{J}_{\Delta\mathbf{Q}}(\delta = 0.0035)$ (dark blue region), illustrating that this region is shaped like a banana, elongated along the direction of the floppy mode. (*Right*) Several contours (iso-surfaces) of $\tilde{\delta}(\Delta\mathbf{Q})$, bounding the banana-shaped regions $\mathcal{J}_{\Delta\mathbf{Q}}(\delta)$.

5.7.3 Pressure Scaling for Hypostatic Jammed Ellipsoid Packings

The observations in Fig. 5.9 are readily generalized to higher dimensions, however, it is no longer easy to determine the volume of $\mathcal{J}_{\Delta\mathbf{Q}}$ (and thus the free energy) in the jamming limit. If we consider the simple two-constraint example in Fig. 5.7, we find that the area A of the feasible (yellow) region scales like $\delta^{3/2}$ instead of δ^2,

$$A = \frac{16}{3}\sqrt{\frac{R_1 R_2}{R_1 + R_2}}\delta^{3/2}.$$

An obvious generalization of this result to higher dimensions can be obtained by assuming that the jamming basin $\mathcal{J}_{\Delta\mathbf{Q}}$ has extent $\sqrt{\delta}$ along all $N_{\text{floppy}} \approx N_f - M$ directions corresponding to floppy modes, where as it has extent δ along all other perpendicular directions. The volume would then scale as

$$|\mathcal{J}_{\Delta\mathbf{Q}}| \sim \delta^M \delta^{(N_f - M)/2} = \delta^{N(d_f/2 + \bar{Z}/4)} = \delta^{N d_f (1+s)/2},$$

where we quantify the *hypostaticity* of the packing by $s = \bar{Z}/2d_f$. The corresponding scaling of the pressure in the jamming limit is

$$p = \frac{PV}{NkT} \approx \frac{d_f(1+s)/2}{1 - \phi/\phi_J}.$$

However, as δ becomes very small, the jamming region becomes so elongated along the space of floppy modes that the time-scales for rattling along the elongated directions becomes much larger than the time for rattling in the perpendicular directions. This manifests itself as a remarkably large and regular oscillation of the "instantaneous" pressure (as measured over time intervals of tens of collisions per particle) during molecular-dynamics runs at a fixed δ, as illustrated in Fig. 5.10. The oscillations are more dramatic the smaller δ is, and can span six or more orders of magnitudes of changes in the instantaneous pressure. The period of oscillation, as measured in numbers of collisions per particle, is dramatically affected by the moment of inertia of the ellipsoids I, most naturally measured in units of mO^2, where m is the particle mass and O is the (say smallest) ellipsoid semiaxis.

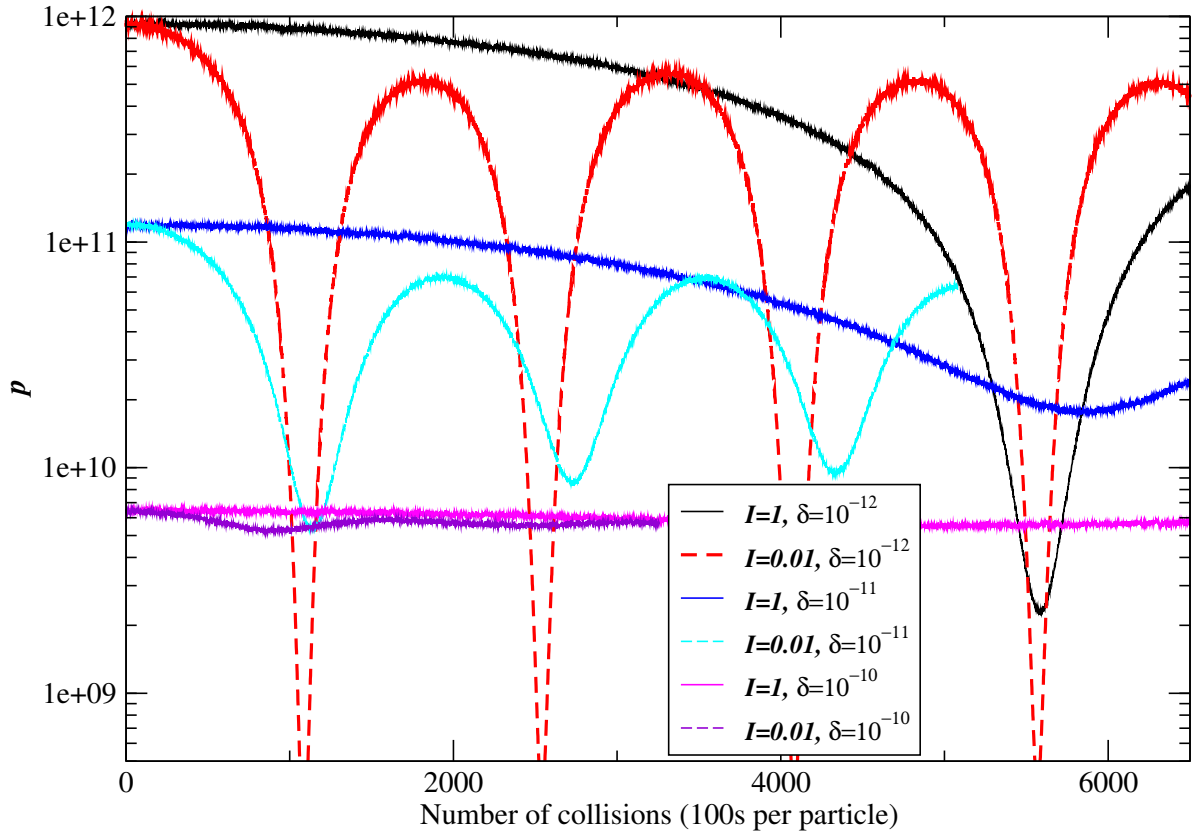

Figure 5.10: The "instantaneous" reduced pressure for a jammed hypostatic packing of three-dimensional ellipsoids with semiaxes ratio $1.025^{-1} : 1 : 1.025$, at different (estimated) distances from the jamming point δ. Molecular dynamics runs using a natural moment of inertia of the particles as well as ones using a much smaller moment of inertia are shown. The pressure oscillations are sustained for very long periods of time, however, it is not clear whether they eventually dissipate.

We do not understand the full details of these pressure oscillations, however, it is clear that dynamics near

the jamming point for the hypostatic ellipsoid packings is not ergodic on small time-scales. In particular, as a packing is compressed during the course of the packing algorithm, the time-scale of the compression may be shorter than the time-scale of exploring the full jamming basin. Over shorter time scales the packing can only explore the directions perpendicular to the floppy modes, and in this case we expect that the pressure would scale as

$$p \approx \frac{d_f s}{1-\phi/\phi_J}.$$

In Fig. 5.11 we show $C = p(1-\phi/\phi_J)$ as a function of the jamming gap for compressions of systems of ellipses of different aspect ratios close to unity. The compression started with a dense liquid and the particles were grown slowly at an expansion rate $\gamma = 10^{-5}$ to a high pressure (jamming) $p = 10^9$. The figure shows for each aspect ratio the lower bound $C_L = d_f s = 3s$ and the upper bound $C_U = d_f(1+s)/2 = 1.5(1+s)$, where s was calculated by counting the almost perfect contacts at the highest pressure [75]. As expected from the arguments above, we see that very close to the jamming point $C \approx C_L$, however, further away from jamming $C \approx C_U$. For packings that are not hypostatic $C_L = C_U = d_f$, and for disks $C_U = C_L = 2$.

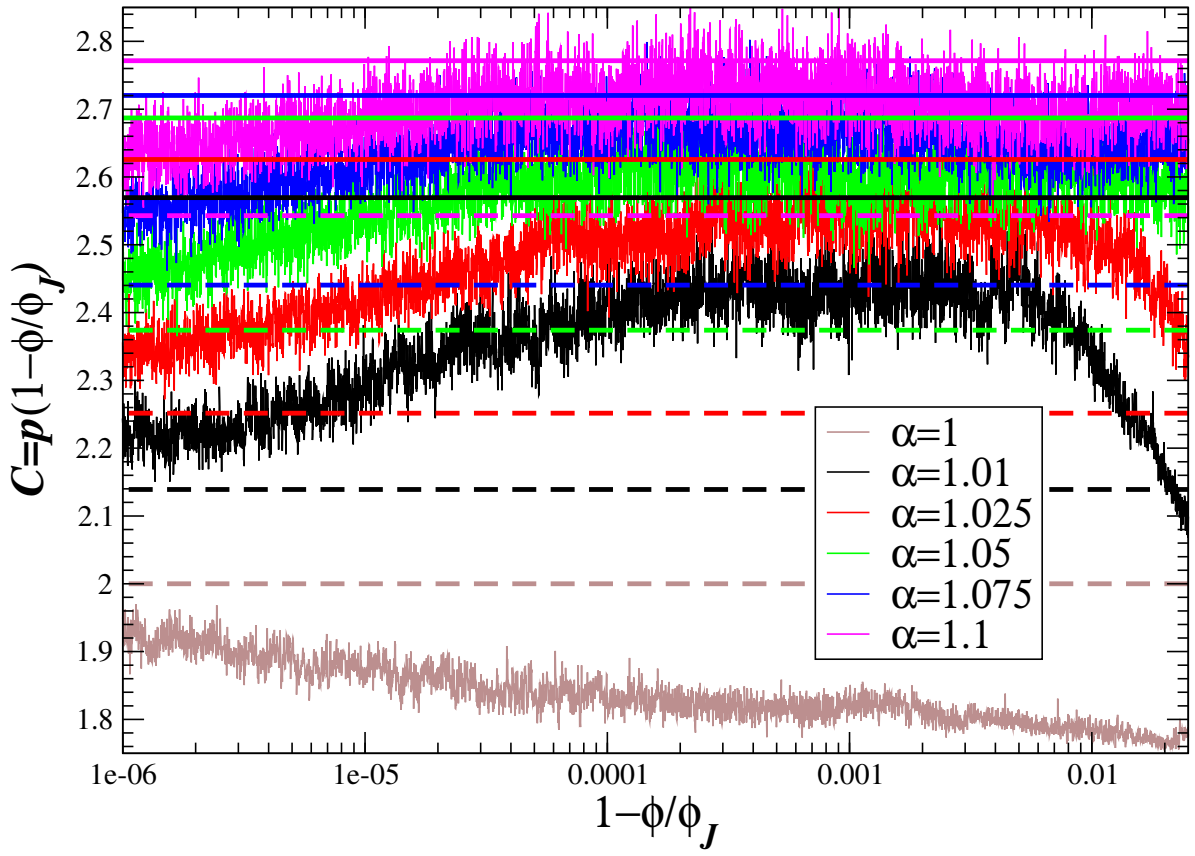

Figure 5.11: The pressure scaling coefficient $C = p(1-\phi/\phi_J)$ as systems of hard ellipses are compressed from a dense liquid to the jamming point. The value of C is not constant however it seems to remain between the bounds C_L (shown with a dashed line in the same color as C) and C_U (shown with a solid line).

5.8 Energy Minima in Systems of Deformable Particles

In this section we consider the connections between jamming in hard particle packings and stable (local) energy minima (inherent structures [29]) for systems of deformable (soft) particles. This has a two-fold purpose. Firstly, in physical systems particles are always deformable, and therefore it is important to

establish that the hard-particle conditions for jamming we established in Section 5.5 are relevant to systems of deformable particles. We expect that if the particles are sufficiently stiff, to be made more quantitative shortly, the behavior of the soft-particle system will approach that of the corresponding hard-particle packing. Secondly, considering the conditions for the existence of a stable energy minimum will enable us to derive in a simpler fashion and better understand the jamming conditions from the previous section.

We consider systems with short-ranged continuous interparticle potentials that are a monotonically decreasing function E of the overlap between particles,

$$U_{ij} = E\left[\zeta\left(\mathbf{q}_i, \mathbf{q}_j\right)\right]. \tag{5.25}$$

That is, we assume that the elastic behavior of the particles is such that the interaction energy only depends on the distance between the particles as measured by the overlap potential ζ. An example of such an elastic potential is an inverse power-law

$$E(\zeta) = (1+\zeta)^{-\nu}, \tag{5.26}$$

which in the limit $p \to \infty$ approaches a hard-particle interaction

$$E_H(\zeta) = \begin{cases} 0 \text{ if } \zeta > 0 \\ \infty \text{ if } \zeta < 0 \end{cases}.$$

For sufficiently large power exponents p the interaction is localized around particles in contact and the overall energy

$$U = \sum_{ij} U_{ij} \to \max_{ij} U_{ij} = \left(1 + \min_{ij} \zeta_{ij}\right)^{-\nu} = \left(1 + \tilde{\delta}^2\right)^{-\nu}$$

is dominated by the most overlapping pair of particles [see Eq. (5.22) for the definition of $\tilde{\delta}$]. Additionally, as p grows the interparticle potential becomes stiff in the sense that small changes in the distance between the particles cause large changes of the interparticle force

$$f = -\frac{dE}{d\zeta} \geq 0,$$

and the stiffness coefficient

$$k = \frac{d^2E}{d\zeta^2} \geq 0$$

becomes very large and positive. This indicates a physical interpretation of the hard-particle interaction potential: It is the limit of taking an infinite stiffness coefficient while the force between particles is kept at some non-negative value, which can be tuned as desired by infinitesimal changes in the distance between the particles (but note that the forces in different contacts are correlated since the motion of particles affects all of them simultaneously).

5.8.1 Stable Energy Minima Correspond to Jammed Packings

Assume that we have a packing of hard particles and that we can find a set of interparticle interaction potentials U_{ij} for the geometric contacts such that the configuration is a stable energy minimum. This means that any motion of the particles leads to increasing the energy U, i.e., to overlap of some pair of particles. Therefore, the packing of hard particles is jammed. This gives a simple way to prove that a given packing is jammed: Find a set of interparticle potentials that makes the configuration a stable energy minimum [165, 73]. We examine the conditions for a stable energy minimum when the interaction potentials are twice differentiable next.

The converse is also true, in the sense that arbitrarily near a jammed packing there is an energy minimum for a sufficiently "hard" interaction potential (in some cases the potential energy U may have to be discontinuous at the origin [73]). We demonstrate this on the examples from Figs. 5.8 and 5.9 for a power-law interaction potential with increasing exponent ν in Figs. 5.12 and 5.13, respectively. It is clear that in the limit $p \to \infty$, the contours of the interaction potential become those of $\tilde{\delta}(\Delta\mathbf{Q})$ and are thus closed near the

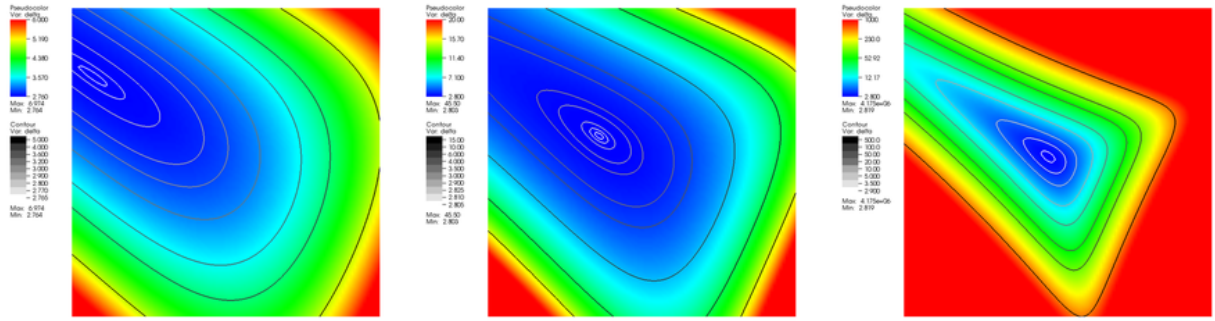

Figure 5.12: The total interaction energy $U(\Delta\mathbf{Q})$ for the example in Fig. 5.8 when the disks are deformable and interact via a power-law potential. We show U as a color plot with overlayed contours for power exponents $\nu = 12, 25$, and 100 (going from left to right). Compare the $\nu = 100$ case to the contours of $\tilde{\delta}(\Delta\mathbf{Q})$ in Fig. 5.8.

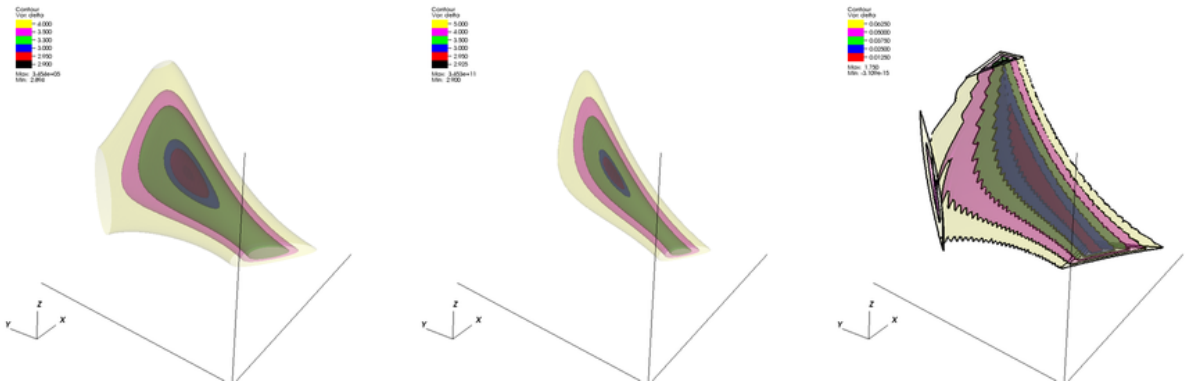

Figure 5.13: The contours (iso-surfaces) of the total interaction energy $U(\Delta\mathbf{Q})$ for the example in Fig. 5.9 when the ellipses are deformable and interact via a power-law potential. Going from left to right, we show $\nu = 12$ and 25, as well as the hard ellipsoid $\tilde{\delta}(\Delta\mathbf{Q})$, corresponding to the limit $\nu \to \infty$.

origin, i.e., the energy has a minimum. The higher the exponent p is, however, the more anharmonic the interaction potential becomes and the contours are no longer ellipsoidal near the energy minimum.

It should be emphasized that not all energy minima correspond to packings. In particular, at large pressures or applied forces the deformability of the particles becomes important and the energy minima no longer have the geometric structure of packings. However, in the limit of no externally-applied forces, i.e., $\mathbf{f} \to 0$, the only interacting particles are those that barely overlap, i.e., that are nearly touching. Therefore energy minima for purely-repulsive interaction potentials and a finite cutoff correspond to jammed packings of hard particles in the limit of zero external pressure (alternatively, one can keep the applied forces constant and make the grains infinitely stiff [161]). Therefore, the packings of soft particles studied in Ref. [164] very slightly above the "jamming treshold" ϕ_c are closely related to collectively jammed ideal packings of spheres of diameter $D = \sigma$ (polydispersity is trivial to incorporate) [176].

5.8.2 Hessian Eigenvalues and Jamming

It is well-known that for smooth interactions a given configuration is a stable energy minimum if the gradient of the energy is zero and the Hessian is positive-definite, and the converse is also true if positive-definite is replaced with positive-semidefinite. This has been used as a criterion for jamming in systems of deformable particles [164, 176].

The gradient of $U = \sum_{ij} U_{ij}$ is

$$\begin{aligned} \boldsymbol{\nabla}_{\mathbf{Q}} U &= \sum_{ij} \frac{dE}{d\zeta_{ij}} \left(\boldsymbol{\nabla}_{\mathbf{Q}} \zeta_{ij}\right) = \left(\boldsymbol{\nabla}_{\mathbf{Q}} \boldsymbol{\zeta}\right) \left(\boldsymbol{\nabla}_{\boldsymbol{\zeta}} E\right) \\ &= \mathbf{A} \left(\boldsymbol{\nabla}_{\boldsymbol{\zeta}} E\right) = -\mathbf{A}\mathbf{f}. \end{aligned}$$

The first-order *necessary* condition for a stable energy minimum is therefore exactly the force/torque balance condition

$$\mathbf{A}\mathbf{f} = 0 \text{ and } \mathbf{f} \geq 0,$$

as we derived using linear programming and duality theory for hard-particle packings. The Hessian is

$$\begin{aligned} \boldsymbol{\nabla}^2_{\mathbf{Q}} U &= \left[\boldsymbol{\nabla}_{\mathbf{Q}} \left(\boldsymbol{\nabla}_{\boldsymbol{\zeta}} E\right)\right] \mathbf{A}^T + \left(\boldsymbol{\nabla}^2_{\mathbf{Q}} \boldsymbol{\zeta}\right) \left(\boldsymbol{\nabla}_{\boldsymbol{\zeta}} E\right) \\ &= \left[\mathbf{A} \left(\boldsymbol{\nabla}^2_{\boldsymbol{\zeta}} E\right)\right] \mathbf{A}^T + \left(\boldsymbol{\nabla}_{\mathbf{Q}} \mathbf{A}\right) \left(\boldsymbol{\nabla}_{\boldsymbol{\zeta}} E\right) \\ &= \mathbf{A}\mathbf{K}\mathbf{A}^T - \mathcal{H}\mathbf{f} = \mathbf{A}\mathbf{K}\mathbf{A}^T - \mathbf{H}, \end{aligned}$$

where $\mathbf{K} = \boldsymbol{\nabla}^2_{\boldsymbol{\zeta}} \epsilon = \text{Diag}\left\{k_{ij}\right\}$ is an $[M \times M]$ diagonal matrix with the stiffness coefficients along the diagonal, and $\mathcal{H} = \boldsymbol{\nabla}_{\mathbf{Q}} \mathbf{A} = \boldsymbol{\nabla}^2_{\mathbf{Q}} \boldsymbol{\zeta}$ is the Hessian of the overlap constraints. Note that more careful notation with derivatives of vectors and matrices can be developed and should in principle be employed in calculations to avoid confusions about the order of matrix multiplications and transpositions [84].

The Hessian

$$\mathbf{H}_U = \boldsymbol{\nabla}^2_{\mathbf{Q}} U = \mathbf{A}\mathbf{K}\mathbf{A}^T - \mathbf{H}$$

consists of two terms, the *stiffness matrix* $\mathbf{H}_K = \mathbf{A}\mathbf{K}\mathbf{A}^T$, and the *stress matrix* $\mathbf{H}$ that we already encountered in the second-order expansion of the impenetrability constraints. The second-order *sufficient* condition for a strict energy minimum is

$$\mathbf{H}_U \succ \mathbf{0}.$$

Since $\mathbf{K} > \mathbf{0}$, the stiffness matrix $\mathbf{H}_C$ is positive-semidefinite: For any vector $\Delta\mathbf{Q}$ that is not a floppy mode, $\Delta\mathbf{Q}^T \mathbf{H}_K \Delta\mathbf{Q} > 0$, while $\Delta\mathbf{Q}^T \mathbf{H}_K \Delta\mathbf{Q} = 0$ if $\Delta\mathbf{Q}$ is a floppy mode (i.e., $\mathbf{A}^T \Delta\mathbf{Q} = \mathbf{0}$). Therefore, for any direction of particle motion that is not a floppy mode, one can make the stiffness coefficients large enough to make $\Delta\mathbf{Q}^T \mathbf{H}_K \Delta\mathbf{Q} > 0$, regardless of the value of $\Delta\mathbf{Q}^T \mathbf{H} \Delta\mathbf{Q}$. Floppy modes, however, correspond to negative curvature directions of the Hessian $\mathbf{H}_U$ if they are positive-curvature directions of the stress matrix, $\Delta\mathbf{Q}^T \mathbf{H} \Delta\mathbf{Q} > 0$. Therefore, *the energy minimum is strict if and only if the stress matrix is negative-definite on the space of floppy modes*. This is exactly the same result as the second-order condition for jamming we derived in Section 5.5 using duality theory.

For deformable particles, the stiffness coefficients are finite. Therefore, for sufficiently large interparticle forces, the stress matrix may affect the eigenspectrum of the Hessian $\mathbf{H}_U$ and therefore the stability of potential energy minima. For spheres, as we derived earlier, $\mathbf{H} \succ \mathbf{0}$ and therefore interparticle forces may only destabilize packings: This is the well known result that increasing the interparticle forces leads to *buckling modes* in sphere packings [160, 177]. Jamming in systems of soft spheres is therefore considered in the limit of $\mathbf{f} \to \mathbf{0}$, i.e., the point when particles first start interacting [164, 173]. For ellipsoids however, the forces can, and in practice they do, provide stability against negative or zero-frequency vibrational modes. The magnitude of the forces becomes important, and will determine the shape of the density of states (DOS) spectrum [177, 173] for small vibrational frequencies. To quote from Ref. [160], "*The basic claim...is that one cannot understand the mechanical properties of amorphous materials if one does not explicitly take into account the direct effect of stresses.*"

5.8.3 An Example of Pre-Stress Stability

Figure 5.14 shows a very simple example in which pre-stressing, i.e., pre-existing forces, stabilize a structure. Although the example is not a packing, it illustrates well some of the essential features. First, the geometry of the system is degenerate, since the two springs are exactly parallel. This degeneracy insures that a self-stress exists, since one can stretch/compress both springs by an identical amount and still maintain force balance.

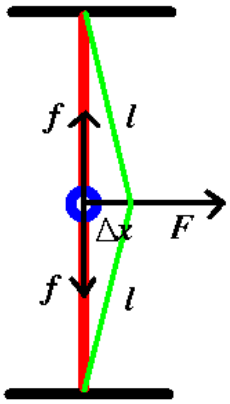

Figure 5.14: An example of a pre-stress stable system. Two elastic springs of stiffness k and length l are connected via a joint that can move in the horizontal direction under the influence of an external force F.

Observe that geometrically the change in the position of the joint Δx causes a *quadratic* change in the length of each spring $\Delta l \approx \Delta x^2$. To balance an applied force F, the force inside each spring f needs to be $f\Delta x = F$. If the system is not pre-stressed, then the potential energy is *quartic* around the origin, $\Delta U = \frac{1}{2}k\Delta l^2 \approx \frac{1}{2}k\Delta x^4$, and the applied force causes a very large deformation of the structure $\Delta x = (F/k)^{1/3}$. The structure is stable (i.e., corresponds to a jammed packing), however, its response to perturbations is not harmonic. If however there is an initial force f in the springs, then the potential energy is *quadratic* around the origin $\Delta U \approx f\Delta l = f\Delta x^2$ and the deformation is linear in the applied force $\Delta x = F/f$. If $f < 0$, then the system is unstable and will buckle, and if $f > 0$ the system is stable and its response to perturbations is harmonic. This is exactly the form of stability that hypostatic ellipsoid packings have.

5.9 Packings of Nearly Spherical Ellipsoids

In this section we will consider nearly spherical ellipsoids, that is, ellipsoids with aspect ratio α close to unity. In particular, we try to understand why these packings are hypostatic and to quantitatively explain the sharp rise in the density and contact numbers of disordered packings as asphericity is introduced. We propose that the packings of nearly spherical ellipsoids should be looked at as continuous perturbations of jammed disordered disk packings, and establish the leading order terms in the expansion around the sphere point.

5.9.1 Rotations and Translations Are Not Equal

One might at first sight expect a discontinuous change in the contact number, and therefore the structure, as asphericity is introduced. After all, the number of degrees of freedom jumps suddenly from $d_f = d$ to (for non-spheroids) $d_f = d(d+1)/2 > d$. However, such an expectation is not reasonable. Firstly, the number of degrees of freedom is $d_f = d(d+1)/2$ even for spheres, since spheres can rotate too. This rotation does not affect the non-overlap conditions and therefore is not coupled to translational degrees of freedom. If the ellipsoids are nearly spherical, particle rotation is only mildly coupled to particle translations and rotation only affects the non-overlap conditions very close to the jamming point. This is seen, for example, through

a violation of the equipartition theorem in non-equilibrium MD simulations of hard ellipsoids, depending on the moment of inertia of the particles and the time-scale of the system evolution. We therefore expect that thermodynamically and kinetically, at least at the level of translations, systems of nearly spherical ellipsoids will behave identically to systems of spheres until the interparticle gaps become comparable to the difference between the semiaxes. It is therefore not really surprising that the properties of the jammed packings such as ϕ_J or $\bar{Z}$ change continuously with α.

What is somewhat surprising however is that ϕ_J and $\bar{Z}$ are not analytic functions of particle shape. In particular, starting with a unit sphere and changing a given semiaxes by $+\epsilon \ll 1$ increases the density linearly in ϵ, and changing it by $-\epsilon$ also increases the density by the same amount, $\Delta\phi_J \sim |\epsilon|$. As we will show through our calculations, this non-analyticity is a consequence of the breaking of rotational symmetry at the sphere point. The particle orientations themselves are not analytic functions of particle shape and change discontinuously as the sphere point is crossed.

Finally, there is little reason to expect packings of nearly spherical particles to be rotationally jammed. After all, sphere packings are never rotationally jammed, since the spheres can rotate in place arbitrarily. Similarly, near the jamming point, it is expected that particles can rotate significantly even though they will be translationally trapped and rattle inside small cages, until of course the actual jamming point is reached, at which point rotational jamming will also come into play. It is therefore not surprising that near the sphere point the parameters inside the packing-generation protocol, such as the moment of inertia of the particles and the expansion rate of the particles, can significantly affect the final results. In particular, using fast particle expansion or too large of a moment of inertia leads to packings that are clearly not rotationally jammed, since the torques are not balanced, however, they are translationally jammed and have balanced centroid forces. We do not have a full understanding of the dynamics of our packing-generation algorithm, even near the jamming point.

We will focus on packings that are also rotationally jammed, however, in general one may need to distinguish between *translational* and *rotational jamming*. For example, the ellipsoid packing produced by simply stretching the crystal packing of spheres along a certain axes by a scaling factor of α is translationally but not rotationally jammed. This is because by changing the axes along which the stretch is performed one gets a whole family of ellipsoid packings with exactly the same density. This is illustrated in Fig. 5.15 in two dimensions.

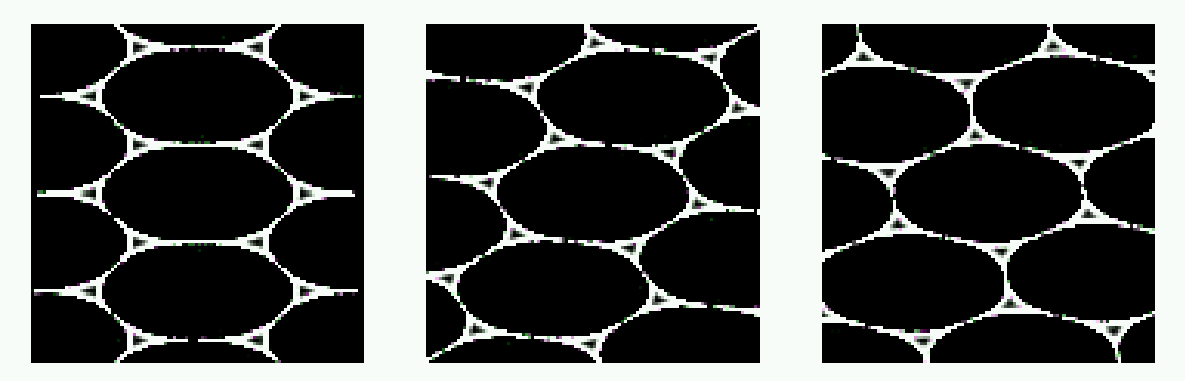

Figure 5.15: The triangular packing of ellipses is not rotationally jammed since one can shear the packing continuously, without introducing overlap or changing the density. The figure shows a sequence of snapshots as this shearing motion proceeds. The packing is however translationally jammed.

5.9.1.1 Isostatic Packings are Translationally Ordered

As we already demonstrated, in order for a hypostatic packing of ellipsoids to be jammed, the packing geometry must be degenerate. The existence of a self-stress $\mathbf{f}$ requires that the orientations of particles be chosen so that the torques are balanced in addition to the forces on the centroids. This leads to a loss of

"randomness" in a certain sense, since the number of jammed configurations is reduced greatly by the fact that geometrically peculiar configurations are needed to balance the torques.

However, it is also important to point out that disordered isostatic packings of nearly spherical ellipsoids are hard to construct. In particular, achieving isostaticity near the sphere point requires translational ordering. In two dimensions, the average number of contacts per particle needed is $\bar{Z} = 6$, however, the maximal kissing number near the sphere point is also $Z_{\max} = 6$. Therefore the only possibility is that every particle have exactly $Z = 6$ contacts. This inevitably leads to translational ordering on a triangular lattice. In other words, the only isostatic packing of ellipses in the limit $\alpha \to 1$ is the hard disk triangular crystal. Similarly, in three dimensions, $\bar{Z} = Z_{\max} = 12$ for non spheroids, and therefore every particle must have exactly $Z = 12$ neighbors. While it not rigorously known what are the sphere packings with all particles having twelve neighbors, it is likely that only stacking variants of the FCC/HCP lattice achieve that property. For spheroids, the isostatic number of contacts is $\bar{Z} = 10$ and the results in Fig. 5.2 indicate that this value is nearly reached for sufficiently large aspect ratios. For nonspheroids, however, we only observe a maximum of 11.4 contacts per particle, consistent with the fact that achieving the isostatic value requires more translational ordering. It is interesting to note that even with exactly $\bar{Z} = 12$ contacts per particle, an ellipsoid packings does not have enough contacts to be first-order jammed, since the total number of degrees of freedom exceeds the number of constraints, $N_f = 6N + 3 > M = 6N$.

In conclusion, near the sphere point there is a competition between translational and rotational jamming and also between translational and translational disorder. At the sphere point $\alpha = 1$, and it is neighborhood, translations win. As one moves away from the sphere point, however, translations and rotations start to play an equal role. For very large aspect ratios, $\alpha \gg 1$, it is expected that rotations will dominate although we do not investigate that region here.

5.9.2 Maintaining Jamming Near the Sphere Point

In this section we will make extensive use of the expansions of the overlap potentials for two nearly spherical ellipsoids given in Section 2.2.6.

Assume now that we have a collectively jammed *isostatic* sphere packing with density ϕ_J^S and that we want to make the disks slightly ellipsoidal by shrinking them along a given set of axes, while still *preserving* jamming. Keeping orientations fixed, one can expand each near-sphere by a scaling factor $\Delta\mu$ *and* displace each centroid by $\Delta\mathbf{r}$, so that all particles that were initially in contact are still in contact. Note that because the matrix $\mathbf{S}$ is proportional to ϵ, so will $\Delta\mu$ and $\Delta\mathbf{R}$. In other words, the change in the density will be linear in asphericity. However, the value of the *slope* depends on the choice of orientations of the ellipsoids. Referring back to Section 5.5.5.2 we see that to first order in ϵ, $\Delta\mu$ is

$$\begin{aligned}\Delta\mu &= \frac{1}{M}\mathbf{f}^T\Delta\boldsymbol{\zeta} = \frac{1}{M}\sum_{\{i,j\}} f_{ij}\mathbf{u}_{ij}^T\mathbf{S}_{ij}\mathbf{u}_{ij} \\ &= \frac{1}{2M}\sum_i \sum_{j\in\mathcal{N}(i)} f_{ij}\mathbf{u}_{ij}^T\mathbf{T}_i\mathbf{u}_{ij},\end{aligned}$$

giving a new jamming density

$$\phi_J/\phi_J^S = (1+\Delta\mu)^d \prod_{k=1}^{d}(1+\epsilon_i^O) \approx 1 + d\Delta\mu + \mathbf{e}^T\boldsymbol{\epsilon}_O.$$

Keeping all ellipsoids aligned produces an affine deformation of the sphere packing that has the same jamming density, but is not (first-order) jammed. Therefore, the true jamming density must be higher, $\phi_J \geq \phi_J^S$. This explains why the jamming density increases with aspect ratio near the sphere point. The added rotational degrees of freedom allow one to increase the density beyond that of the aligned (nematic) packing, which for ellipsoids has exactly the same density as the sphere point.

Can we find a set of orientations for the ellipsoids so that the resulting packing is jammed? The first condition for jamming is that there exist a self-stress that balances both forces and torques on each particle. Just from the force-balance condition, one can already determine the magnitudes of the interparticle

forces $\mathbf{f}$. These will change little as one makes the particles slightly aspherical, because the normal vectors barely change. Therefore, the self-stress is already known a priori, without regard to the choice of particle orientations. The orientations must be chosen so that the *torques* are also balanced. The torque exerted by the contact force $\mathbf{f} = f\mathbf{n}$ on a given particle, to first order in asphericity ϵ, comes about because the normal vector no longer passes through the centroid of the particle (as it does for spheres). One can ignore the small changes in the magnitude of the normal force or the change in the contact point $\mathbf{r}_C$, and only consider the change in the normal vector

$$\mathbf{n} \approx \mathbf{X}\mathbf{u} \approx (\mathbf{I} - 2\mathbf{T})\,\mathbf{u} = \mathbf{u} - 2\mathbf{T}\mathbf{u},$$

giving a torque

$$\boldsymbol{\tau} = \mathbf{r}_C \times \mathbf{f} \approx 2Of\,(\mathbf{T}\mathbf{u}) \times \mathbf{u}.$$

Therefore, to first order in asphericity ϵ, the torque balance condition for particle i is

$$\sum_{j\in\mathcal{N}(i)} f_{ij}\,(\mathbf{T}_i\mathbf{u}_{ij}) \times \mathbf{u}_{ij} = \sum_{j\in\mathcal{N}(i)} f_{ij}\mathbf{U}_{ij}\mathbf{T}_i\mathbf{u}_{ij} = 0. \tag{5.27}$$

This gives for each particle a set of possible orientations, given the contact network of the isostatic sphere packing. The torque balance condition (5.27) is in fact the first-order optimality condition for maximizing the jamming density, as expected. It is worth pointing out that for a random assignment of orientations to ellipses the expected change in density is identically zero; in order to get an increase in the density one must use orientations correlated with the translational degrees of freedom.

5.9.2.1 Ellipses

In two dimensions, for a particular contact with $\mathbf{u} = \langle \cos\theta, \sin\theta \rangle$ we have the simple expressions

$$\begin{aligned} \mathbf{u}\mathbf{T}_\phi\mathbf{u} &= \sin^2(\phi - \theta) \\ \mathbf{u} \times (\mathbf{T}_\phi\mathbf{u}) &= \frac{1}{2}\sin\left[2(\phi - \theta)\right]. \end{aligned}$$

Considering 2ϕ as the variable, one easily finds the solution to Eq. (5.27)

$$2\phi = \arctan(\pm\sum_i f_i \sin 2\theta_i, \pm\sum_i f_i \cos 2\theta_i). \tag{5.28}$$

If we calculate the second derivative for the density increase we find that

$$\frac{d^2}{d\phi^2}\left[\sum_i f_i \sin^2(\phi - \theta_i)\right] = \pm 1,$$

and therefore in order to maximize the jamming density we need to choose the minus signs in Eq. (5.28). Once we find the unique orientation of each ellipse that ensures torque balance, we can calculate the jamming density

$$\phi_J/\phi_J^S \approx 1 + s_\phi\epsilon, \tag{5.29}$$

where

$$s_\phi = 2\frac{\sum_i \sum_{j\in\mathcal{N}(i)} f_{ij}\left(\mathbf{u}_{ij}^T\mathbf{T}_i^\phi\mathbf{u}_{ij}\right)}{\sum_i \sum_{j\in\mathcal{N}(i)} f_{ij}} - 1.$$

We have calculated the slope s_ϕ for disordered binary disk packings (with $\phi_J^S \approx 0.84$) numerically, and find a value $s_\phi \approx 0.454$. We compare this theoretical value with numerical calculations in Fig. 5.16. The first comparison is directly to the packing fractions obtained using the Lubachevsky-Stillinger algorithm, which do not have anything to do with perturbing a sphere packing. Although the simulation jamming densities are not linear over a wide range of aspect ratios, near $\alpha = 1$ they are and the slope is close to the theoretically-predicted s_ϕ. We also compare to results obtained by perturbing a jammed disk packing using MD. Specifically, we start with a jammed disk packing at a relatively high pressure ($p = 1000$) and assign

an orientation according to Eq. (5.28) to every disk, and then we start growing the large semiaxes slowly while performing a form of constant-pressure MD. The density changes automatically to keep the pressure constant, and from the instantaneous density we estimate the jamming density using Eq. (5.24). In Fig. 5.16 we show how the (estimated) jamming density changes with aspect ratio. If we freeze the orientations (i.e., use an infinite moment of inertia), we obtain results that follow the theoretical slope prediction closely. Very good agreement with the results from the LS algorithm is obtained over a wide range of α if we start with the correct orientations and then allow the ellipse orientations to change dynamically. For comparison, in the inset we show that the packing density actually decreases if we use the LS algorithm and freeze orientations at their initial (random) values, demonstrating that balancing the torques and (maximally) increasing the density requires a particular value for the particle orientations.

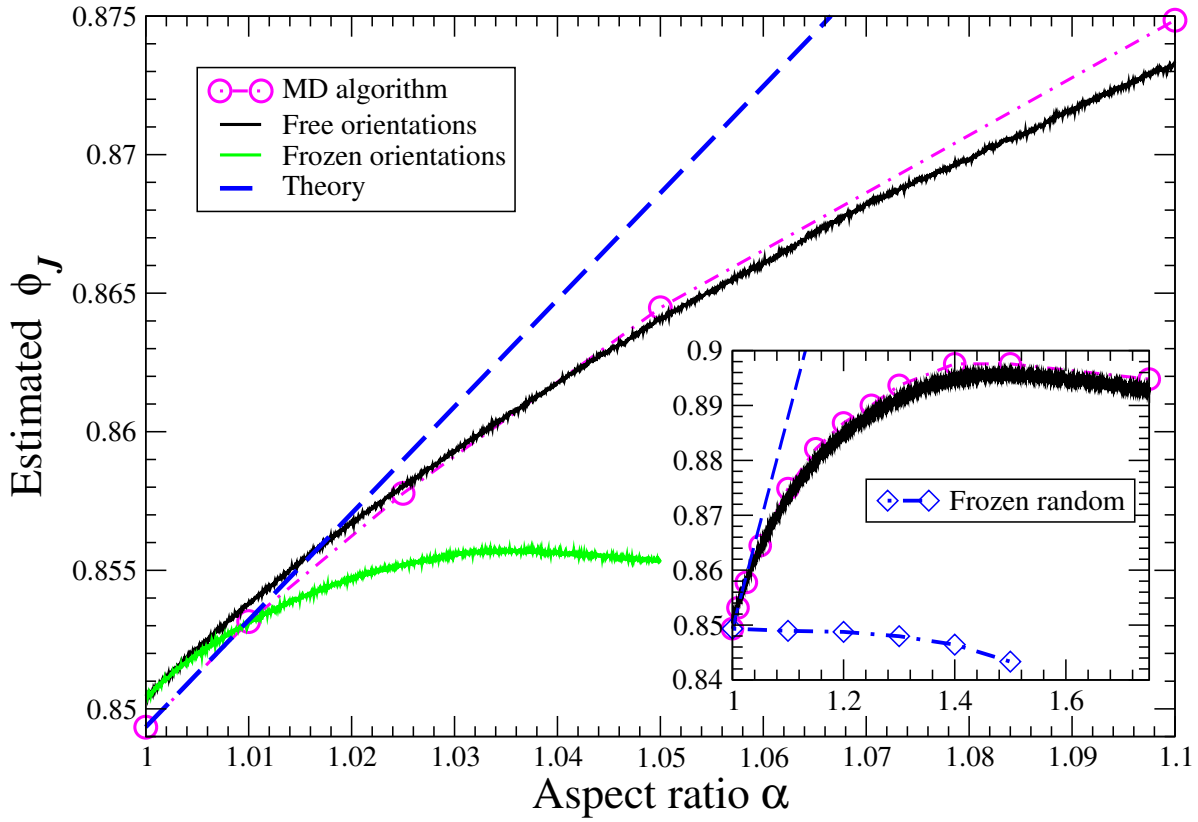

Figure 5.16: The estimated jamming density near the disk point for binary packings of hard ellipses, as obtained from the LS packing algorithm, from perturbing the disk packing using constant-pressure MD, and from the first-order perturbation theory. The inset shows some of the data over a larger range of aspect ratio and also shows the packing densities obtained when the ellipses have infinite moment of inertia in the LS algorithm.

For ellipses, there are unique orientations that guarantee the existence of self-stresses near a given isostatic jammed disk packing. Do these orientations actually lead to jammed packings, that is, are the second-order conditions for jamming also satisfied? If one starts with a jammed disk packing and transforms the disks into ellipses of aspect ratio sufficiently close to unity, the packing will remain translationally jammed [165]. Subsequent increase in the size of the particles must eventually lead to a packing of maximal density. It is not however *a priori* whether this packing is rotationally and translationally jammed or has some kind of peculiar unjamming motions that preserve the density, such as the ones shown in Fig. 5.15. For small disk packings, we have found the perturbed ellipse packings to be second-order jammed sufficiently close to the sphere point. For larger systems, even for very small asphericities, it is difficult numerically to perturb a given disk packing into an ellipse packing without leading to new contacts or breaking of old ones, as

discussed shortly. An analytical investigation may be able to prove that the perturbed packings are actually second-order jammed, and therefore prove that there exist (large) jammed ellipse packings with $\bar{Z} = 4$, the absolute minimum contact number possible for a jammed packing.

Finally, we note that in three dimensions the torque balance equations (5.27) involve quaternions and are quartic, and it does not seem an analytical solution is possible as it is in two dimensions. We however expect that the calculations performed here in $d = 2$ can be generalized to higher dimensions as well. One interesting question to answer theoretically in $d = 3$ is whether the middle axes (β) affects the slope of the density s_ϕ or whether only the ratio of the largest to the smallest semiaxes (α), matters. In Ref. [78] we proposed that the rapid increase in packing fraction could be attributed to the need to increase the contact numbers, since forming more contacts requires a denser packing of the particles. This is supported by the observation that the maximal packing density is achieved for the most aspherical shape ($\beta = 1/2$), however, the numerical results very close to the sphere point shown in Fig. 5.4 are also consistent with a slope s_ϕ independent of β. The arguments of this section indicate that the density rise is independent of the rise of the coordination number, at least near the sphere point.

5.9.3 Contact Number Near the Sphere Point

In our perturbation approach to ellipsoid packings near the sphere point, we assumed that the contact network remains that of the disk packing even as the aspect ratio moves away from unity. However, as the aspect ratio increases and the packing structure is perturbed more and more, some new contacts between nearby particles will inevitably close, and some of the old contacts may break. In Fig. 5.17 we show a system that the linear perturbation prediction produces at $\alpha = 1.025$. While the original contacts in the jammed disk packing are maintained relatively well, we see that many new overlaps form that were not contacts in the disk packing. This means that the contact number will increase from $\bar{Z} = 4$ as asphericity is introduced.

These observations suggest a way to calculate the leading order term of $\bar{Z}(\alpha) - 2d$: We simply count the overlaps introduced by orienting and displacing the centroids of the ellipsoids according to the linear perturbation theory. It is well-known that jammed disordered sphere packings have an unusual multitude of nearly-touching particles, as manifested by a nearly inverse-square-root divergence in the pair correlation function near contact [75]. These near contacts will close to form true contacts and cause the rapid increase in $\bar{Z}(\alpha)$, and we expect that the growth will be of the form

$$\bar{Z}(\alpha) - 2d \approx Z_\alpha \sqrt{\alpha - 1}. \tag{5.30}$$

A more rigorous analysis is difficult since we do not really have an understanding of the geometry of the near contacts. We have numerically estimated the coefficient Z_α and plotted the prediction of Eq. (5.30) in Fig. 5.3. It is seen that the prediction matches the actual simulation results well sufficiently close to the sphere point.

5.10 Conclusions

We presented in detail the mathematical theory of jamming for packings of ellipsoids and tried to understand the properties of jammed packings of ellipsoids of aspect ratio close to unity. In this section we summarize our findings and also point to directions for future investigation.

In Section 5.2 we applied the modified LS algorithm to generate disordered jammed packings of hard ellipsoids. The simulations showed that ellipsoids can randomly pack more densely, approaching $\varphi \approx 0.74$ for non-spheroidal ellipsoids. Both the density of the random packings ϕ_J and the average contact number $\bar{Z}$ rise sharply from their sphere values as asphericity is introduced, showing that the sphere point is a singular point and that orientational degrees of freedom dramatically impact the properties of jammed packings. The contact numbers reach values close to isostatic only for very aspherical particles, while the packings of nearly spherical ellipsoids are rather hypostatic, in sharp contrast with random sphere packings, which have been shown to be isostatic.

In Section 5.4 we discussed the conjecture that large disordered jammed packings of hard particles are isostatic, i.e., that they have an equal number of constraints and degrees of freedom, $\bar{Z} = 2d_f$. It is not possible to make this conjecture into a theorem since the term "disordered" is highly nontrivial to define

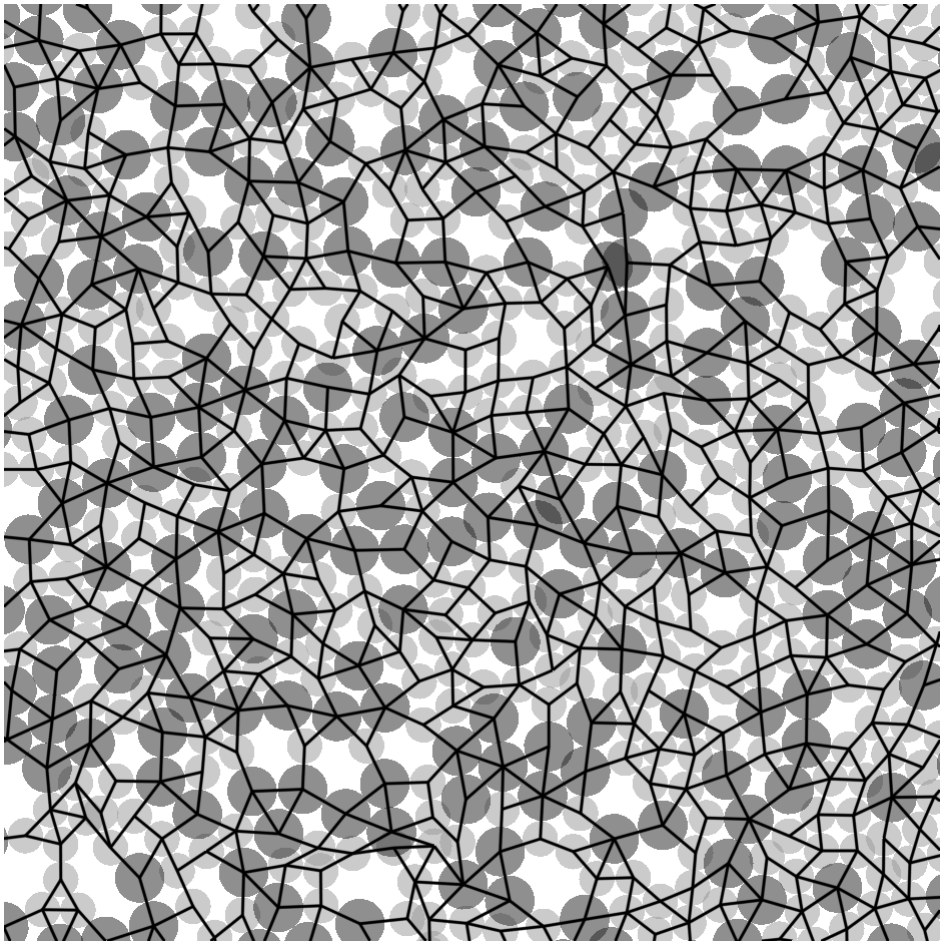

Figure 5.17: Overlaps introduced at $\alpha = 1.025$ by the naive linear perturbation theory which only takes into account the original contact network of the disk packing (black lines). We see many overlaps forming between particles that were nearly touching when $\alpha = 1$.

[58], however, arguments have been made in the literature in support of the isostatic property. We showed that this conjecture can be supported with reasonable arguments only for spheres, where particle rotations are not considered. In particular, while it is expected that $\bar{Z} \leq 2d_f$ for "random" packings, the converse inequality $\bar{Z} \geq 2d_f$ only applies to spheres. Packings of nonspherical particles can be jammed and have less than $2d_f$ contacts per particle, i.e., be hypostatic. A minimally rigid ellipsoid packing, i.e., a packing that has the minimal number of contacts needed for jamming, satisfied only the inequality $\bar{Z} \geq 2d$, since at least $2d$ contacts per particle are needed to block particle translations. Particle rotations, however, and combined rotation/translation motions, can be blocked by the curvature of the particle surfaces at the point of contact. In essence, if the radii of curvatures at the point of contact are sufficiently large, i.e., the particle contact is sufficiently "flat", rotation of the particles is blocked. This can be visualized by considering the limit of infinite radii of curvatures, when have a contact between two flat surfaces. Such contacts, in a certain sense, count as several "contact points" and block several degrees of freedom.

In Section 5.5 we generalized the mathematics of first and second-order rigidity for tensegrity frameworks developed in Ref. [73] to packings of nonspherical particles. We proved that in order for a packing to be jammed there must exist a set of (nonzero) non-negative interparticle forces that are in equilibrium, i.e., the packing must have a self-stress. Furthermore, we considered second-order terms for hypostatic packings that do have a self-stress but also have floppy modes, that is, particle motions that preserve interparticle distances to first order. The second-order analysis showed that jammed packings of strictly convex particles cannot

have less than $2d$ contacts per particle. We found that floppy modes involving particle rotations can be blocked (rigidified) by the stress-matrix, which includes second-order information about the particle surfaces at the point of contact. We proposed that this is exactly the type of jamming found in disordered ellipsoid packings near the sphere point, and in Section 5.6 we presented a numerical algorithm for testing hypostatic ellipsoid packings for jamming and applied it to some computer-generated samples. We demonstrated that the packings are indeed jammed even very close to the sphere point, where they have close to $2d$ contacts per particle.

In Section 5.7 we considered the thermodynamics of packings that are close to, but not exactly at, the jamming point, so that particles have some room to rattle (free volume). We found that for hypostatic packings the jamming basin $\mathcal{J}_{\Delta\mathbf{Q}}$, which is localized around the jamming point in configuration space, is very elongated along the space of floppy modes. For iso- or hyper-static packings, as jammed sphere packings always are, the jamming basin approaches a polytope in the jamming limit, whereas for hypostatic packings it approaches a (hyper) banana. The latter leads to very large oscillations of the instantaneous pressure near the jamming point and a violation of the asymptotic free-volume equation of state (pressure scaling).

Real packings are always made from deformable (albeit very stiff) particles, i.e., particles that interact via some elastic interaction potential. The analog of a jammed hard-particle packing for deformable particles are strict energy minima (inherent structures), i.e., structures where any motion of the particles costs energy (quadratic in the displacements). In Section 5.8 we analyzed the first- and second-order conditions for a strict energy minimum for twice-differentiable interaction potentials. We found that the first-order condition is exactly the requirement for the existence of a self-stress, and that the second-order condition is exactly the condition that the stress-matrix blocks the floppy modes. This deep analogy between jamming in hard-particle packings and energy minima in soft-particle packings is not unexpected since a "soft" potential can approximate the singular hard-particle potential arbitrarily closely. As the potential becomes stiffer, the energy minimum will become highly anharmonic and its shape will closely resemble that of the jamming basin $\mathcal{J}_{\Delta\mathbf{Q}}$ (even at very small temperatures).

Finally, in Section 5.9 we developed a first-order perturbation theory for packings of nearly spherical ellipsoids, expanding around the sphere point. The theory is based on the idea that packings of ellipsoids with aspect ratio $\alpha = 1 + \epsilon$ near unity have the same contact network as a nearby jammed isostatic packing of hard spheres. In order for the ellipsoid packing to also be jammed, the orientations of the ellipsoids must be chosen so as to balance the torques on each particle. These orientations also maximize the jamming density, increasing it beyond that of the disk packing, and we analytically calculated the linear slope of the density increase with ϵ for binary ellipse packings. The calculated coefficient is in good agreement with numerical results. The perturbation of the sphere packing also leads to a rapid increase in the average particle coordination $\bar{Z}$, which we attributed to the closing of the multitude of near contacts present in disordered disk packings. The predicted $\bar{Z} \sim \sqrt{\epsilon}$ is also in good agreement with numerical observations.

The observed peculiar behavior of ellipsoid packings near the sphere point is a consequence of the breaking of rotational symmetry. Near the sphere point the coupling between particle positions and orientations is weak and translations dominate the behavior of the system. In this sense sphere packings are a good model system, and particle shapes close to spherical can be treated as a continuous perturbation of sphere packings. However, even for aspect ratios relatively close to unity, the perturbation changes the properties of the system such as density and contact number in a sharp fashion, making sphere packings a quantitatively unreliable reference point for packings of more realistic particle shapes. Furthermore, even qualitative understanding of jamming and mechanical rigidity for packings of nonspherical particles requires consideration of phenomena that simply do not have a sphere equivalent.

We hope that through our work it has become clear that the study of nearly jammed packings requires the use of mathematical programming techniques. Unjamming motions at very high densities are collective and involve careful correlated motion of most if not all of the particles in the packing, and such collective motions cannot be efficiently found using Monte Carlo or Molecular Dynamics methods because of the long separation of time scales between local rattling motion and collective rearrangements. It remains a challenge for the future to develop quadratic programming techniques for studying collective rearrangements in packings of hard ellipsoids. Future work should consider the mathematics of jamming for packings of hard particles that are convex, but not necessarily smooth or strictly convex. In particular, particles with sharp corners and/or flat edges are of interest. We also believe that jamming in frictional packings, even for the case of spheres, is not understood well-enough. It is also important to consider packings of soft ellipsoids and in

particular develop algorithms to generate them computationally and to study their mechanical properties. Investigations of the thermodynamics of nearly jammed ellipsoid packings also demand further attention.

Chapter 6

Free Energy of Nearly Jammed Packings

In this chapter we present in considerable detail an event-driven molecular dynamics algorithm for measuring the free energy of nearly jammed hard-particle packings [82]. This Bounding Cell Molecular Dynamics (BCMD) algorithm calculates exactly the free-energy of a single-occupancy cell (SOC) model in which each particle is restricted to the neighborhood of its initial position using a hard-wall bounding cell. It is based on previous MD algorithms appearing in the literature, however, several small but important modifications enable us to apply it to nonspherical particles as well as to measure the free-energy change during continuous irreversible transformations. Additionally, we point connections to the well-studied problem of computing the volume of convex bodies in high dimensions using random walks. We test and verify the numerical accuracy of the method by comparing against rigorous (asymptotic) results for the free energy of isostatic disordered packings of both hard spheres and ellipsoids, for which the free energy can be calculated directly as the volume of a high-dimensional simplex. We also compare our results to previously published Monte Carlo results for hard-sphere crystals and find excellent agreement.

6.1 Introduction

Calculating the absolute free energy of atomic and molecular systems is a fundamental problem in computational physics [178]. Knowledge of the absolute free energy of various phases gives the phase diagram and identifies the phase transitions in the system. In recent years several problems for calculating free energies have received great attention, and conflicting claims have emerged. For example, the determination of the thermodynamically stable solid phase of the hard-sphere system [179, 45, 46] and the determination of the configurational entropy of supercooled liquids [180, 181, 182, 183] are still not settled. The latter problem in particular is of significant importance since it provides insight into the validity of ideal glass transition theories, thus addressing one of the most fundamental open problems in statistical mechanics.

Hard-particle systems are an excellent substrate for such studies because of the balance they offer between a simple model and resulting complex behavior. They have been shown to exhibit a variety of phases including liquid, solid, liquid crystal and glassy. Furthermore, they can be simulated with very fast specialized simulation techniques. Calculating the free energy of structured phases such as solid, glassy and liquid crystal phases is a particular challenge, and various approaches have been developed, mostly using Monte Carlo (MC) algorithms [178]. Hard-particle systems are however best simulated using event-driven Molecular Dynamics (MD) algorithms [11]. We present in detail an algorithm for computing the free energy (equivalently, entropy) of nearly jammed hard particle packings, i.e., hard-particle systems where diffusion can be ignored and particles remain localized in the vicinity of their initial configuration for long times. Note that (nearly) jammed packings are not in thermodynamic equilibrium and therefore the free energy we calculate is not the equilibrium free energy at the given packing fraction (density), but rather, it is the *free-volume contribution* to the thermodynamic free energy. We have successfully applied the algorithm to disordered (glassy) jammed hard-sphere packings and demonstrated that previous claims of an ideal glass transition in binary hard-sphere

systems need to reevaluated [57, 83].

The algorithm we develop fits directly into the collision-driven MD algorithm which we developed in detail in Ref. [11] and reviewed in Chapter 3. Our Bounding Cell MD (BCMD) algorithm is based on the tether method of Speedy [184] and calculates the free-energy of a single-occupancy cell (SOC) model [179] in which each particle is restricted to the neighborhood of its initial position using a hard-wall bounding cell. The BCMD algorithm can be applied to nonspherical particles, and it can measure the free-energy change during continuous irreversible transformations, unlike previous MD algorithms. We test and verify the numerical accuracy of the method by comparing against rigorous (asymptotic) results for the free energy of isostatic disordered packings of both hard spheres and ellipsoids, for which the free energy can be calculated directly as the volume of a high-dimensional simplex. Previous algorithms have only been verified by comparing to earlier results obtained with similar methods, or other indirect methods, and this is the first time that the free-energy of a nontrivial hard-particle system has been calculated with a deterministic method. We also compare our results to previously published Monte Carlo results for hard-sphere crystals and find excellent agreement. We also discuss the connections of our algorithm to the best-known algorithms for computing the volume of convex bodies in high dimensions, as well as the effects of boundary conditions.

In Section 6.2 we provide a brief mathematical introduction to the problem, focusing on the jamming limit for hard particle packings. We then describe the BCMD algorithm in Section 6.3 in considerable detail, and give an illustrative example in Section 6.3.5. We apply the algorithm to hard-sphere crystals as well as isostatic jammed packings of spheres and ellipses in Section 6.4, and verify its high numerical accuracy.

6.2 Background

We consider a (thermal) system of hard particles with covering fraction (or density) ϕ, characterized by the particle displacements $\Delta\mathbf{Q} = (\Delta\mathbf{q}_1, \ldots, \Delta\mathbf{q}_N)$ from an ideal *collectively jammed* [71] configuration $\mathbf{Q}_J$ with jamming density ϕ_J. For simplicity we will focus on sphere packings and denote the configuration with $\mathbf{R}$ unless specifically focusing on nonspherical particles, in which case we will revert to $\mathbf{Q}$. Depending on the boundary conditions, $\mathbf{Q}$ may include additional degrees of freedom, for example, the lattice vectors defining the periodic unit cell, and also certain additional constraints, such as freezing the center of mass, may be imposed, leading to a total of $N_f \sim Nd_f$ degrees of freedom.

We are concerned here with *nearly jammed* packings, where the particles are shrunk from the terminal jamming point $(\mathbf{R}_J, \phi_J)$ to a reduced density $\phi = \phi_J (1-\delta)^d \approx \phi_J (1 - d\delta)$, where $\delta \approx \Delta\mu$ is a small *jamming gap*. The configuration point $\mathbf{R} = \mathbf{R}_J + \Delta\mathbf{R}$ remains trapped in a small neighborhood $\mathcal{J}_{\Delta\mathbf{R}}$ around $\mathbf{R}_J$ [165], called a *jamming basin*. For the purposes of this work, it is not necessary that the packing be rigorously trapped inside $\mathcal{J}_{\Delta\mathbf{R}}$. Instead, we can consider $\mathcal{J}_{\Delta\mathbf{R}}$ to be the region of configuration space explored by the packing on the time scale of the observation, that is, the region of configuration space which makes an appreciable impact on the *measured* thermodynamic properties of the system. Let the number of particle pairs which participate in trapping the configuration inside $\mathcal{J}_{\Delta\mathbf{R}}$ be M. That is, $\mathcal{J}_{\Delta\mathbf{R}}$ is bounded by M surfaces determining the impenetrability condition between nearby particles. In the jamming limit M becomes the number of true particle contacts.

Two fundamental assumptions are that $\mathcal{J}_{\Delta\mathbf{R}}$ is bounded (i.e., the displacements of the particles from the jammed point are finite), and that $\mathcal{J}_{\Delta\mathbf{R}}$ is explored ergodically on the time scale of the measurement of thermodynamic properties such as pressure. In particular the last part enables us to separate the configurational from kinetic portions of phase space. These assumptions allow us to consider both nonequilibrium metastable systems such as disordered packings and equilibrium ones such as perfect crystals. Even in the equilibrium crystal there are occasional very large density fluctuations and therefore particle displacements, and additionally, there are diffusing defects such as vacancies and dislocations. However, at sufficiently high densities, the thermodynamic properties are primarily determined by the configurations close to the perfect crystal. For metastable systems sufficiently long timescales will eventually lead to large particle rearrangements into the equilibrium configuration, however, such systems can be observed in metastable disordered packing configurations for sufficiently long times as to make meaningful measurements. In a sense, the techniques presented in this work can be applied to any system where diffusion can be neglected. Free energy is a concept that strictly speaking applies only to systems in thermodynamic equilibrium, and the free energy we compute is in fact only the free-volume contribution to the thermodynamic free energy of the system.

However, phrases such as "free energy of the FCC/HCP lattice" or "free energy of a glass" are used commonly. One can formally make the concept of free energy well-defined for non-equilibrium states by restricting the parition function to configurations in the vicinity of a reference jammed packing; we will do this by use of a cell model in Section. 6.3.

The fundamental problem considered here is computing the logarithm of the volume of the body $V = |\mathcal{J}_{\Delta\mathbf{R}}|$ in configuration space. In a slight abuse of (equilibrium) thermodynamic nomenclature, we will refer to this as the (non-normalized) free-energy

$$F = -\ln |\mathcal{J}_{\Delta\mathbf{R}}| = Nf,$$

where f is the sought-after free-energy per particle, presumed to be independent of N for sufficiently large systems. Note that since there is no internal energy for hard-particle systems we can fix $kT = 1$ and also the free energy is simply the negative of the entropy. We want to alert the reader that the choice of f to denote (intensive) free energy conflicts with our notation for forces. In this Chapter we denote interparticle forces via $\vec{f}$, in the spirit of forces being vectors.

6.2.1 Jamming Polytope

As discussed in Sections 4.3.1.1 and 5.7.1, for first-order jammed packings, as sphere packings always are, in the jamming limit $\delta \to 0$ the set of displacements that are accessible to the packing $\mathcal{J}_{\Delta\mathbf{Q}}$ approaches a convex *polytope* $\mathcal{P}_{\Delta\mathbf{Q}}$ [52, 145], given by Eq. (5.23). Complications to this simple and intuitive jamming polytope picture arise for hypostatic packings of nonspherical particles, as discussed in Section 5.7.2. In particular, $\mathcal{J}_{\Delta\mathbf{Q}}$ does not approach a polytope in the jamming limit, but rather, it can have some curved (quadratic) surfaces which cannot be linearized no matter how close to the jamming point. Due to this additional layer of complexity, we will not try to study such jammed hypostatic ellipsoid packings analytically in this work. Instead, we will produce first-order jammed packings of ellipses for which the polytope picture applies. Since the translational degrees of freedom have a length scale while the rotational do not, we renormalize all distances so that the linear extents of each particle are dimensionless, for example, the smallest semiaxes of an ellipsoid is unity[1]. Note that the molecular dynamics numerical algorithm works even for second-order jammed packings.

The polytope $\mathcal{P}_{\Delta\mathbf{Q}}$ is determined by the system of linear inequalities $\mathbf{A}^T\Delta\mathbf{Q} \geq -\delta\mathbf{e}$, where we have assumed that all spheres have diameter $D = 1$ (using the notation of Chapter 4) or that the smallest semiaxes of all ellipsoids is unity[2] (using the notation of Chapter 5). Therefore, when studying the jamming limit we can focus on the normalized *jamming polytope*

$$\mathcal{P}_{\mathbf{x}} : \ \mathbf{A}^T\mathbf{x} \geq -\mathbf{e}, \tag{6.1}$$

which can be scaled by a factor of δ along all coordinate directions in order to obtain $\mathcal{P}_{\Delta\mathbf{R}}$. In particular, the volume scales in a predictable manner, $|\mathcal{P}_{\Delta\mathbf{Q}}| = \delta^{N_f} |\mathcal{P}_{\mathbf{x}}|$. The free energy thus diverges in the well-known free-volume [52] logarithmic fashion,

$$f = -d_f \ln\delta - \frac{\ln |\mathcal{P}_{\mathbf{x}}|}{N} = -d_f \ln\delta - f_J, \tag{6.2}$$

where the fundamental constant f_J is determined by the geometry of the jammed state. Computing f_J reduces the well-known #P-hard problem of computing the volume of a polyhedron in high-dimensional spaces. In general this problem is very difficult and takes exponentially long in the number of dimensions [185], and our attempts to use state-of-the-art software for computing the volume of polyhedra in high dimensions [186] have failed even for relatively small system sizes due to computational limitations.

[1] A more consistent choice would have been to follow the sphere example and make the smallest axes (diameter) unity, however, the usual overlap function used for spheres (the interparticle gap) differs by a factor of 2 (near jamming) from the PW overlap potential.

[2] A more consistent choice would have been to follow the sphere example and make the smallest axes (diameter) unity, however, the usual overlap function used for spheres (the interparticle gap) differs by a factor of 2 (near jamming) from the PW overlap potential [c.f. Eq. (5.23)], and this normalization of lengths makes the jamming polytopes have right hand sides of $-\delta\mathbf{e}$ for both spheres and ellipsoids.

6.2.2 Isostatic Packings

An especially tractable case arises when the jammed configuration is *isostatic*, in the sense that $M = N_f + 1$. In this particular case the jamming polytope is a *simplex*, whose volume can be computed exactly easily (see Fig. 4.1 for an illustration). We have derived a formula for computing the volume directly from the rigidity matrix [187], specifically,

$$|\mathcal{P}_{\mathbf{x}}|^{-1} = N_f! \left|\tilde{\mathbf{A}}\right| \prod_{ij} \vec{f}_{ij}, \tag{6.3}$$

can be calculated using common (sparse) linear-algebra operations. Here the augmented rigidity matrix

$$\tilde{\mathbf{A}} = \begin{bmatrix} \mathbf{A} \\ \mathbf{e} \end{bmatrix} \tag{6.4}$$

is invertible, and the vector

$$\vec{\mathbf{f}} = \tilde{\mathbf{A}}^{-1} \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix} \geq \mathbf{0},$$

has the physical interpretation of normalized (to a unit sum, rather than unit mean) interparticle forces appearing in the force-chains near the jamming point [75]. The derivation of this fact for a general simplex is given in Section 6.2.2.2. In Section 6.2.2.3 we discuss the behavior in the thermodynamic limit $N \to \infty$ and find that the (intensive) free energy per particle is well defined even for infinite packings, and is related to the density of vibrational states of the isostatic solid.

6.2.2.1 The Augmented Rigidity Matrix

In this section we give some useful relations concerning the augmented rigidity matrix (6.4) and its relation to the rigidity matrix $\mathbf{A}$. One can view $\tilde{\mathbf{A}}$ as a generalized rigidity matrix that includes the density (particle growth) as an extra degree of freedom, as discussed in Chapter 5. The rigidity matrix $\mathbf{A}$ for an isostatic packing is not invertible, however, if we remove a contact $\{i, j\}$, i.e., we remove the (last) column $\mathbf{a} = \mathbf{A}_{ij}$ from the rigidity matrix, the remaining square submatrix $\mathbf{B}$ is invertible. The augmented rigidity matrix becomes a $[2 \times 2]$ partitioned matrix

$$\tilde{\mathbf{A}} = \begin{bmatrix} \mathbf{B} & \mathbf{a} \\ \mathbf{e}^T & 1 \end{bmatrix},$$

and its inverse is

$$\tilde{\mathbf{A}}^{-1} = \alpha \begin{bmatrix} \mathbf{B}^{-1}\left[\alpha^{-1}\mathbf{I} + \left(\mathbf{a}\mathbf{e}^T\right)\mathbf{B}^{-1}\right] & -\mathbf{B}^{-1}\mathbf{a} \\ -\mathbf{e}^T\mathbf{B}^{-1} & 1 \end{bmatrix},$$

where

$$\alpha = \left(1 - \mathbf{e}^T\mathbf{B}^{-1}\mathbf{a}\right)^{-1}.$$

Note that $\left|\tilde{\mathbf{A}}\right| = \alpha^{-1}\left|\mathbf{B}\right|$. The interparticle forces are

$$\vec{\mathbf{f}} = \alpha \begin{bmatrix} -\mathbf{B}^{-1}\mathbf{a} \\ 1 \end{bmatrix}.$$

It is useful to explicitly calculate the vertex $\mathbf{v} = \mathbf{x}_{ij}$ of the simplex $\mathbf{A}^T\mathbf{x} \geq -\mathbf{e}$, corresponding to the contact ij. This vertex is the solution to the linear system of equations $\mathbf{B}^T\mathbf{v} = -\mathbf{e}$, $\mathbf{v} = -\mathbf{B}^{-T}\mathbf{e}$, which is exactly the last row in $\tilde{\mathbf{A}}^{-1}$, scaled by α. Therefore, an efficient way to calculate all of the vertices of a simplex is to calculate $\tilde{\mathbf{A}}^{-1}$ and then normalize each row so that the last column has all unit entries. The rows of the renormalized inverse of the augmented rigidity matrix are the vertices of the simplex. Note that the distance from the vertex $\mathbf{v}$ to the ij face of the polytope is $h_{ij} = 1 + \mathbf{v}^T\mathbf{a} = \alpha^{-1} = 1/\vec{f}_{ij}$.

6.2.2.2 The Volume of a Simplex in the $\mathcal{H}$-representation

Consider the simplex $\mathcal{S}$ given in the (hyperplane) $\mathcal{H}$-representation,

$$\mathbf{A}\mathbf{x} \leq \mathbf{b},$$

where $\mathbf{b} > \mathbf{0}$ (this simply guarantees that the simplex has a positive volume and $\mathbf{0}$ is an interior point), and $\mathbf{x} \in \mathcal{R}^n$. Here $\mathbf{A}$ itself is not invertible (it is not square), however,

$$\tilde{\mathbf{A}} = \left[\begin{array}{cc} \mathbf{A} & \mathbf{b} \end{array} \right]$$

is invertible. Therefore, let's append an artificial variable x_0,

$$\tilde{\mathbf{x}} = \left[\begin{array}{c} \mathbf{x} \\ -x_0 \end{array} \right],$$

and go to $(n+1)$-dimensional space, where $\mathcal{S}$ is given by

$$\tilde{\mathbf{A}}\tilde{\mathbf{x}} \leq \mathbf{0} \text{ and } x_0 = 1.$$

Consider now the transformation of variables

$$\tilde{\mathbf{x}} = -\tilde{\mathbf{A}}^{-1}\mathbf{y}$$

with Jacobian $J = \left|\tilde{\mathbf{A}}\right|^{-1}$, which transforms $\mathcal{S}$ into a slightly modified version of the *standard simplex* $\tilde{\mathcal{S}}$,

$$\mathbf{y} \geq 0 \text{ and } \mathbf{a}^T\mathbf{y} = 1.$$

Here $\mathbf{a}$ is the last row of $\tilde{\mathbf{A}}^{-1}$, that is

$$\mathbf{a} = \tilde{\mathbf{A}}^{-T} \left[\begin{array}{c} \mathbf{0} \\ 1 \end{array} \right].$$

It is easy to see that the volume of $\tilde{\mathcal{S}}$ is

$$\tilde{V}^{-1} = n! \prod_{i=1}^{n+1} a_i,$$

which directly gives the volume of $\mathcal{S}$ by virtue of the constant Jacobian

$$V^{-1} = n! \left|\tilde{\mathbf{A}}\right| \prod_{i=1}^{n+1} a_i. \tag{6.5}$$

All of the quantities involved here can be calculated very quickly from an LU decomposition of $\tilde{\mathbf{A}}$, the determinant from the diagonal of L and U, and $\mathbf{a}$ from the solution to $\tilde{\mathbf{A}}^T\mathbf{a} = \left[\begin{array}{c} \mathbf{0} \\ 1 \end{array} \right]$. The calculation can be made performed in sparse mode if $\mathbf{A}$ is sparse, by using common tricks to handle the dense column $\mathbf{b}$.

It is unusual that we have not found in the literature any direct prescription for calculating the volume of a simplex in the $\mathcal{H}$-representation. The standard formula quoted in most elementary textbooks is expressed in very simple form in the (vertex-based) $\mathcal{V}$-representation. It is however not computationally efficient to transform to the $\mathcal{V}$-representation just for the purposes of computing the volume, instead, Eq. (6.5) should be used.

6.2.2.3 Simplex Volume in the Limit $N \to \infty$

It is worth considering the limit $N \to \infty$ for an isostatic jammed packing more closely. We can rewrite the an explicit formula for the volume of the simplex, Eq. (6.5), as

$$f_J = \frac{\ln |\mathcal{P}_{\mathbf{x}}|}{N} = d_f - d_f \frac{\ln \left|\tilde{\mathbf{A}}\right|}{M} - d_f \frac{\sum_{i=1}^{M} \ln \vec{f}_i}{M}, \tag{6.6}$$

where $\tilde{\mathbf{A}}$ is the (augmented) rigidity matrix, and $\vec{\mathbf{f}}$ are the compressive contact forces between the particles at jamming [75]. These forces are in static equilibrium, $\mathbf{A}\vec{\mathbf{f}} = \mathbf{0}$, normalized to unit mean, $\sum_{i=1}^{M} \vec{f}_i = M$, and are positive, $\vec{\mathbf{f}} > \mathbf{0}$. As $N \to \infty$ the magnitudes of these forces follow a probability distribution $P_f(\vec{f})$, so that the last term in Eq. (6.6) above becomes an integral $d_f \left\langle \ln \vec{f} \right\rangle = d_f \int P_f(\vec{f}) \ln \vec{f} d\vec{f}$. We have studied this probability distribution in Ref. [75] and found that $P_f(0) > 0$, indicating that the polytope picture does not apply strictly since the simplex is not closed. This does not pose problems when calculating f_J since the logarithmic function is integrable, giving a finite contribution to f_J independent of N.

It would also be useful to give a physical interpretation to the second term in Eq. (6.6). We can do this by noting that

$$\frac{\ln|\mathbf{A}|}{M} = \frac{1}{2M} \ln\left|\mathbf{A}\mathbf{A}^T\right| = \frac{1}{M} \sum_{i=1}^{N_f} \ln \sqrt{\lambda_i(\mathbf{H})} = \frac{1}{M} \sum_{i=1}^{N_f} \ln \omega_i = \int P_\omega(\omega) \ln \omega d\omega, \tag{6.7}$$

where $\mathbf{H} = \mathbf{A}\mathbf{A}^T$. As we derived in Section 5.8, the dynamical matrix (Hessian) for energy minima in soft-sphere systems in the high stiffness limit, or equivalently, in the limit of no internal stresses $\vec{\mathbf{f}} \to \mathbf{0}$, is $\mathbf{A}\mathbf{K}\mathbf{A}^T$, where the diagonal matrix $\mathbf{K}$ contains the stiffness coefficients along its diagonal. Therefore, $\mathbf{H}_1$ is the dynamical matrix in the case of all stiffnesses being unity, and its eigenvalues are the squares of the (harmonic) vibrational eigenfrequencies ω_i . Essentially, the problem reduces to calculating the density of states $P_\omega(\omega)$ for a system of oscillators of unit mass coupled with springs of unit stiffness, where springs are placed between all pairs of touching spheres. It is exactly this density of states which has been calculated in Ref. [188], where it was found that $P_\omega(0) > 0$, indicating an excess of low-frequency vibrational modes (so-called Boson peak).

We can even combine the last two terms in Eq. (6.6) if we consider the density of states $P_\omega(\omega)$ when the stiffness constants of the harmonic springs are made proportional to the square of the compressive force, $\mathbf{K} = \vec{\mathbf{F}}^2$, to obtain

$$f_J = \frac{\ln|\mathcal{P}_\mathbf{x}|}{N} = d_f - d_f \int P_\omega(\omega) \ln \omega d\omega. \tag{6.8}$$

This equation is surprisingly similar to that used for the vibrational entropy of glasses around their inherent structures in the so-called harmonic approximation [178]. This shows an interesting and potentially deep connection between packings of hard spheres and inherent structures (glasses) of soft spheres, at least under the identification of spring constants proportional to the square of the contact force, as is the case, for example, for logarithmic springs. Such highly nonuniform interactions are not usually used for soft sphere glasses, for example, Ref. [188] uses springs of uniform stiffness, and it may be the very high anharmonicity of the hard-sphere potential makes such stiff interactions more appropriate for hard sphere systems. In particular, preliminary calculations show that the resulting density of states $P_\omega(\omega)$ with nonuniform stiffnesses diverges at $\omega = 0$ as $\omega^{-1/2}$. Alternatively, it may be that the limit of zero compressive force is not appropriate and one must take into consideration the internal stresses in the packing [177]. It has been proposed in Ref. [173] that the "effective" potential between hard spheres in isostatic packings is logarithmic, with stiffness constant proportional to the square of the interparticle force, although the argument used there does not actually prove that the free energy is actually connected to the vibrational free energy of the corresponding logarithmic spring soft-particle system. The $P_\omega(\omega)$ observed in Ref. [173] using such logarithmic springs does appear to diverge at $\omega = 0$ as $\omega^{-1/2}$, although the exact exponent is not reported.

The important point here is that the simplex picture enables us to derive Eq. (6.8), which can be easily applied to infinite systems even though the simplex picture itself breaks down for very large disordered packings [75]. Additionally, it reveals surprising analogies with the vibrational entropy of inherent structures of soft-sphere glasses. However, Eq. (6.8) does not directly apply to hyperstatic systems such as the triangular or the FCC crystals, and it is an interesting open question to generalize it to such systems.

6.3 Molecular Dynamic Method for Measuring f

Traditional molecular simulation methods can only measure the difference in free energy between two states, typically by constructing a reversible path between the two states, one of which has known free energy [178].

The problem is usually in finding a state for which the free-energy is known analytically, and which can be connected to the state of interest by a reversible path free of (first-order) phase transitions. The classical molecular dynamics (MD) method for hard-sphere crystals is to use the so-called *single-occupancy cell* (SOC) model, in which space is partitioned into N cells, typically the Voronoi polyhedra of the unit cell, and the centroid of each particle is constrained to remain fully within its (polyhedral) cell [179]. In the high density limit the system is (virtually) indistinguishable from the unconstrained crystal, and in the low-density limit the system is an SOC ideal gas, the free-energy of which can be calculated easily. The classical Monte Carlo (MC) method on the other hand uses a crystal state in which the particles are constrained to remain within a small neighbourhood of the equilibrium crystal sites by virtue of strong harmonic springs (the so-called Einstein crystal) [44]. In the limit of very stiff springs the hard-particle interactions become negligible and the free-energy can be calculated exactly. Both of these methods have major disadvantages. The MD method requires a complicated polyhedral cell complex, which would be rather non-trivial to construct for non-crystal packings or for systems of nonspherical particles. The MC method on the other hand requires introducing "soft" (spring) interactions into an otherwise hard-particle problem, so that special efficient hard-particle techniques such as event-driven MD cannot be substituted for MC. Additionally, both methods have an approximate (extrapolated) termination criterion, one in a limit of zero density, the other in the limit of sufficiently strong interactions.

The MD method developed here essentially combines the two methods to obtain a method which uses only hard-particle interactions and is also readily applicable to disordered configurations and to nonspherical particles. Additionally, it has a well-defined termination criterion. It is based on the *tether-method* that Speedy introduced for spheres [184], however, with some important improvements and generalizations, as we explain shortly. In particular, the term tether-method is not appropriate for nonspherical particles, instead, we prefer to continue to think of this as a cell-method, but with a wiser choice of cell than Voronoi polyhedra. We will call the method the Bounding Cell MD (BCMD) algorithm. It fits directly into the collision-driven MD algorithm which we developed in detail in Ref. [11].

6.3.1 Basic Algorithm

The algorithm for performing hard-particle MD described in Ref. [11] is collision-driven (event-driven), meaning that the algorithm makes predictions on when two particles collide and then jumps from collision to collision asynchronously. For the purposes of improving the computational efficiency of the algorithm, near-neighbor lists (NNLs) were introduced in Ref. [11] through the concept of a *bounding neighborhood*, which we will call *bounding cells* in this context. Namely, for a given snapshot configuration of the packing, each particle is surrounded by a cell which has exactly the same shape as the particle itself, but is scaled uniformly by a scaling factor $\mu = 1 + \Delta\mu$. The volume of the cell is thus $V_c = \mu^d V_p$, where V_p is the volume of the particle. Each particle only predicts collisions with the particles whose bounding neighborhoods/cells overlap with its own, and also predicts collisions with its own bounding cell. The BCMD algorithm uses exactly this same machinery, with one important difference: When a particle collides with its cell, instead of rebuilding its list of near neighbors as done in Ref. [11], the particle bounces off the cell wall as if the cell has hard walls.

This is just like the single-occupancy cell (SOC) system used in Ref. [179], however, with cells which do not necessarily cover space and which have the same shape as the particles, rather then being complex polyhedra. We now focus on solid-like systems, meaning that there is no or very little free diffusion, so that over long periods of time the particles do not move far away from their initial positions (i.e., the centers of the cells). When the cells are very large, that is, $\mu = \mu_{\max} \gg 1$, the SOC system is indistinguishable in its thermodynamic properties from the unconstrained system. In the limit $\mu \to 1$, the cells will become disjoint and the system becomes a collection of N independent particles, which can be treated analytically. We will assume that there exists a $\mu_{\min} > 1$ for which the cells are fully disjoint. This can always be assured by preparing the initial state more carefully or by shrinking the particles slightly.

The basic idea of the BCMD algorithm is to connect the two states, one with large-enough and the other with small-enough cells, via a thermodynamically reversible (quasi-equilibrium) path. We do this in the our MD simulation by simply slowly reducing the scaling factor during the course of the MD simulation

$$\mu = \mu_{\max} - \gamma_\mu t$$

with a constant *cell reduction rate* γ_μ. The algorithm developed in Ref. [179] already has allowance for a potentially changing shape of colliding particles, and therefore the introduction of a bounding cell with a changing shape poses no additional technical issues. During the course of the MD we can measure the average reduced pressure on the walls of a cell $p_c = P_c V_c/kT$ and then obtain the change in free energy as the work done in shrinking the cells

$$f_c\left(\Delta\mu_{\min}\right) - f = \int_{V_c^{\max}}^{V_c^{\max}} p_c \frac{dV_c}{V_c}. \tag{6.9}$$

We will assume that one can calculate $f_c\left(\Delta\mu\right)$ theoretically, and thus Eq. (6.9) gives us f. For example, for spheres of unit diameter we have the trivial result

$$f_c\left(\Delta\mu\right) = -d_f \ln \Delta\mu - \ln V_p, \tag{6.10}$$

and in general for particles with unit dimensions

$$f_c\left(\Delta\mu\right) = -d_f \ln \Delta\mu - f_c^J, \tag{6.11}$$

where f_c^J is a constant which depends on the exact particle shape, $f_c^J = \ln(\pi/6)$ for spheres.

The BCMD method is inspired by and closely related to the tethered-spheres method of Speedy [184]. Speedy describes this method for spheres as having the center of each particle tied with a tether to its initial position, and then considers changing the length of the tether in fixed finite increments, performing MD, and measuring the average tether force. He also points out the implicit presence of cells. One naturally extends the algorithm to nonspherical particles by adopting the cell picture. Additionally, we continuously change the cell size with time and adjust our MD algorithm accordingly, rather than using fixed finite steps. This offers more than just computational convenience. Namely, it gives a control parameter, γ_μ, which can be used to control the accuracy of the results: Reducing γ_μ improves the accuracy by both allowing for better equilibration and by increasing the number of collisions processed (and thus the overall statistical accuracy). More importantly, as we will explain in the next section, one can directly obtain the change in free energy from the change in kinetic energy of the system, thus obviating the need to define and measure[3] "pressure" and integrate it numerically after the simulation has finished, using potentially artificial interpolation. This not only improves the accuracy, but also enables one to calculate change in "free energy" in nonequilibrium processes where dynamics does matter and one cannot stop the simulation at specific points, such as the production of glasses by (relatively rapid) quenching.

We note in passing that recent Monte Carlo switch methods can more directly (and thus accurately) measure the difference in free energies between two states by directly switching from one state to the other [189]. It would be interesting to investigate such methods where the two states are the large and small cell states.

6.3.2 Elastic Collision Law

An important detail which has not been discussed carefully enough in the literature is the collision law for particle-particle and particle-wall collisions in the presence of a deforming particle shape or deforming wall. The collision law that we use is an elastic one, in the following sense: Linear and angular momentum are conserved, and the direction of the *relative* normal surface-to-surface velocity $v_\perp$ at the point of contact is reversed upon collision (but the magnitude is unchanged). It is most convenient to consider binary collisions only and then just take the limit that the "mass" of a hard wall goes to infinity. This collision law corresponds to a non-dissipative collision during which a purely normal exchange of momentum $\Delta\pi_c$ occurs and an additional work $W_c = v_\perp \Delta\pi_c$ is performed in order to maintain the rate of deformation of the particle/wall shape during the collision. Therefore, even though at first sight such collisions are not conserving kinetic energy, they are in fact conserving if one takes into account the work done by external "agents" to change the shape of the particles, cells, or container. The first work we are aware of where such collisions are considered is Ref. [12], and there an adjustable parameter h is left in the collision law, and

[3]For example, for nonspherical particles the force exerted by the particle on the walls of the cell is nonuniform, so that in principle a pressure tensor needs to be considered.

its choice is considered free. In fact, the actual choice made there corresponds to the above "conserving" (elastic) collision law, and is thus the best choice possible.

Using energy-conserving collisions enables us to calculate the change in free energy during a continuous transformation of the system by simply considering the change in kinetic energy of the particles. For example, the change in kinetic energy as the cells are shrunk during the cell method, or as particles are grown in size in order to produce jammed/glassy packings, or as the particles go from nonspherical to spherical, or as the shape of the simulation box shears from cubic to sheared-cubic, etc. Virtually any quasi-equilibrium transformation can be studied by simply performing it using collision-driven MD and letting the kinetic energy K increase or decrease spontaneously, to give the change in free energy

$$\Delta f = \frac{d_f}{2} \ln \frac{K}{K_0}. \tag{6.12}$$

In Section 6.3.3 we derive this formula for the BCMD method by looking at the algorithm in high-dimensional configuration space, giving us a unique perspective on the key elements of the algorithm that enable measurement of configurational volume to be done by performing (molecular) dynamics.

6.3.3 Billiards Algorithm for Volume Calculation

In this section we look at the BCMD algorithm in high-dimensional configuration space in order to understand how it actually measures the free volume and also how it can be applied as a general algorithm for measuring the volume of a (nearly) convex body in high dimensions. We also compare it to known efficient algorithms for measuring the volume of a convex body in high dimensions.

The motion of a nearly jammed hard-particle system in N_f-dimensional configuration space corresponds to a (point) billiard ball $\mathcal{B} \in \mathcal{R}^{N_f}$ elastically bouncing inside the nearly convex closed body $\mathcal{J}_{\Delta\mathbf{R}}$ whose volume $V = |\mathcal{J}_{\Delta\mathbf{R}}|$ we wish to calculate. We assume here that the dynamics of the ball is ergodic and can be analyzed within classical equilibrium thermodynamics. That is, the ball exerts a *uniform* pressure $P = kT/V$ on the walls when averaged over sufficiently many collisions with the wall, where the temperature $kT = 2K/N_f$ measures the kinetic energy of $\mathcal{B}$. The assumption of ergodicity is non-trivial and is usually assumed to hold when all of the walls of $\mathcal{J}_{\Delta\mathbf{R}}$ are semi-dispersing (concave) [190], which is true only for sphere packings. Physically, ergodicity is often attributed to the presence of a large number of particles, however, even with a large number of nonspherical particles the dynamics of $\mathcal{B}$ can be highly nontrivial and nonuniform.

The elastic nature of the ball implies that both the kinetic energy and the component of the momentum parallel to the wall are conserved upon collisions with a stationary wall of $\mathcal{J}_{\Delta\mathbf{R}}$. If a wall is moving due to, for example, growth in the size of the particles, when $\mathcal{B}$ collides with the wall it gets a velocity boost in the direction normal to the wall such that the normal velocity of the ball relative to the wall is reversed,

$$v_\perp^{\text{after}} - v_\perp = -\left(v_\perp^{\text{before}} + v_\perp\right),$$

where $v_\perp$ denotes the (local) normal velocity of the wall. This kind of dynamics ensures that the ball never sticks to the wall and also implies energy conservation in the following sense. The change in kinetic energy of the ball is

$$\Delta K_c = \frac{m}{2}\left(v_\perp^{\text{after}} - v_\perp^{\text{before}}\right)\left(v_\perp^{\text{after}} + v_\perp^{\text{before}}\right) = v_\perp \Delta\pi_c,$$

where $\Delta\pi_c$ is the momentum exchanged between the ball and the wall. Therefore ΔK_c is exactly the work done by the moving wall against the force exerted on it by the bouncing ball.

Now consider adding constraints on the displacements of the particles in addition to the nonoverlap conditions, with some parameter ξ determining how strong these additional constraints are. In isolation, the additional constraints limit the configuration to some neighbourhood $\tilde{\mathcal{J}}_{\Delta\mathbf{R}}(\xi)$ around $\mathbf{R}_J$, so that $\tilde{\mathcal{J}}_{\Delta\mathbf{R}}(\xi \to 0) = \{\mathbf{R}_J\}$ and $\tilde{\mathcal{J}}_{\Delta\mathbf{R}}(\xi \to \infty) = \mathcal{R}^n$. For example, in our MD algorithm, we add the constraints that each particle remain within its bounding cell, and can identify $\xi = \Delta\mu$. In a more general situation we could simply add the constraint that the configuration remain within a hypersphere of radius ξ centered at $\mathbf{R}_J$. Here we will assume that the volume $\left|\tilde{\mathcal{J}}_{\Delta\mathbf{R}}(\xi)\right|$ is known. The resulting available region of configuration space $\hat{\mathcal{J}} = \tilde{\mathcal{J}}_{\Delta\mathbf{R}} \cap \mathcal{J}_{\Delta\mathbf{R}}$ interpolates between $\mathcal{J}_{\Delta\mathbf{R}}$ for sufficiently large $\xi = \xi_{\max}$ and $\tilde{\mathcal{J}}_{\Delta\mathbf{R}}$ for sufficiently small $\xi = \xi_{\min}$. A two-dimensional illustration is provided in Fig. 6.1.

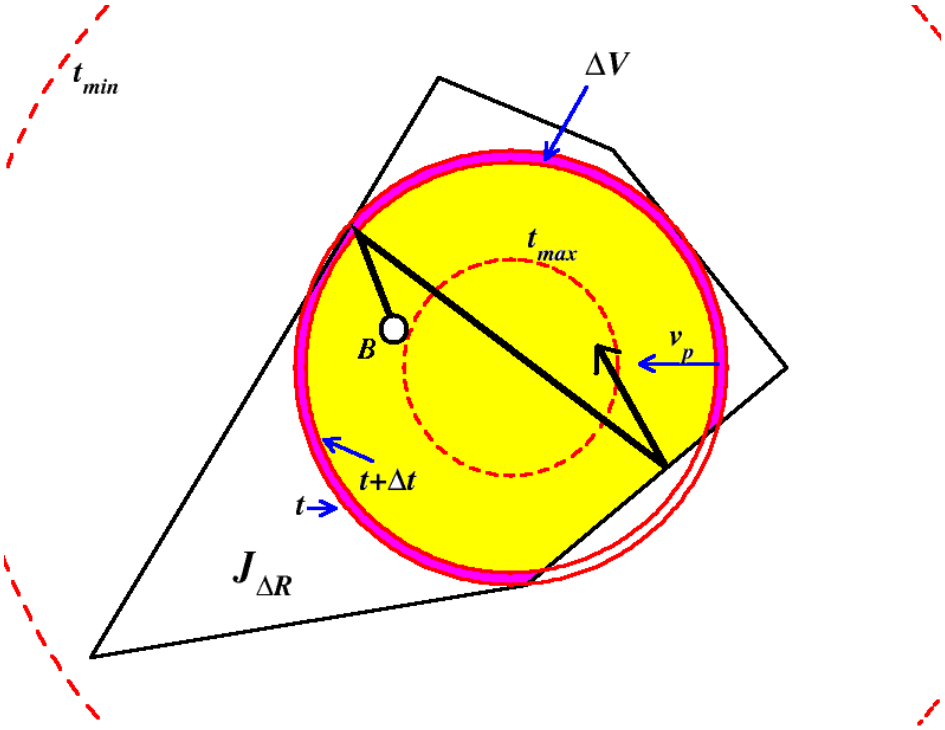

Figure 6.1: An illustration of the BCMD algorithm in a fictitious two-dimensional configuration space. The surface area of the polygon ($\mathcal{J}_{\Delta\mathbf{R}}$) is obtained by intersecting it with a shrinking disk of radius $\xi(t)$ centered inside the polygon [$\tilde{\mathcal{J}}_{\Delta\mathbf{R}}$], and integrating the surface area (colored purple) of the intersection region ($\hat{\mathcal{J}}$, shaded in yellow). This surface area is measured from the pressure exerted on it by a ball $\mathcal{B}$ bouncing elastically inside the intersection region.

Our algorithm for measuring $|\mathcal{J}_{\Delta\mathbf{R}}|$ essentially consists of dynamically decreasing $\xi(t)$ from $\xi_{\max}$ to $\xi_{\min}$ while following the collision dynamics of the billiard ball, and simply measuring the relative change in kinetic energy K/K_0 during the process. This change in kinetic energy comes because of the motion of the walls of $\tilde{\mathcal{J}}_{\Delta\mathbf{R}}$ as ξ changes. Namely, consider the short time interval from t to $t+\Delta t$ during which the volume $V(\xi) = \left|\hat{\mathcal{J}}(\xi)\right|$ decreases by $\Delta V = S v_\perp \Delta t$, where S is the surface area of $\tilde{\mathcal{J}}_{\Delta\mathbf{R}}$ contained within $\mathcal{J}_{\Delta\mathbf{R}}$. The increase in kinetic energy of $\mathcal{B}$ during this interval is

$$\Delta K = v_\perp \Delta\pi = \frac{\Delta\pi}{S\Delta t}\Delta V = P\Delta V = kT\frac{\Delta V}{V(\xi)} = \frac{2K}{N_f}\frac{\Delta V}{V(\xi)}.$$

Taking the limit $\Delta t \to 0$ we get a differential equation whose solution gives the volume $V = V(\xi_{\max})$,

$$\ln\frac{V}{V_{\min}} = \frac{N_f}{2}\ln\frac{K}{K_0}, \tag{6.13}$$

where $V_{\min} = V(\xi_{\min}) = \left|\tilde{\mathcal{J}}_{\Delta\mathbf{R}}(\xi_{\min})\right|$ is known. Equation (6.13) was used to obtain Eq. (6.12).

The basic elements of the algorithm are essentially identical to well-known polynomial algorithms for computing the volume of a convex body in high dimensions [185]. The majority of these algorithms (other than the best known algorithm from Ref. [191]) use a random walk through the convex body $\hat{\mathcal{J}}$ to essentially estimate the fraction of the surface area of $\tilde{\mathcal{J}}_{\Delta\mathbf{R}}$ that is contained in $\hat{\mathcal{J}}$. The total volume is then calculated as a product of such fractions, that is, the logarithm of the volume is calculated as a sum (i.e., an approximation to the integral of the surface area). In our algorithm we use (ergodic) molecular dynamics to estimate and integrate the fractions as ξ is changed dynamically. We believe that the powerful mathematical techniques

developed for the analysis of random walk algorithms together with techniques from theory of ergodic billiards could be applied to analyze the BCMD algorithm and come to quantitatively understand its trade-offs between accuracy and complexity.

6.3.4 Algorithmic Details

We now turn to some technical details needed when using the BCMD algorithm in practice.

The first concerns the choice of boundary conditions. We have employed periodic boundary conditions here, which gives d trivial translational degrees of freedom to the unconstrained system. However, when cells are present, they are fixed in space, breaking the translational symmetry. The usual approach to eliminating trivial translations is to freeze the center of mass of the particles. In traditional conservative MD this is done easily by simply ensuring that the initial velocities add to zero. However, collisions with the hard walls of the bounding cells do not conserve linear momentum. This leads to an artificial increase of the collision rate of the particles with the cells, particularly at high densities (i.e., near jamming), because the center of mass oscillates as the excess linear momentum is distributed among the particles through binary collisions. Speedy proposed to artificially enforce a frozen center of mass by correcting the velocity of all of the particles whenever a particle collides with a bounding cell. This is rather inefficient; instead, we have chosen to evaluate the kinetic energy K in the center-of-mass system. For sufficiently large systems, the corrections due center-of-mass oscillations is small. Another way to eliminate trivial translations is to freeze one particle (give it infinite mass), which is especially useful near jamming. The two methods are not equivalent in terms of free energy. The difference in free energy due to this change of boundary conditions is shown to be $d \ln N/N$ in Section 6.3.4.1.

The second technical issue concerns that of adjusting the reduction rate γ_μ. The rule of thumb is to keep it small enough so that the actual linear velocity of the moving cell walls is small (say 10^{-3} times smaller) compared to the average particle velocity. Keeping γ_μ constant, however, uses too much computer time for large μ, a region which actually contributes little to the total configurational volume, and too little on small μ, where more precision is actually required. We have chosen to implement an periodic adaptive change of the reduction rate[4]

$$\gamma_\mu(\mu) = \gamma_\mu(\mu_{\max}) \left(\frac{\Delta\mu}{\Delta\mu_{\max}} \right)^{\vartheta},$$

where ϑ is an exponent, which we have usually kept in the range $0.25-1$, depending on the size of the system and the initial $\gamma_\mu(\mu_{\max})$ as well as the ratio $\Delta\mu_{\min}/\Delta\mu_{\max}$. It is important to consider the convergence of the BCMD method in more detail using existing powerful mathematical techniques and try to determine a theoretically sound $\gamma_\mu(\mu)$ that would provide a *desired* total error in f with reasonable confidence[5]. At present, we do not have such a theory and therefore are unable to quote rigorous error bars on our results other than statistical fluctuations among samples and also estimated errors from different runs with successively smaller $\gamma_\mu(\mu_{\max})$ or larger ϑ.

The final technical point concerns the actual practical implementation of the BCMD algorithm for ellipses/ellipsoids. As explained in the second part of Ref. [11], the numerical calculation of overlap potentials in the case of one ellipsoid contained inside another is numerically unstable and causes difficulties in the correct prediction and processing of collisions with the bounding cells. While the occasional numerical failures can be tolerated when the bounding neighborhoods are used only for speeding up the neighbor search in the MD algorithm, they become an insurmountable obstacle with the cell method, especially when the cells are disjoint or nearly disjoint. This is because a wrong collision prediction can lead to a particle leaving its cell permanently. We have not really found a satisfactory solution to this problem, and have been forced to use reduction rates γ_μ sufficiently fast so that such failures are unlikely to appear.

An additional technical difficulty when dealing with ellipsoids is the fact that there is no analytical equivalent to Eq. (6.10). One possibility is to calculate f_c for a particle inside its cell numerically by following a reversible path to a known system. Two possibilities include the case of an infinitesimally small particle, reducing to a one-particle ideal gas inside a container of volume V_c, or the case of a spherical particle inside a spherical cell. The first system can be reached by shrinking the size of the particle slowly,

[4]Note that in the actual implementation γ_μ is updated periodically, rather than continuously.

[5]It is reasonable to assume that an error of the order $\ln N/N$ is acceptable since finite-size effects are of this order [192]. This means that the configurational volume only needs to be determined to within a factor of $\ln N$.

and the second by reducing the asphericity. We have implemented both options in our MD codes, however, as explained above, numerical difficulties are a problem. For this reason we (semi)analytically consider the case of an ellipse of aspect ratio $\alpha = 2$ inside a bounding cell in Section 6.3.4.2, to obtain $f_c^J \approx 2.269$. We have verified this result (to within the precision possible) with the MD algorithm by calculating the change in free energy as the ellipses and the cells were slowly transformed into disks while keeping $\Delta\mu$ fixed by shrinking the larger semiaxis to unity. Namely, for small-enough $\Delta\mu$ the change in free energy during such a transformation is known analytically to be

$$\Delta f_{\text{sphere}} = -\ln\left[\pi\left(\pi\Delta\mu^2\right)\right] - \left[-3\ln\Delta\mu - f_c^J\right] = \ln\Delta\mu + f_c^J - 2\ln\pi,$$

where we have used the fact that a disk has full rotational freedom (contributing a factor of π to the configurational volume).

6.3.4.1 Frozen Particle vs. Frozen Center of Mass

The easiest way to eliminate the d trivial translational degrees of freedom is to freeze one of the particles, say the first one, $\Delta\mathbf{R}_1 = \mathbf{0}$, which simply amounts to deleting the d rows of the rigidity matrix corresponding to that particle. Another method is to freeze the center of mass, $\sum_{i=1}^{N}\Delta\mathbf{R}_i = \mathbf{0}$, which amounts to adding d columns to the rigidity matrix. When performing MD in the NVT ensemble, the center of mass is frozen by virtue of momentum conservation. However, it is also possible to freeze a particle, so that collisions with it no longer conserve momentum.

The two different boundary conditions do not give the same free energy. However, the difference can be readily calculated. When one particle is frozen the natural set of independent coordinates are the displacements of the other $N-1$ particles (coordinate system F). When the center of mass is frozen we can choose the same set of independent coordinates since the displacement of the frozen particle is just the negative sum of the other displacements (coordinate system CM). The transformation between these two coordinate systems is

$$\Delta\mathbf{R}^{CM} = \Delta\mathbf{R}^F - \frac{1}{N}\sum_{i=2}^{N-1}\Delta\mathbf{R}_i^F = \mathbf{J}\Delta\mathbf{R}^F$$

where the Jacobian $\mathbf{J} = \mathbf{I} - \Delta\mathbf{J}$ and $\Delta\mathbf{J}$ is a block matrix made up of $d\times d$ diagonal blocks which have $-1/N$ on the diagonal. The determinant of this Jacobian can be calculated to be

$$|\mathbf{J}| = \frac{1}{N^d},$$

which relates the ratios of configurational volumes in the two different coordinate systems. For the free energy per particle we thus have the transformation

$$f^{CM} = f^F - \frac{d\ln N}{N},$$

the difference of course vanishing in the infinite system limit. A term of the order $\ln N/N$ is the leading finite-size correction in other free-energy methods as well [192].

6.3.4.2 The Free Energy of an Ellipse in a Cell

In this Section we consider analytic expressions for the free energy of an ellipse of aspect ratio α enclosed in a fixed cell that is of the same shape but $1+\Delta\mu$ times larger, in the limit of small $\Delta\mu$. In this limit we can linearize the nonoverlap condition between the inner and outer ellipses. However, unlike the case of two disjoint ellipsoids, there is no unique contact point around which to linearize the nonoverlap. Instead, every point on the ellipse is a potential contact point, and therefore the linearization of the nonoverlap condition consists of infinitely many linear inequalities. For a given potential contact $\mathbf{r}_c$ on the ellipse with a (normalized) normal vector $\mathbf{n}$, the linearized nonoverlap condition is

$$\mathbf{n}^T\left(\Delta\mathbf{r} + \Delta\boldsymbol{\varphi}\times\mathbf{r}_c\right) \le \Delta l = \Delta\mu\left(\mathbf{n}^T\mathbf{r}_c\right), \tag{6.14}$$

where $\Delta\mathbf{r}$ and $\Delta\boldsymbol{\varphi}$ are the translational and orientational displacement of the ellipse, and Δl is the gap between the "surfaces" of the inner and outer ellipse. Since the ellipse is a smooth convex shape the mapping between $\mathbf{n}$ and $\mathbf{r}_c$ is unique [11],

$$\mathbf{r}_c = \frac{\mathbf{X}^{-1}\mathbf{n}}{\sqrt{\mathbf{n}^T\mathbf{X}^{-1}\mathbf{n}}},$$

where $\mathbf{X} = \mathrm{Diag}\{1, \alpha^{-2}\}$. The three dimensional convex set of allowed displacements of the inner ellipse $\mathcal{J}_{\Delta\mathbf{q}}$ is bounded by the infinite collection of linear inequalities in $\Delta\mathbf{q} = (\Delta\mathbf{r}, \Delta\boldsymbol{\varphi})$ as given by Eq. (6.14), for all $\mathbf{n} = (\cos\theta, \sin\theta)$. We are interested in calculating its volume. It is clear that its extents scale linearly with $\Delta\mu$, and its volume is therefore $V(\Delta\mu) = \Delta\mu^3 V(\Delta\mu = 1)$. We can thus focus on the case $\Delta\mu = 1$.

It is not trivial to calculate the volume $V(\mathcal{J}_{\Delta\mathbf{q}})$. We approached the problem by considering a fixed $\Delta\varphi$ and then attempting to find an analytic representation of the resulting planar set of feasible $\Delta\mathbf{r}$, $\mathcal{J}_{\Delta\mathbf{r}}$, so that its area can be calculated. For example, for the trivial case $\Delta\varphi = 0$, $\mathcal{J}_{\Delta\mathbf{r}}$ is simply an ellipse of aspect ratio α (this is easy to see geometrically). The maximum possible $\Delta\varphi$ corresponds to only $\Delta\mathbf{r} = 0$ being feasible, and in the particular case $\alpha = 2$ it is $\Delta\varphi_{\max} = 4/3$. Unfortunately, for sufficiently large $\Delta\varphi$ the convex $\mathcal{J}_{\Delta\mathbf{r}}$ is not smooth, so that it cannot be parameterized by a continuous family of normal vectors. We therefore partially resorted to numerical calculations of the area of $\mathcal{J}_{\Delta\mathbf{r}}(\Delta\varphi)$, and then numerically integrated

$$V = \int_{-\Delta\varphi_{\max}}^{\Delta\varphi_{\max}} \mathcal{J}_{\Delta\mathbf{r}} d(\Delta\varphi).$$

For the particular case of an ellipse with semiaxes 1 and 2 we obtained the final result

$$f_c = -2.268 - 3\ln\Delta\mu.$$

It is desirable that the general case be worked out analytically, particularly for ellipsoids in three dimensions.

6.3.5 Illustrative Example: Dense Hard-Sphere Liquid

As an illustration of the BCMD algorithm we apply it to a system which is at the limit of the method's applicability: A dense hard-sphere liquid near the melting point, at $\phi = 0.5$. This is still a liquid and thus the particles diffuse freely given sufficient time, and the cell method is not rigorously applicable. However, adding the cells stabilizes the liquid in the neighbourhood of its initial configuration for long periods of time, thus allowing the measurement of the free-energy for a metastable cell-constrained liquid (CCL). This free-energy is certainly larger than that of the unconstrained liquid, which has more free volume available to explore. If one assumes that one can divide this liquid free volume into more-or-less N_b equivalent "basins", each basin corresponding to a single "glassy" jammed configuration, then the loss of entropy due to the presence of the cells would be of the order of $N^{-1}\ln N_b$. This assumption about dividing configuration space into statistically equivalent compartments or basins, though commonly used in the literature, has not really been justified for the hard-sphere system (where entropy cannot be "turned off" by going to zero temperature).

In Fig. 6.2 we show the results of the cell method as applied to a stable liquid. We show both the pressure on the cell walls p_c, which dominates at $\mu = \mu_{\min}$, and the internal pressure p, which dominates at $\mu = \mu_{\max}$, when it approaches the true (unconstrained) liquid pressure. It is seen that the cell-wall pressure shows a minimum, as first observed in Ref. [184], when $\mu \approx 2$, that is, when a bounding cell corresponds to the exclusion sphere of its particle in the initial liquid configuration. For $\mu > 2$ it appears that particle rearrangements take place which increase p_c. Even though p_c has a minimum, it remains positive throughout due to the diffusion of the particles, and therefore one cannot measure the true free-energy of the liquid using the BCMD algorithm. Instead, we show in Fig. 6.2 the excess free-energy of the cell-constrained liquid relative to the ideal gas,

$$\Delta f_{CCL}(\mu) = f(\mu) - f_{\text{ideal}} = \Delta f_c(\Delta\mu_{\min}) - d\int_{\mu_{\min}}^{\mu} p_c \frac{d\mu}{\mu}.$$

Here the ideal-gas free energy is trivially

$$f_{\text{ideal}} = -\frac{1}{N}\ln\frac{V^N}{N!} \approx -\ln\frac{V}{N} - 1 = -\ln\frac{V_p}{\phi} - 1 \tag{6.15}$$

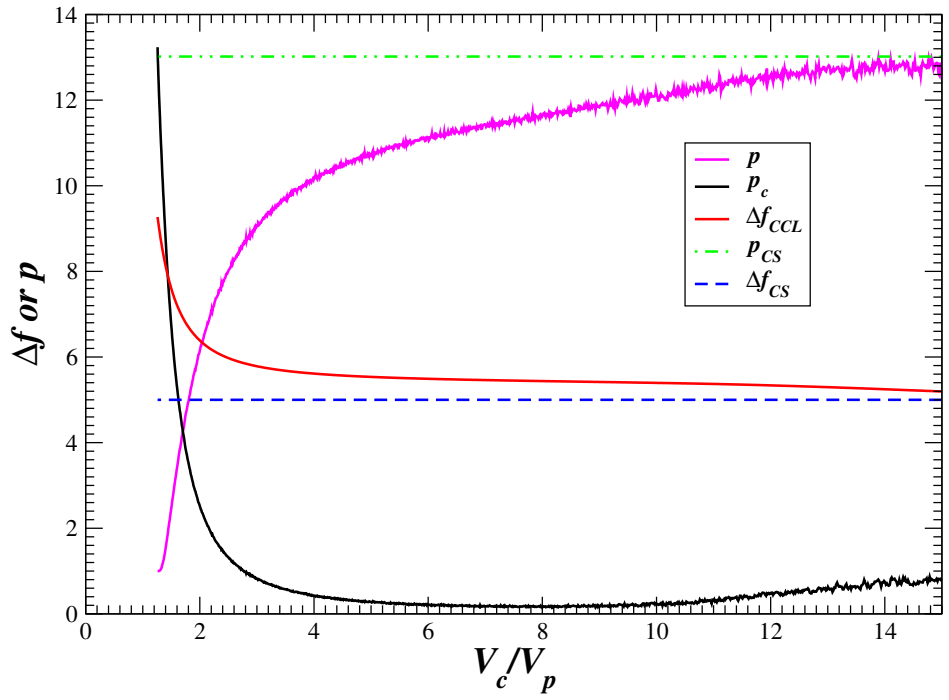

Figure 6.2: Illustration of the BCMD algorithm for a dense equilibrium hard-sphere liquid close to the melting point, $\phi = 0.5$. We show as a function of the relative size of bounding cells the internal pressure p as well as the cell-wall pressure p_c, along with the calculated excess free energy of the cell-constrained liquid $\Delta f_{CCL}(\mu)$. The results from the semi-empirical CS EOS is shown for comparison. It is seen that the pressure on the cell walls goes through a minimum but is always positive due to the free diffusion of the particles. The values of p_c for large μ are more difficult to measure accurately since they depend on particle diffusion and they have not been studied carefully.

and therefore

$$\Delta f_c(\Delta\mu) = -\ln \Delta\mu^d V_p + \ln \frac{V_p}{\phi} + 1 = -\ln\left(\phi \Delta\mu^d\right) + 1. \tag{6.16}$$

In the figure we compare $\Delta f_{CCL}(\mu)$ to the results predicted by the relatively-accurate Carnahan-Starling (CS) equation of state (EOS) for the liquid [193],

$$p_{CS} = \frac{1+\phi+\phi^2-\phi^3}{(1-\phi)^3} \text{ and } f_{CS} = \int_o^\phi \frac{p-1}{\phi} = \frac{\phi(4-3\phi)}{(1-\phi)^2}.$$

While $\Delta f_{CCL}(\mu \approx 2)$ is a reasonable approximation to the true liquid free energy at such high liquid densities, the difference between the two, called *communal entropy* by Kirkwood for his single-occupancy cell liquid model, is not zero and has been estimated a long time ago [194]. Additionally, due to somewhat arbitrary choice of the cell size (unlike in Kirkwood's model where the cell partitioning is fixed), the exact value is not well-defined, and therefore interpretations in terms of number of "inherent-structures" that the liquid samples are questionable. We will discuss these issues in significant detail in Chapter 13.

6.3.5.1 Rattlers

Computer-generated packings often have *rattlers*, particles which do not participate in the jammed backbone of the packing and are free to rattle inside a cage formed by the jammed backbone. Near the jamming limit, systems with rattlers fall out of equilibrium on the time-scale of the BCMD algorithm because the rattlers heat up less than the other particles and the equipartition theorem no longer holds. In essence, two time-scales emerge, a short time scale for the backbone and a long timescale for the rattlers, and the theory in

Section 6.3.3 no longer applies. The usual "velocity-rescaling" thermostat, i.e., a uniform cooling of the whole system [11], is not appropriate since it lowers the rattler temperature more than the backbone. A better thermostat to completely reinitialize all the particle velocities to random values as a way to cool down the system. This strategy affects the natural dynamics of the system but is necessary because the separation in time-scales between the dynamics of the rattlers and the jammed particles makes the system nonergodic on the time-scale of the simulation.

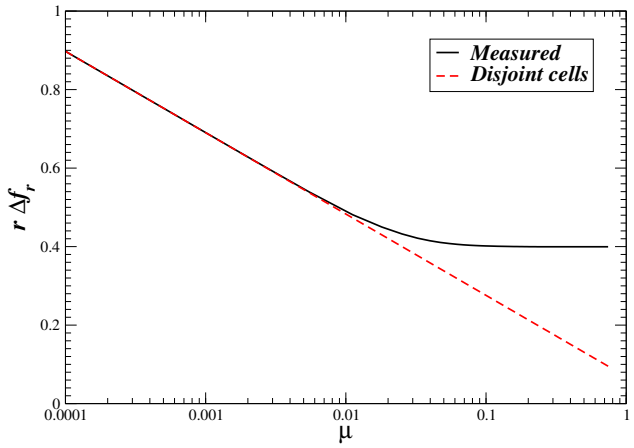

Figure 6.3: The measured contribution of the 30 rattlers to the excess free energy for a jammed disordered packing of $N = 1000$ spheres, $r\Delta f_r(\mu)$, as a function of the bounding cell relative size. The non-rattling particles, i.e., the particles in the jammed backbone, were frozen. For small μ the rattler cells are disjoint from the backbone in which case the free energy is known analytically (red line). For large μ the presence of the cells does not make a difference and one can extract the appropriate $r\Delta f_r(\mu) \approx 0.4$ to use in Eq. (6.17).

For hard spheres in three dimensions, Eqs. (6.2) and (6.15) suggests that in the neighbourhood of the jamming point the excess free energy (relative to the ideal gas) increases with density as

$$\Delta f_J(\phi) = b\left\{-3\ln\left(1-\frac{\phi}{\phi_J}\right) - \left[\ln\left(\frac{2b\phi_J}{9\pi}\right) - 1\right] - f_J\right\} + (1-b)\Delta f_r, \tag{6.17}$$

where $b \approx 1 - 0.025$ is the fraction of particles participating in the jamming backbone, and Δf_r is the contribution due to the trapping of rattlers inside the cage of jammed particles (assumed to be independent of ϕ sufficiently close to the jamming point). We can measure the rattler contribution $(1-b)\Delta f_r = r\Delta f_r$ by using the BCMD method as follows. We freeze all non-rattling particles and then perform the BCMD algorithm (in reverse) on just the rattlers, starting with bounding cells small enough so that they are disjoint with other particles, and ending with cells which are large enough so that they do not affect the rattler free volume. The resulting excess free energy for the cell-constrained rattler system, scaled by $r = 1-b$, is shown in Fig. 6.3. It is seen that even though the fraction of rattlers is small, their excess free energy relative to the ideal gas is rather high (i.e., the rattler cages are rather tight) and thus the contribution to $\Delta f_J(\phi)$ is significant and cannot be neglected if one wants to calculate free energies to accuracy better than $0.1Nk_BT$.

6.4 Results

In this section we present the first testing of the BCMD algorithm against rigorously-known free energies by applying it to isostatic jammed systems. For these systems, as explained in Section 6.2, the free-energy in the jamming limit can be found from the volume of a simplex polytope, which can be calculated exactly easily even for large systems. In the jamming limit we choose to freeze one particle in order to eliminate trivial translations, so that the effective number of particles N is one less than the true number of particles. Assuming that we start the cell method at exactly the jammed configuration $\mathbf{R}_J$, the cells will become

disjoint when $\Delta\mu_{\min} = \delta$. The total change in free energy during the course of shrinking the cells is

$$\Delta f = \frac{d_f}{2} \ln \frac{K}{K_0} = \left(-f_c^J - d_f \ln \delta\right) - \left(-d_f \ln \delta - f_J\right) = f_J - f_c^J, \tag{6.18}$$

where we recall that $f_J = (N-1)^{-1} \left[\ln |\mathcal{P}_{\mathbf{x}}|\right]$, which can be calculated easily if $\mathcal{P}_{\mathbf{x}}$ is a simplex, and f_c^J is known for a given particle shape. We use Eq. (6.18) to rigorously validate and test the accuracy of the BCMD algorithm. As an initial test, we compare against the highest-precision MC results for hard-sphere crystals that have appeared in the literature. In subsequent work we will apply the validated BCMD algorithm to dense liquid and glassy systems further from the jamming point in order to better understand the thermodynamics of hard-particle metastable liquids and glasses [57, 83].

6.4.1 Hard-Sphere Crystals

The entropy of hard-sphere crystals as a function of density has approximately been known from the early days of computer simulations [194]. Interest in the problem has surged recently trying to access whether the FCC or the HCP crystal has lower free energy (and is thus the stable solid phase) [44, 179, 54, 45, 46]. The difference between the two is very small and the literature is filled with contradicting claims and underestimated error bars. The most accurate results are those produced by MC switch algorithms, which do not determine an absolute free-energy but rather directly measure the entropy difference between FCC/HCP, producing a difference of about $1.1 \cdot 10^{-3} NkT$ in favor of FCC configurations in the jamming limit [45].

The results for the pressures and absolute value of the free energies have been summarized and processed in Ref. [54], and the most precise quoted MC result for the excess free energy of the FCC crystal over the ideal gas is close to the melting point, $\Delta f_{FCC}(\phi = 0.545) = 5.91916(1)$. At the same density, the BCMD algorithm with $\Delta\mu_{\max} = 1$, $\gamma_\mu\left(\mu_{\max}\right) = 0.001$ and $\vartheta = 1$ produces 5.919(0), which is in excellent agreement. In the jamming limit, from the data in Ref. [54] we estimate

$$\left[f_J - f_c^J\right]_{FCC} = 2.160 \pm 0.001,$$

where the exact error bars are difficult to cite since most of the focus has been near the melting point rather than close packing. Our most precise BCMD runs have involved about $10,000$ spheres at $\delta = 10^{-6}$ and $\Delta\mu_{\max} = 10^{-5}$ with $\gamma_\mu\left(\mu_{\max}\right) = 0.001$ and $\vartheta = 0.5$, giving $\left[f_J - f_c^J\right]_{FCC} = 2.1599 \pm 0.0005$ and $\left[f_J - f_c^J\right]_{HCP} = 2.1593 \pm 0.0005$, where as pointed out previously the true error bars are not known. The excellent agreement between our results and the published ones serves as a validation of the algorithm and our implementation. Unfortunately, the true volumes of the jamming polytopes corresponding to the jamming limit cannot be calculated exactly for anything but small systems [46], since their combinatorial complexity is exponential.

The difference in free energy between the FCC and HCP lattices as calculated with the BCMD algorithm is of the same sign and order of magnitude as the literature results. However, the error bars are significant and although it is possible to run longer and larger simulations to improve them we believe that without a rigorous error analysis one should not really engage in trying to decide the "winner" among these two lattices. Additionally, the impact of vacancies needs to be accessed more carefully. Furthermore, the HCP lattice does not have full cubic symmetry and therefore the crystal of minimum free energy is slightly compressed along the hexagonal symmetry axis. Preliminary calculations indicate that both of these effects produce corrections about an order of magnitude smaller than the reported difference between FCC and HCP, however, more rigorous results are lacking. Finally, the importance of randomly stacked arrangements is not obvious, despite the natural assumption that one of the two extremal stacking arrangements (ABA and ABC) will produce the minimal free energy.

It is interesting to observe that the BCMD method works remarkably well for a solid like the FCC crystal, even at densities as low as the melting density. This is because the displacements of the particles from the ideal lattice positions are small (usually Gaussian in density functional approximations) and thus well-localized around the equilibrium locations. This makes $p_c(\Delta\mu)$ a rapidly (faster than exponential) decaying function which is easy to integrate accurately with little computational effort, as illustrated in Fig. 6.4. This is not the case for disordered packings of monodisperse spheres, where the particle displacements

are significant and $p_c(\Delta\mu)$ decays slowly in a power-law manner. Two-dimensional crystals which lack long-range translational order are also expected to show significant particle delocalization at densities away from the jamming density.

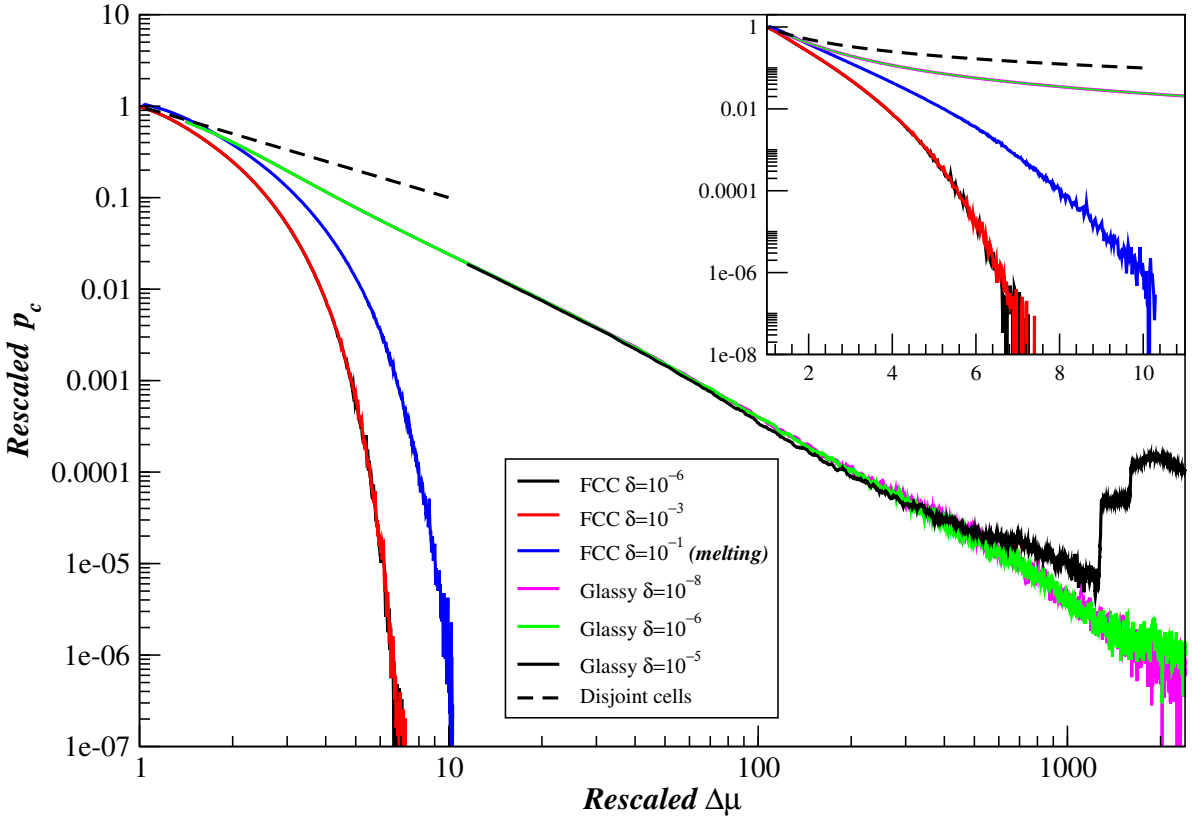

Figure 6.4: Rescaled values of the cell-wall pressure $\tilde{p_c} = \delta p_c$ as a function of rescaled $\Delta\tilde{\mu} = \Delta\mu/\delta$ for several different values of δ for both an FCC crystal packing ($N \approx 10^4$) and an isostatic disordered packing ($N \approx 10^3$, see Section 6.4.2).We also show for comparison the expected asymptotic dependence in the limit of disjoint cells, $\tilde{p_c} = \Delta\tilde{\mu}^{-1}$, as well as a logarithmic-scale view of the data in the inset. In the jamming limit, where the polytope picture applies, we expect these rescaled curves to follow a master curve independent of δ.

6.4.2 Isostatic Jammed Packings: Spheres

We have generated isostatic disordered jammed sphere packings of $N = 1000$ spheres using the procedures described in Ref. [75]. These packings have a number of *rattlers*, particles which are free to move inside a cage of jammed particles. We remove these rattlers and freeze one particle before applying the BCMD method in order to make the packings fit the simplex picture exactly, and then we compute the exact free-volume in the jamming limit from the volume of the simplex using Eq. (6.5). As an illustration, we consider one of these packings of $N = 978$ particles (1000 particles with the rattlers), and get $f_J - f_c^J \approx 4.9479$ (in the frozen center-of-mass coordinate system $f_J - f_c^J \approx 4.9690$). Table 6.1 shows the results obtained from the BCMD algorithm at $\delta = 10^{-8}$ with several different choices of the main parameters, illustrating the excellent agreement with the exact result, especially in the limit of infinitely long runs[6]. The primary error is likely due to statistical fluctuations in the kinetic energy due to the relatively small number of particles. It is

[6]One could potentially improve the precision by extrapolating to, for example, $\gamma_\mu\,(\mu_{\max}) \to 0$, however here we are primarily interested in testing the method rather than obtaining the most accurate result.

important to note that the simplex polytopes for disordered packings are very elongated in certain directions and therefore require significantly larger $\Delta\mu_{\max}/\delta$ than crystal packings.

δ	$\Delta\mu_{\max}$	$\gamma_\mu(\mu_{\max})$	ϑ	$f_J - f_c^J$
10^{-8}	$2.5 \cdot 10^{-5}$	0.1	1.00	4.9476
10^{-8}	$2.5 \cdot 10^{-5}$	0.01	0.25	4.9480
10^{-8}	$1.5 \cdot 10^{-5}$	0.001	0.25	4.9458
10^{-8}	$2.5 \cdot 10^{-5}$	0.01	0.50	4.9493
10^{-8}	$2.5 \cdot 10^{-5}$	0.01	0.75	4.9485
10^{-8} (frozen CM)	$2.5 \cdot 10^{-5}$	0.01	0.50	4.9624
10^{-7}	$2.5 \cdot 10^{-4}$	0.01	0.25	4.9498
10^{-6}	$2.5 \cdot 10^{-3}$	0.01	0.25	4.9414

Table 6.1: Results of the BCMD algorithm for an isostatic disordered jammed sphere packings of $N = 1000$ spheres, with different parameters for the algorithm. The rattlers were removed and one particle was frozen, except in one case where the center of mass was frozen. In principle the exact answer $f_J - f_c^J = 4.9479$, as calculated from the volume of the jamming simplex, should be reached for infinitely long runs as $\delta \to 0$. Typical running times are of the order of several hours to a day on an 1666MHz AMD Athlon PC running Linux.

In Fig. 6.4 we show the cell-wall pressure as a function of the cell size, illustrating the fact that it decays very slowly. In fact, this particular isostatic packing (recall, after rattlers have been removed) unjams when $\delta \approx 10^{-5}$ without the presence of the cell walls to stabilize it. The difference between the crystal and glassy packings is striking. For the crystal packings $\tilde{p_c}$ decays at least exponentially, even at densities as low as the melting point ($\phi \approx 0.545$), making it easy to obtain the free energies. However, for the disordered packings the decay is power-law like, with long tails which are more difficult to integrate accurately. Studies of glasses further away from the jamming point will be presented in separate work [83], here we focus on the region where the simplex picture applies.

6.4.3 Isostatic Jammed Packings: Ellipses

As we discussed already, nonspherical particles pose problems to the polytope picture in general and cannot usually be analyzed within the polytope picture even in the jamming limit. However, it is possible to (artificially) make isostatic jammed packings which do fit the polytope picture (i.e., first-order rigid packings [73]). The idea here is to obtain a packing which has $6 - 2/N \approx 6$ contacts per particle, since this is the number required for an isostatic packing ($M = N_f = 3N - 2 + 1 = 3N - 1$). We cannot use a crystal packing here since those have exactly 6 contacts per particle and are thus hyperstatic by one contact, and typical disordered packings of ellipses are usually rather hypostatic, having too few contacts for first-order rigidity. We therefore started with a triangular lattice of $N = 100$ disks, replaced each disk with an ellipse of aspect ratio 2, and then compressed the packing using MD to jamming to obtain a partially disordered packing. Even these packings had $M < 3N - 1$, however, by freezing 4 of the particles we were able to obtain packings which had matched number of degrees of freedom and contacts. We then calculated the volume of the simplex, obtaining $f_J - f_c^J = 3.6693$ for the packing illustrated in Fig. 6.5. Running the BCMD algorithm at $\delta = 10^{-4}$ with $\Delta\mu_{\max} = 1.5 \cdot 10^{-2}$ with $\gamma_\mu(\mu_{\max}) = 0.001$ and $\vartheta = 0.5$ gave $f_J - f_c^J = 3.61 \pm 0.01$, in reasonable agreement with the known answer. Higher accuracy is hard to achieve at present due to numerical difficulties in the implementation discussed previously. We note that when using a frozen particle to eliminate trivial translations only its centroid must be fixed, but not its orientation, since the orientation of the coordinate system is fixed by the orientation of the periodic unit cell.

An additional possibility which deserves further investigations is using transformations of the particle shape without any cell models in order to measure free energies. Examples would include superellipses transforming into ellipses, or ellipses transforming into spheres, etc., while measuring the change in kinetic energy. The difficulty in such simulations is the identification of reversible paths free of discontinuous phase transitions, however, the possible increased computational efficiency as well as the ability to study a whole

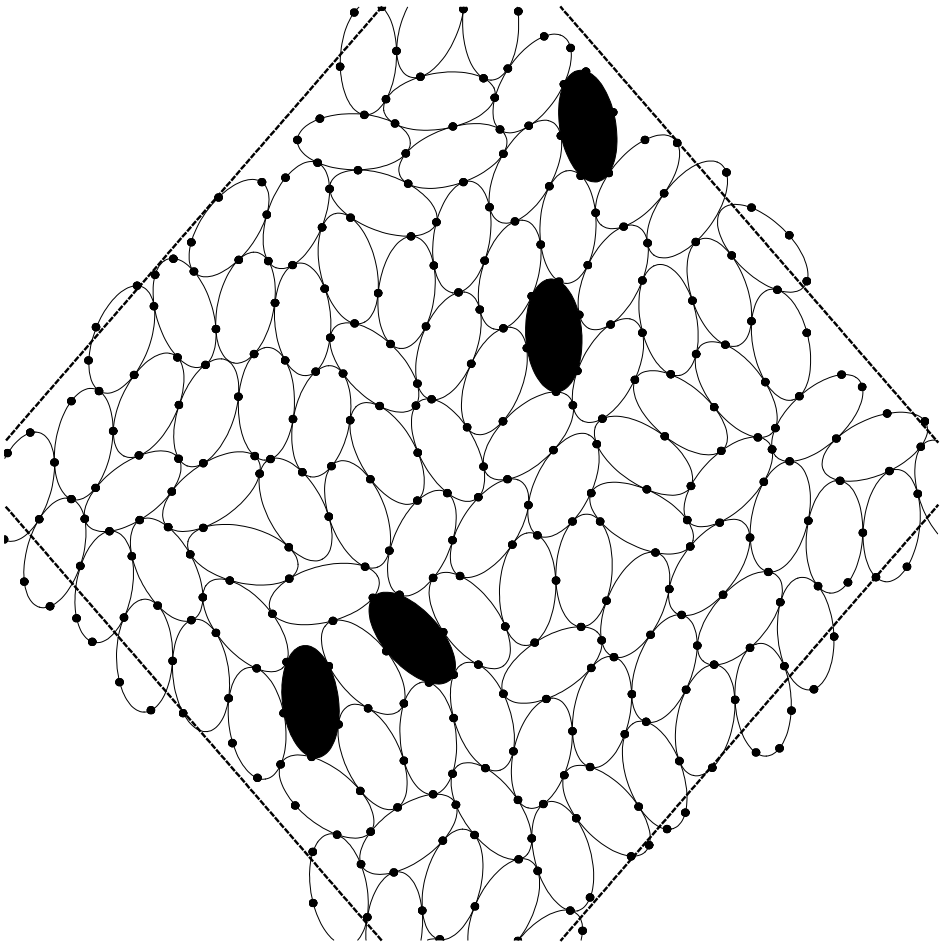

Figure 6.5: A sample isostatic jammed (first-order rigid) packing of $N = 96$ free and four frozen ellipses of aspect ratio $\alpha = 2$, used to test the BCMD algorithm on nonspherical particles. This packing was produced using a complicated artificial procedure as described in the body of the paper, and is not typical of disordered ellipse packings, which have fewer contacts per particle.

range of particle shapes in one simulations make the method very attractive.

6.5 Conclusions

We developed an MD algorithm for measuring the free energy of systems of hard particles where diffusion is negligible. By exploiting free-volume theories for the jamming limit of hard-particle packings we were able to exactly calculate the asymptotic behavior of the free energy of model disordered systems, namely isostatic jammed sphere and ellipsoid packings. This provided us with a rigorous verification of the numerical accuracy of the BCMD algorithm. However, a quantitative understanding of the statistical and systematic errors in the BCMD algorithm is missing. Additionally, quantitative understanding of the connection between hard-particle systems (where free energy is the volume of a bounded high-dimensional body) and low-temperature soft-particle systems (where free energy is a weighted volume integral) is important to elucidating the usefulness and limitations of hard-particle models of real materials. The algorithm presented here is successfully applied to the study of the glass transition in hard-sphere systems in subsequent work [57, 83], as described in Chapter 13.

Part II

(Nearly) Jammed Packings: Applications

Chapter 7

Packings of Spheres and Ellipsoids in Finite Containers

In this Chapter we compare the properties of computer-generated packings of hard spheres and ellipsoids with experimental results for packings of ball bearings, M&M'S Candies®, as well as manufactured ellipsoids [78, 79]. For M&M'S Candies® we can experimentally attain sufficiently large packings to obtain reasonable estimates of the bulk packing fractions, which can easily be obtained in simulations by using periodic boundary conditions [78]. The simulations in Ref. [78] predicted the ellipsoid shape which gives the highest random packing density, namely, that ellipsoids with axes ratios near $1.25 : 1 : 0.8$ form bulk amorphous packings as dense as the densest crystal packing (FCC) of spheres, $\phi_{\rm bulk} \approx 0.735$. We demonstrate that such dense packings are realizable by fabricating ellipsoids using stereolitography and packing them inside spherical containers [79]. The analysis of the packings required better understanding of finite-size effects, and we use both simulations and novel experimental methods to minimize surface effects and to obtain good estimates of bulk packing densities. We show that, in a sphere, the radial packing fraction $\phi(r)$ can be obtained from $V(h)$, the volume of added fluid to fill the sphere to height h. We also obtain $\phi(r)$ from a magnetic resonance imaging (MRI) scan. The measurements of the overall density $\phi_{\rm avg}$, the radial density distribution $\phi(r)$, and the core density $\varphi_o \approx 0.740 \pm 0.005$ agree with simulations. This verifies that idealized computer packings of frictionless hard particles are relevant to the quantitative understanding of real-life packings of frictional macroscopic particles in a gravitational field.

7.1 Introduction

Hard-particle packings have been intensely studied theoretically, computationally, and also experimentally [195, 79, 40, 16, 196, 70, 3, 154, 197, 155, 198, 199, 200, 201]. Packings of frictionless hard spheres or ellipsoids cannot be realized experimentally because hard frictionless particles are an idealization of realistic particles. However, colloids and granular materials can be synthesized so that the constituent particles are very good approximations of hard spheres, either by tuning the interaction potential between the colloidal particles, or by using manufactured particles such as ball bearings, marbles, beads, etc. Obtaining a high fidelity for the shape of the particle and small polydispersity between different particles is not trivial, and interaction forces in addition to the hard repulsive core, such as friction or hydrodynamic forces, cannot be completely eliminated. Boundary conditions, such as the effects of walls and gravity, also cannot be eliminated [70]. Finally, studying the internal structure of the packings, such as the positions of the particles and the interparticle contacts and forces, is challenging.

Packings can be studied using both high-school and high-tech approaches. Determining the packing fraction of a packing in a hard-wall container is as trivial as measuring the volume of liquid that can be poured into the container. The structure of packings can also be studied by using sophisticated X-ray micro-tomography [198, 199] or confocal microscopy [16]. In this chapter we apply a variety of methods to study the packing of hard ellipsoids, with the ultimate goal of verifying the realizability and practical relevance of the computer-generated disordered packings we studied in Chapter 5. Experiments have been

done with hard spheres for more than half a century, and it is well-known that experimental packings closely resemble the MRJ packings obtained through a variety of simulation methods, if the packing is vibrated or tapped sufficiently to overcome friction and to jam the particles into mechanically stable configurations [195]. Ellipsoids are however hard to manufacture, and to our knowledge ours are the first thorough experimental studies of disordered packings of hard ellipsoids.

In Section 7.2 we perform experiments on large packings of M&M'S Candies®, and compare with the results obtained in simulations using periodic boundary conditions. In Section 7.3 we test whether a random packing of ellipsoids as dense as $\phi \approx 0.74$ could be realized experimentally by performing experiments on relatively small packings of fabricated ellipsoids in spherical flasks, and compare the results against direct computer simulations of packings inside a spherical hard-wall container.

7.2 Bulk Packing of M&M'S Candies®

In this first set of experiments, we study large packings of M&M'S Candies®, a widely-available commercial product that is an excellent substrate for studying packings of hard ellipsoids, with the goal of obtaining a good estimate of the packing density of very large packings, i.e., the *bulk density* ϕ_{bulk}.

The experiments used two varieties of M&M'S Candies®, regular and baking or "mini" candies[1]. Both are oblate spheroids, with small deviations from true ellipsoids, $\Delta r/r < 0.01$, as illustrated in Figure 7.1. Additionally, M&M'S Candies® have a very low degree of polydispersity with principle axes $2O_1 = 1.34 \pm 0.02$ cm, $2O_2 = 0.693 \pm 0.018$ cm, $O_1/O_2 = 1.93 \pm 0.05$ for regular, and $2O_1 = 0.925 \pm 0.011$ cm, $2O_2 = 0.493 \pm 0.018$ cm, $O_1/O_2 = 1.88 \pm 0.06$ for minis. A set of experiments was performed by filling 0.5 , 1 and 5 liter round flasks (to minimize ordering due to wall effects) with candies by pouring them into the flasks while tapping (5 liters corresponds to about 23,000 minis or 7500 regulars). Figure 7.1 shows an illustration. The volume fractions found were $\varphi = 0.685 \pm 0.01$ for both the minis and regulars[2]. The same procedure for 30,000 ball bearings in the 0.5 liter flask gave $\varphi = 0.635 \pm 0.01$, which is close to the accepted MRJ density. A five liter sample of regular candies similar to that shown in Fig. 7.1 was scanned in a medical MRI at Princeton Hospital. For several planar slices, the direction θ (with respect to an arbitrary axis) of the major elliptical axis was manually measured and the two-dimensional nematic order parameter, $S_2 = \langle 2\cos^2\theta - 1 \rangle$ computed with the result $S_2 \approx 0.05$. This is consistent with the absence of orientational order in the packing [43]. Compare these results to those in Chapter 5 for the computer-generated packings of $N = 1000$ oblate ellipsoids with aspect ratio $\alpha = a/b = 1.9$, as illustrated in Fig. 5.1, with density of about $\varphi \approx 0.70$, and nematic order parameter $S \approx 0.02 - 0.05$.

Experimentally, the average contact number $\bar{Z}$ for spheres was determined by Bernal and Mason by coating a system of ball bearings with paint, draining the paint, letting it dry, and counting the number of paint spots per particle when the system was disassembled [201]. Their results gave $\bar{Z} \approx 6.4$, surprisingly close to isostaticity for frictionless spheres. We performed the same experiments with the M&M'S Candies®, counting the number of true contacts between the particles. Near neighbors (even when very close) leave a spot, touching neighbors leave a spot with a hole in the middle at the contact point. A histogram of the number of touching neighbors per particle for the regular candies is shown in Fig. 7.2. The average number is $\bar{Z} = 9.82$. In simulations a contact is typically defined by a cutoff on the gap between the particles. Fortunately, over a wide range ($10^{-9} - 10^{-4}$) of contact tolerances, $\bar{Z}$ is reasonably constant. Superposed in Fig. 7.2 is the histogram of contact numbers obtained for simulated packings of oblate ellipsoids for $\alpha = 1.9$, from which we found $\bar{Z} \approx 9.80$. The minimal number of contacts per particle observed in the simulation is 6 (for only three particles), and 7 in the experiment (thus, there are no rattlers).

7.3 Finite Packings

The simulations presented in Chapter 5 predicted that ellipsoids that do not have an axis of rotational symmetry and have skewness of $\beta = 0.5$ (self-dual ellipsoids) produce the densest random packings, and in particular, it was predicted that ellipsoids with axes ratios of $1.25 : 1 : 0.8$ ($\alpha \approx 1.6$) pack densest, reaching

[1] M&M'S Candies® are a registered trademark of Mars, Inc.

[2] We estimate the correction due to the lower density at the surface of the flasks to be about 0.005, and a more thorough discussion of this issue will be given shortly.

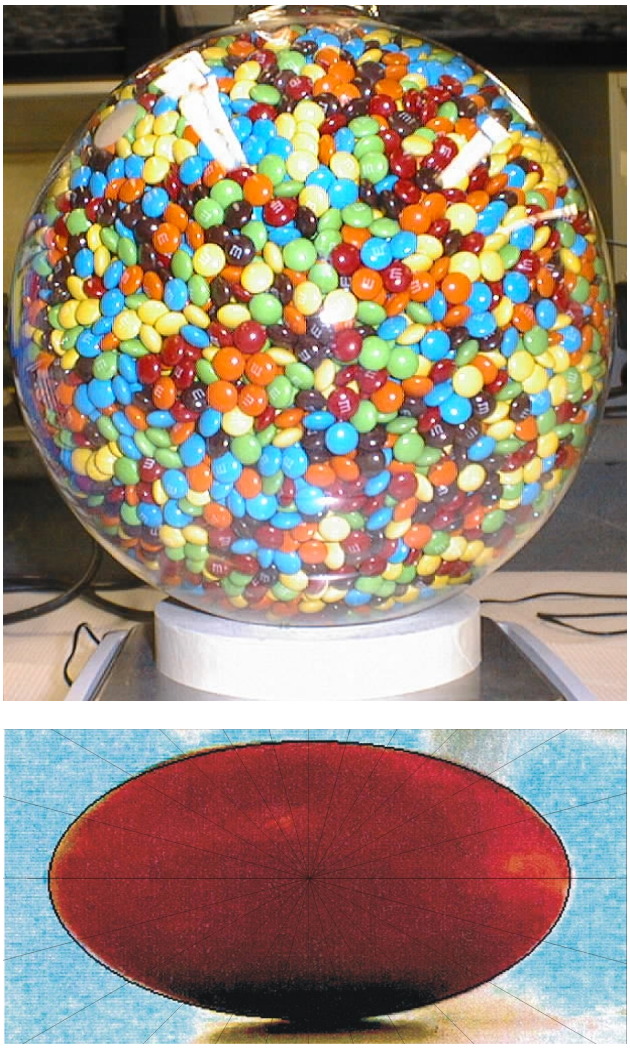

Figure 7.1: (*Top*) An experimental packing of about 7500 regular M&M'S Candies®. (*Bottom*) A cross section of one of the candies showing its high degree of fidelity to a true ellipse.

a random packing density as high as the highest packing density achievable with spheres, $\phi \approx 0.74$. We will call this special ellipsoid shape *ollipsoid*, short for "optimal ellipsoid". In this section we test whether such a dense random packing of ollipsoids could be realized experimentally.

For this purpose, we fabricated ollipsoids using a stereolitography machine 3D Systems Model SLA 250. The machine uses a UV laser with an absorption length of $\sim 100\mu$m in monomer CIBATOLL® SL 5170. The computer controlled laser scans the surface leaving a solid polymerized layer 100μm deep. This patterned layer is then lowered and the next layer written by polymerizing the covering liquid. The beam width, which determines the xy resolution, is comparably scaled at 150μm allowing feature size of 150μm in each direction. The size of our ellipsoids was $2.344 : 1.875 : 1.5$cm for $<1\%$ resolution. A picture of an experimental packing of these ollipsoids inside a spherical flask is shown in Fig. 7.4. The fabrication process is involved and costly, and we could only produce about a thousand ollipsoids. Additionally, due to the resolution limitations of stereolitography, the particles are rather large and therefore producing very large packings is difficult.

It is therefore important to look at ways of extracting information about bulk packings from very small packings (a few hundred particles), in finite hard-wall containers. In particular, in order to minimize the effects of the hard walls, it is best to choose a spherical packing container, since a sphere minimizes the surface area for a given volume, and additionally, effects due to sharp corners are avoided. In this section we

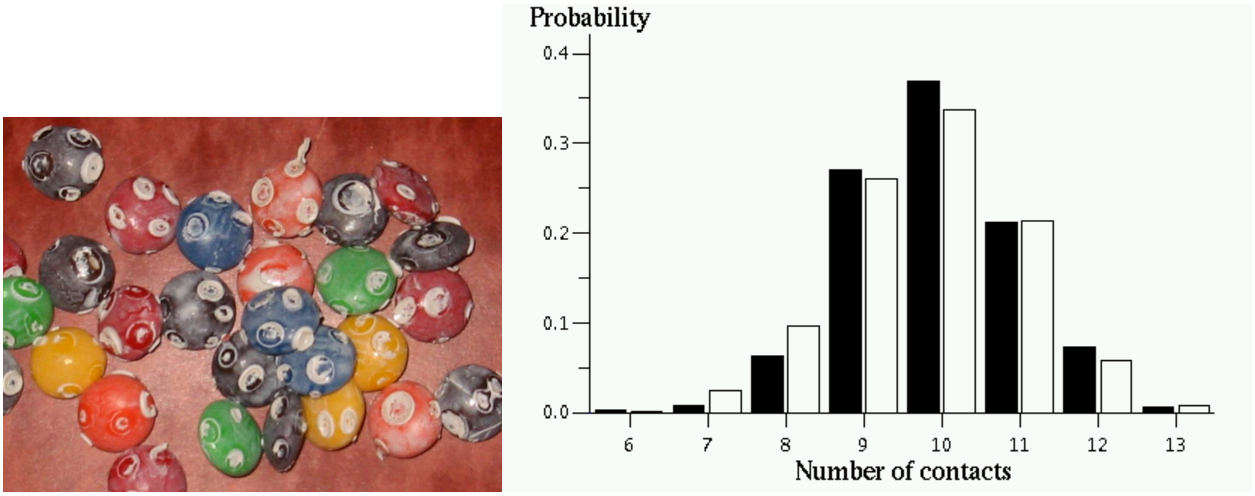

Figure 7.2: (*Left*) The paint marks left on the candies after draining the paint. (*Right)* Comparison of experimental (black bars, from 489 M&M'S Candies®) and simulated (white bars, from 1000 particles) distribution of particle contact numbers.

study finite size and boundary effects for packings in a spherical flask via simulation, theory, and experiments, and demonstrate that it is indeed possible to design experimental methods that largely eliminate the effects of the boundary.

7.3.1 Simulations

In this section we use the modified LS algorithm described in Chapter 3 to study the effects of finite system size and boundary conditions on the density and internal structure of packings of hard ellipsoids.

7.3.1.1 Finite-Size Effects

Systems studied in simulations or experiments are always finite, and it is important to access the *finite-size* and *boundary effects* on the measured packing densities. At the most basic level, even in an infinite boundary-free packing, the density that would be measured inside a finite window has a certain statistical variance from the mean bulk density. Quantifying the variance of the density, i.e., the density fluctuations, inside a packing is a non-trivial task. In Chapter 10 we will study number density fluctuations in large disordered jammed hard sphere packings, and we will see that for jammed packings the variance of the number of particles found inside a spherical window of radius R, as the center of the window is moved through an infinite packing, grows like the surface-area of the sphere R^2, rather than the volume of the sphere (this is the so-called *hyperuniformity* property [202]). This suggests that the standard deviation of the packing fraction $\phi(R)$ inside a randomly positioned spherical window of radius R (i.e., the fraction of the sphere covered by particles), decays like R^{-2}.

Indeed, this is observed computationally by moving a spherical window inside a large packing of hard spheres and monitoring the root mean square deviation $\sigma_\phi(R)$ of the calculated density, as shown in Fig. 7.3. The numerical results are well fitted by

$$\sigma_\phi(R) \sim 0.1(2R/D)^{-2} \approx 0.1(\phi/N)^{2/3} \approx 0.07N^{-2/3},$$

where we used the fact that $\phi \approx N(D/2R)^3$. This sets a bound on the achievable accuracy in measuring bulk packing densities using finite samples, even if boundary effects have been eliminated, for example, if one is far from the walls of the packing container. As an example, if one wants to estimate the density to within $\sigma_\phi \approx 0.001$, one needs to have a spherical sample of at least $N \approx 0.02\sigma_\phi^{-3/2} \approx 650$ particles. Of course, averaging over N_s independent samples will improve the accuracy by a factor of $1/\sqrt{N_s}$. However, it

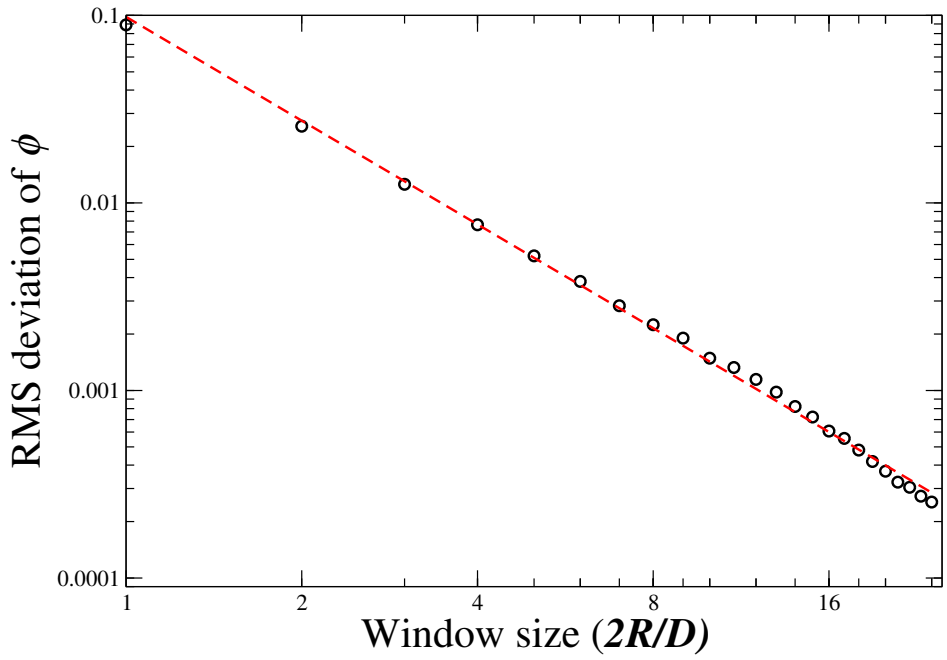

Figure 7.3: A log-log plot of the standard deviation (root mean square deviation) of the packing fraction $\phi(R)$ inside a spherical container of radius R randomly placed inside a jammed disordered packing of $N = 10000$ hard spheres. Also shown is a power-law fit $\sigma_\phi(R) \sim 0.1(2R/D)^{-2}$.

should be noted that it is difficult to exactly reproduce packing protocols such as shaking/tapping intensity between different packings.

7.3.1.2 Boundary Effects: Spherical Containers

We have also implemented a hard-wall spherical boundary in the molecular dynamics algorithm described in Chapter 3, for the purpose of direct comparison between computer-generated packings with experimental packings of M&M'S Candies® and/or manufactured ellipsoids in spherical containers. To save time and implementation effort, we used lattice boundaries (but without periodic images) and employed a trick to implement the spherical boundary. Specifically, we put a spherical container inside a cube and then added special code in the handling of the boundary conditions to predict and process collisions with the hard walls. This was not difficult to do because the spherical container is a special case of an ellipsoid and we already have well-developed tools to deal with collisions between an ellipsoid contained within a larger one, as presented in Section 2.2.2. In fact, implementing true flat hard walls is more difficult for ellipsoids as it necessitates the development of new overlap potentials (see Section 2.2.5.1). Figure 7.4 shows one of our ollipsoid packings. As we will demonstrate shortly, the properties of this computer-generated packing compare very well to actual experimental data.

In Fig. 7.5 we show how the packing density varies with the radial distance for a jammed packing of $N = 5000$ hard spheres inside a spherical hard-wall container or radius R. More specifically, we show the radial density profile $\phi(r)$ as determined by calculating the packing fraction inside a thin spherical shell[3] between spheres of radii r and $r + \Delta r$. The radial distance is normalized with the size of the particles, specifically, with the geometric mean of the ellipsoid semiaxis $a = \sqrt[3]{O_1O_2O_3}$, which for spheres is simply the radius, $a = O = D/2$. The figure clearly illustrates the layering of the particles near the hard walls. As

[3]Note that for ellipsoids what we actually calculate is a numerical approximation to the true radial density profile since it is rather nontrivial to exactly evaluate the volume of intersection of an ellipsoid and a sphere. Out calculation of $\phi(r)$ for ellipsoids is based on a local approximation of the sphere with its tangent plane, and is rather accurate for sufficiently large r (as compared to the size of a particle).

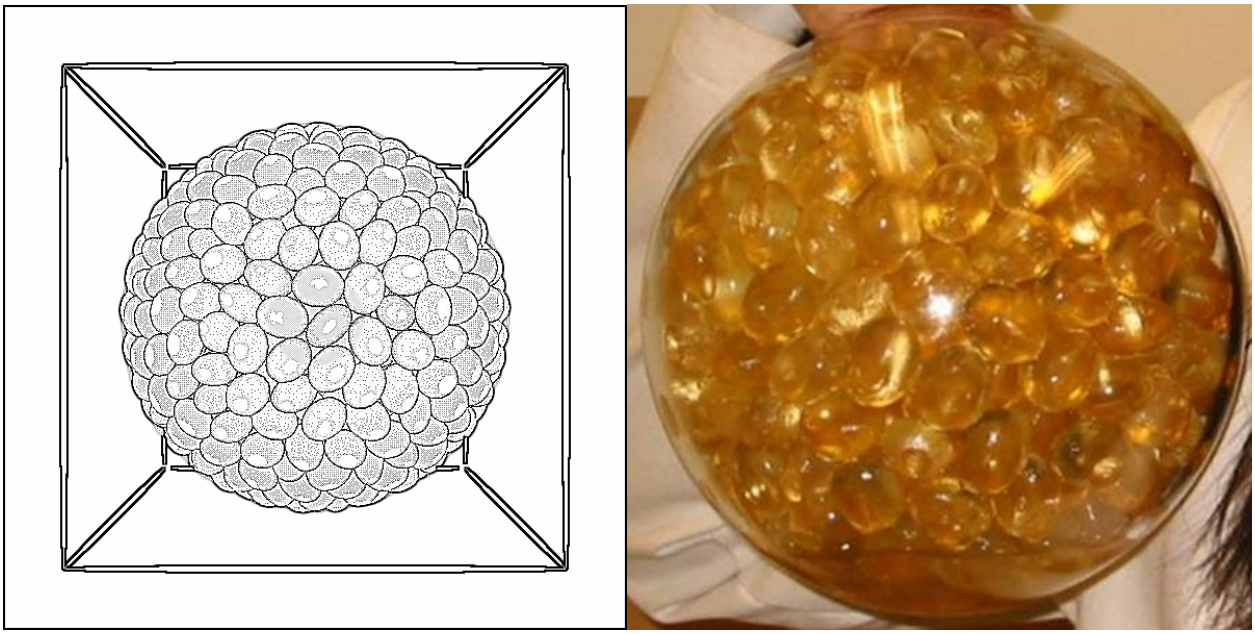

Figure 7.4: (*Left*) A packing of $N = 1000$ ellipsoids with aspect ratios $0.8 : 1 : 1.25$ (ollipsoids) inside a spherical container. A cube enclosing the sphere is used as a pseudo-boundary for the purposes of the MD algorithm. (*Right*) An experimental packing of ellipsoids manufactured using stereolitography, to be compared with the computational packing on the left.

expected, the density is zero next to the wall $r = R$, and has large fluctuations due to layering near the hard wall. The fluctuations are much larger near the walls and slowly decay toward the core. Due to the lower density and ordering of the particles next to the hard wall, one needs to eliminate several layers of particles close to the wall before $\phi(r)$ reaches the core density of about $\phi_0(r) = 0.64$, however, it should be noted that there are oscillations of the density even far away from the walls due to the finite size of the sample.

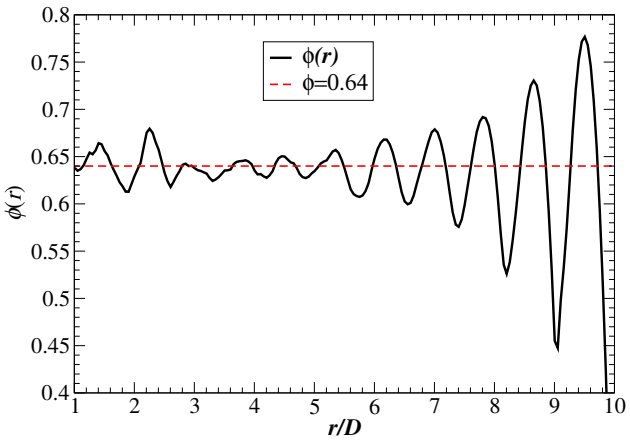

Figure 7.5: The (approximate) fraction $\phi(r)$ of a sphere of radius r covered by the particles for a packing of $N = 5000$ hard spheres inside a spherical container of radius R. The dashed line shows the expected (bulk) packing density far from the hard walls, $\phi_o \approx 0.64$.

We note a few observations about the nature of the ordering (layering) near hard walls. The layering is very reproducible between samples, likely due to the high degree of ordering (crystallization) near the walls.

The density oscillations are largest for spheres, as can be expected due to the sharp localization of length scale, and are smaller for ellipsoids with three different semiaxes (and thus a broad range of length scales). Larger aspect ratios lead to smaller oscillations, again, likely due to the spread in length scales between the different semiaxes. The period of oscillations, or more generally, the nature of the layering near the hard wall, seems nontrivial to predict, and in fact we have not attempted a full study of ordering near flat surfaces in random jammed packings of ellipsoids. Our current implementation of the MD packing algorithm only implements hard-wall boundary conditions with spherical containers for ellipsoids, and implements both spherical and flat hard walls for spheres.

7.3.2 Experiments

In this section we try to design experiments that try to eliminate the effects of the hard wall boundaries and obtain a good estimate of the bulk packing properties of hard ellipsoids. We develop methods in order of increasing complexity and technical sophistication.

7.3.2.1 Extrapolation to $R \to \infty$

The simplest way to measure packing fraction is to pack a container with the particles and determine the volume of fluid which fills the voids. The average (overall) packing density ϕ_{avg} of a finite particle packing in a spherical container of radius R is significantly smaller than the bulk density of equivalent particles ϕ_{bulk}, even for large containers. This is primarily due to the lower density $\phi(r)$ in the first layer of particles near the wall [next to the wall the density is identically zero $\phi(R) = 0$]. To get an estimate of what the correction to the measured packing density is for a large spherical container, we can assume that the packing is radially symmetric and that the radial density profile has the form

$$\phi(r) = \phi_0 - \Delta\phi\left(\frac{R-r}{a}\right), \tag{7.1}$$

where a is a length-scale for the particles, and $\Delta\phi(x)$ is an approximately universal (i.e., independent of R) rapidly decaying (on the scale of a few a's) function representing the correction to the density due to layering near the wall. For large containers, the core density ϕ_0 will approach the bulk packing density, $\phi_0 = \phi_{\text{bulk}}$. If we use (7.1) in

$$\phi_{\text{avg}} = \frac{3\int_0^R r^2\phi(r)dr}{R^3}$$

we obtain the expansion of ϕ_{avg} in powers of $\delta = a/R$,

$$\phi_{\text{avg}}(\delta) = \phi_0 - 3\left\langle\Delta\phi\right\rangle_1\delta + 6\left\langle\Delta\phi\right\rangle_2\delta^2 - 3\left\langle\Delta\phi\right\rangle_3\delta^3, \tag{7.2}$$

where $\left\langle\Delta\phi\right\rangle_n = \int_0^\infty x^n\Delta\phi(x)dx$ are moments of $\Delta\phi$. This argument demonstrates that it is very reasonable to expect the leading-order correction to the density to be given with a term proportional to a/R, so that a linear fit to $\phi_{\text{avg}}(\delta)$ for a range of (sufficiently large) container sizes and extrapolation to $a/R \to 0$ should give a very good estimate of ϕ_0.

In Fig. 7.6, we present the results of experiments with spherical glass marbles of diameter 1.1cm, mini M&M'S Candies® (oblate spheroids), and our fabricated ollipsoids, in spherical flasks of several sizes. The ellipsoids or spheres were poured into glass flasks which were then shaken by hand. More particles were added until further shaking failed to produce space for the addition of another particle. In order to overcome friction, we added a small amount of isopropanol as a lubricant for the marbles and ellipsoids. Also shown are computer generated jammed configurations of frictionless spheres and ellipsoids in geometrically similar spherical containers [11]. From this data, it is clear that the density is the largest for the fabricated ellipsoids and smallest for the spheres. Extrapolating the simulation results to $\delta \to 0$, we find $\phi_0(\text{spheres}) = 0.642 \pm 0.002$, $\phi_0(\text{MMs}) = 0.700 \pm 0.002$ and $\phi_0(\text{ollipsoids}) = 0.741 \pm 0.002$. Our experimental data for glass marbles and M&M'S Candies® agree well with simulated results of frictionless particles. However, for the fabricated ellipsoids, the measured packing densities are noticeably lower than the simulated result for $a/R > 0.1$, most likely due to the influence of friction near the walls of the container (the ollipsoids have a rather rough surface). A meaningful extrapolation to an infinite container for the ollipsoids is difficult for the number of particles we have studied; we need to further minimize surface effects.

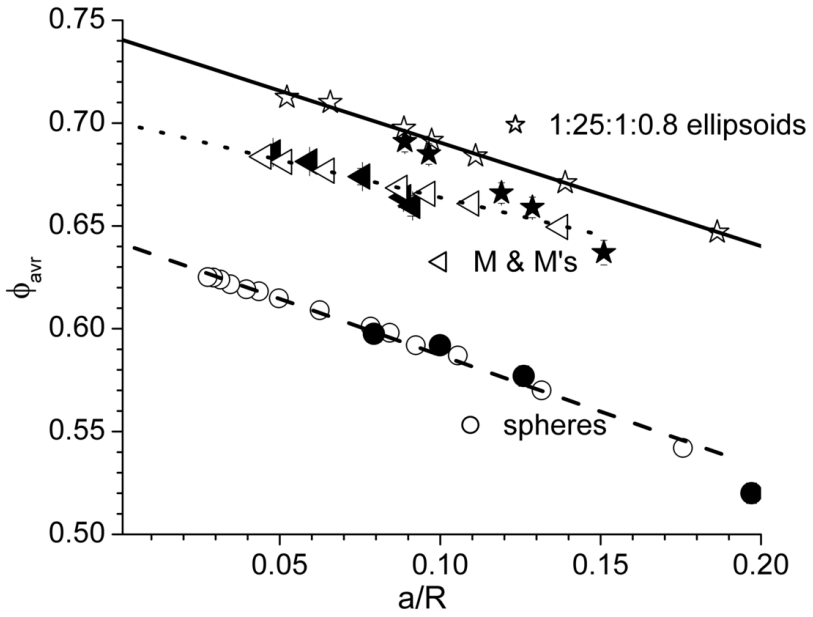

Figure 7.6: Packing fraction $\phi_{\rm avg}(\delta)$ as a function of inverse spherical container size $\delta = a/R$. Empty symbols and lines are data and extrapolation for simulated packing. Solid symbols are the corresponding experimental results.

7.3.2.2 Extracting the Average Core Density

The average density of a spherical core, excluding several particle layers near the container wall, should provide a very good approximation to the bulk density. We therefore devised a different experiment to allow a measurement of the radial density profile $\phi(r)$. A spherical container of radius 10cm was prepared with two holes in the bottom. Tubes were attached to the holes with epoxy. The sphere was filled with particles and shaken until no further particles could be added. The sphere was then set on a digital scale and one tube was attached to a water reservoir. The other tube was attached to a pressure transducer with a sensitivity of 10^{-5}bar (equivalent to a water height of about 0.1mm). The experiment proceeded by allowing the water to slowly enter the container from the bottom while recording the weight [i.e., volume $V(h)$] of fluid added and the pressure (i.e. fluid level h). This is illustrated in the inset in Fig. 7.7.

As liquid is poured into the container, the total volume of liquid V as a function of the height of the liquid column h, measured from the center of the flask, $-R \le h \le R$, is given by the integral

$$V(h) = \int_R^h 2\pi r^2 (1 - \frac{h}{r}) \, [1 - \phi(r)] \, dr. \tag{7.3}$$

Eq. (7.3) is an integral equation for the radial density profile $\phi(r)$, and its solution is trivially obtained by differentiating $V(h)$ twice:

$$1 - \phi(r) = -\frac{1}{2\pi r} \left. \frac{\partial^2 V(h)}{\partial h^2} \right|_{h=r}. \tag{7.4}$$

Using Eq. (7.4) to obtain $\phi(r)$ from an experimental $V(h)$ is difficult. Not only is second-order differentiation a numerically unstable procedure, especially with noisy data, but also the natural fluctuations in $V(h)$ due to the discrete character of the actual density mask the global fluctuations.

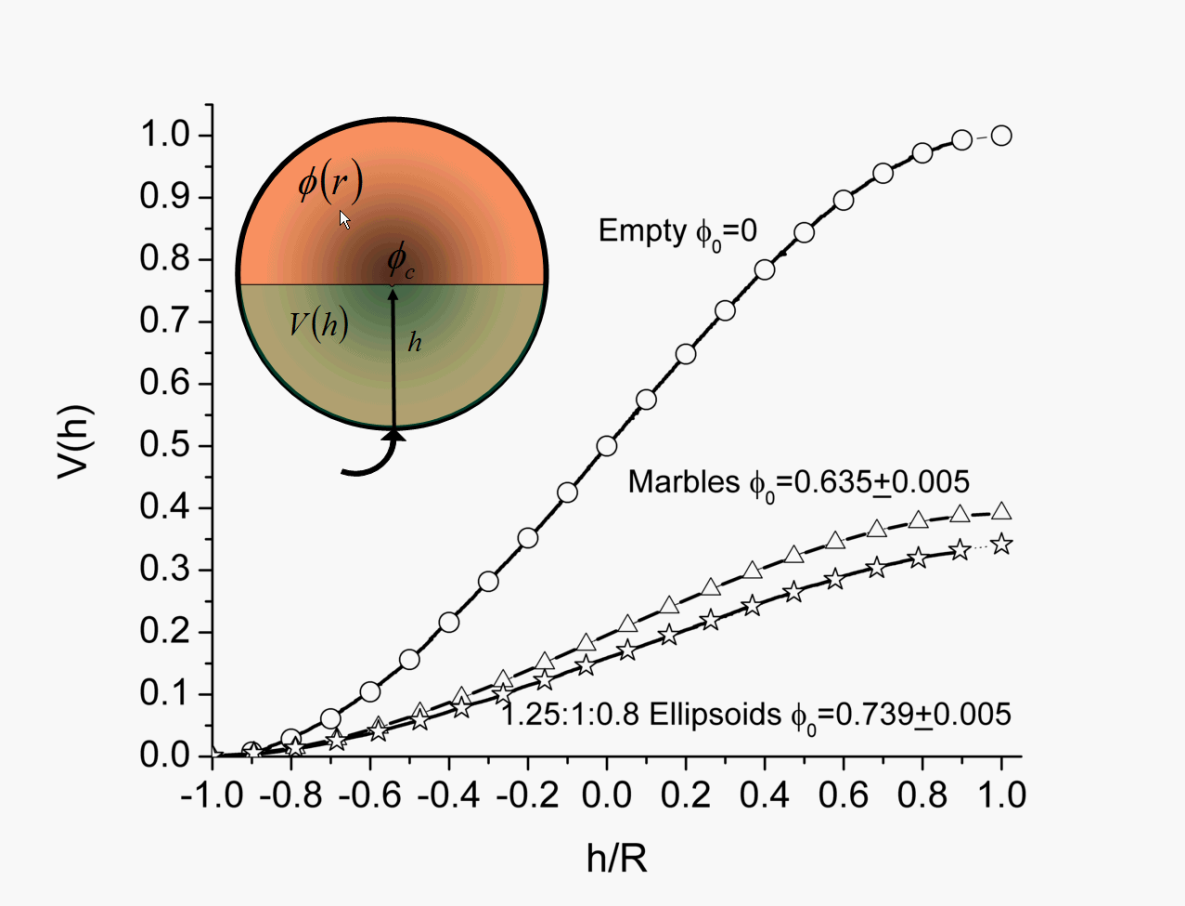

Figure 7.7: Volume $V(h)$ of liquid to fill a spherical container to height h from center. Circles are for an empty container, triangles for a marble packing, and star for an ollipsoid packing. Lines are representative of 35,000 data points. Symbol for empty sphere is from Eq. (7.5) with $\phi_0 = 0$, while the symbols for marbles and ollipsoids are cubic fits.

It is instructive to consider the case when $\phi(r) = \phi_0$ is constant, to obtain a $V(h)$ which is cubic in h,

$$V(h) = -(1-\phi_0)\left[\frac{\pi}{3}h^3 - \pi R^2 h - \frac{2\pi R^3}{3}\right]. \tag{7.5}$$

An important observation is that the coefficient of h^3 does not depend on R and is thus independent of system size. In fact, this coefficient is rather robust to fluctuations in the density $\phi(r)$ as well. For example, it is easy to show that if $\phi(r) = \phi_0$ is constant only inside a core of radius $\tilde{R} < R$, but is not constant outside the core (which is a good approximation for real packings), then $V(h)$ is again cubic for$-\tilde{R} \le h \le \tilde{R}$, with an unaltered coefficient $-\frac{\pi}{3}(1-\phi_0)$ in front of h^3, but with altered coefficients in front of lower powers of h. Therefore, a quick and rather accurate way of obtaining an estimate for the core (bulk) density in a packing is to fit a cubic polynomial through $V(h)$ and extract ϕ_0 from the coefficient of h^3. In Fig. 7.7, we show the measured $V(h)$ together with the corresponding cubic fits. We find $\phi_0(\text{spheres}) = 0.635 \pm 0.005$, and $\phi_0(\text{ollipsoids}) = 0.735 \pm 0.01$, in excellent agreement with simulation estimates of the bulk packing densities.

7.3.2.3 MRI Scans

We also performed MRI scans of the ellipsoids in the same spherical container to assure that the packing did not show signs of crystal or orientational order. The space between ellipsoids was filled with water. MRI

scans were collected on a Siemens 3T Allegra scanner. The images were acquired at 1mm^3 isotropic resolution per voxel using a conventional 3D FLASH sequence in a standard birdcage head coil. The manufacturer's 3D large field of view geometric distortion correction algorithm was applied to the images. Fig. 7.8 insert shows a gray scale (raw) slice from the MRI data and no apparent ordering is present.

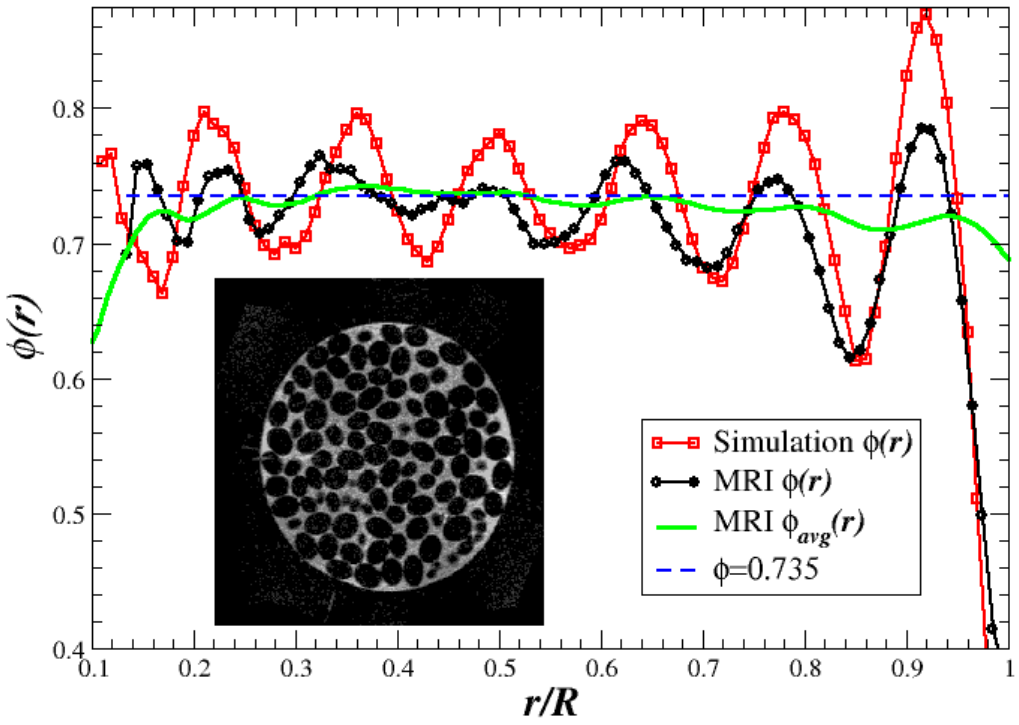

Figure 7.8: The inset shows one of the 256 grayscale slices of an MRI scan of the packing of ollipsoids shown in Fig. 7.4. The radial density profile $\phi(r)$ and its integral, obtained by converting the grayscale images into binary ones, are shown and compared against simulation results. Note that the experimental results have geometrical distortions, especially near the bottom and top of the experimental sample, making it difficult to exactly identify the container boundary.

We can also use the MRI data for an independent measurement of $\phi(r)$. Thresholding based on Sobel derivative boundary detection (we used IDL's image-processing routines) converts the gray scale images into binary ones, and using the binary data, it is straightforward to calculate $\phi(r)$, as shown in Fig. 7.8, along with results from a simulated packing with the same geometry. We also show the average density up to a certain radius r,

$$\phi_{\text{avg}}(r) = \frac{3\int_0^r x^2\phi(x)dx}{r^3},$$

which tells us how well we could estimate the bulk packing density if we could experimentally perfectly remove the particle layers from radius r to the wall container. The average density inside the core $r < 0.6R$ is 0.735, and the error due to finite pixel size, error in threshold and geometry distortions in the MRI scan is in total ± 0.01. While the agreement between the simulation results and the MRI results is surprisingly good, some systematic differences are observed, for example, the amplitude of the density oscillations is smaller

for the MRI data. The largest source of error is the geometry distortion in the MRI scanner, as well as the difficulty of exactly identifying the particle boundaries (which are rough) due to the finite voxel size. These sources of errors will be corrected for in future experiments in order to enable a more accurate description of the microstructure of the experimental packings. Such improved accuracy is necessary, for example, in order to determine particle contacts reliably.

We have also tried to develop procedures that extract the positions and orientations of the ellipsoids from the MRI data. For spheres, this can easily be done using Fourier transform techniques [199]. For ellipsoids this cannot be done and we developed a real space method instead. Using this method we obtained the orientations of the particles and for the sample core we find the nematic order parameter $S_2 \approx 0.13$, which is higher than that expected for a random sample [43], $S_2 = 0.05 \pm 0.03$. This may be due to small amounts of layering present in the packing due to gravity or the experimental packing protocol. The method is not yet accurate enough to identify particle contacts reliably.

7.4 Conclusions

We performed simulations to generate packings of hard ellipsoids inside spherical hard-wall containers. The simulations helped us better understand the impact that finite-size and boundary effects have on the packing density and structure of the packings. We found that it is relatively straightforward to experimentally realize random ellipsoid packings whose packing density matches simulations of hard frictionless particles. We have also demonstrated a new way to measure the bulk density of an infinite random system, using a limited number of particles, which should be useful for packing measurements for other particle shapes. While we will pursue further MRI experiments which can provide more detailed information about inter-particle correlations, the new technique can be used for all types of granular material whether amenable to MRI investigation or not, and is experimentally simple. The development of better experimental methods to study packings is necessary in order to study packing conditions that are not as easily simulated. In particular, friction and gravity are an important factor for the mechanical stability of packings that are not jammed in the frictionless sense we study in this work. We also note that work is underway in the research group of Dr. Paul Chaikin to produce colloids with monodisperse ellipsoidal particles, which would enable a whole series of experiments aimed at studying the thermodynamics of hard ellipsoid systems, something not possible with granular packings.

We find very good agreement between results obtained for computer-generated and experimental (strongly tapped) packings of hard particles. This verifies that idealized computer packings of frictionless hard particles are relevant to the quantitative understanding of real-life packings of frictional macroscopic particles in a gravitational field. It is surprising that such good agreement is observed since the effect of friction is very strong, especially for our fabricated ollipsoids. Also, the experimental packing protocol, in which particles are added and sediment on the bottom of the container as the packing is shaken or tapped, has little resemblance to the LS algorithm, in which all of the particles are in the container and grow while bouncing elastically. This similarity between jammed packings produced by a variety of simulation or experimental protocols points to a fundamental meaning of the MRJ state, and in particular, to maximal disorder. In Chapter 13 we will study binary hard-disk packings and suggest that the MRJ packings have maximal entropy, in the sense that an exponential majority of jammed packings has the same statistical properties (density, pair correlations, etc.) as the MRJ state. If this interpretation of entropy, or degeneracy, as a fundamental order metric is correct, it would provide a deep insight: Almost all jammed packings look alike and therefore it has historically been assumed that "random packing" is a unique and well-defined concept [58]?

Chapter 8

Unusually Dense Crystal Packings of Ellipsoids

Extensive experience with spheres has shown that for reasonably large packings, sufficiently slowing down the growth of the density, so that the hard-particle system remains close to the equilibrium solid branch of the equation of state, leads to packings near the FCC lattice. This however requires impractically long simulation times for large ellipsoid packings. In this chapter, by running the MD packing algorithm for very small deformable unit cells, we are able to identify crystal packings of ellipsoids significantly denser than the FCC lattice [80]. An analysis of the generated packings and an analytical calculation leads us to discover ellipsoid packings with a remarkably high density of $\phi \approx 0.7707$. The family of new packings we discovered are crystal (periodic) arrangements of nearly spherically-shaped ellipsoids, and always surpass the densest lattice packing. The maximum density of $\varphi \approx 0.7707$ is achieved for aspect ratio of $\sqrt{3}$, and in this densest-know ellipsoid packing each ellipsoid has 14 touching neighbors. These results do not exclude the possibility that even denser crystal packings of ellipsoids could be found, and that a corresponding Kepler-like conjecture could be formulated for ellipsoids. We briefly explore the thermodynamical properties of our unusually dense crystal and also propose rigorous ways to look for the densest packing in the neighborhood of the sphere point by using global optimization.

8.1 Introduction

Packing problems, how densely objects can fill a volume, are among the most ancient and persistent problems in mathematics and science [8]. The simulation results for disordered packings of hard ellipsoids [78] suggested that disordered packings can be as dense as the FCC sphere packing, $\varphi \approx 0.735$. It has generally been assumed in the literature that the crystal phase for not very aspherical ellipsoids is a stretched (affinely deformed) FCC sphere packing (i.e., a nematic FCC crystal) or a stacking variant thereof [43], with the same density $\varphi \approx 0.7405$. A system in which the density of a disordered phase surpasses the density of the ordered solid would be a candidate for the elusive thermodynamic glass, since at very high densities the densest phase has the largest free volume and is thus always thermodynamically favored. In light of the possibility that adding orientational degrees of freedom could lead to stable glasses, it is important to determine what the densest packing of ellipsoids in the vicinity of the sphere point really is.

In addition to being important for understanding the physics of complex materials, finding the densest packing for a given particle shape is a basic problem in geometry. The famous Kepler conjecture postulates that the densest packing of spheres in three-dimensional Euclidean space has a packing fraction (density) $\phi = \pi/\sqrt{18} \approx 0.7405$, as realized by stacking variants of the face-centered cubic (FCC) lattice packing. It is only recently that this conjecture has been proved [39]. Very little is known about the most efficient packings of convex congruent particles that do not tile three-dimensional space [9]. The only other known optimal three-dimensional result involves infinitely long circular cylindrical particles: the maximal packing density $\phi_{\max} = \pi/\sqrt{12}$ is attained by arranging the cylinders in parallel in the triangular lattice arrangement [203]. Of particular interest are dense packings of congruent ellipsoids (an affine deformation of a sphere) with

semi-axes a, b and c or, equivalently, with aspect ratio $\alpha = b/a$ and skewness β, where we have assumed that the c semiaxes is the middle semiaxes.

In two dimensions, it can easily be shown that the densest packing of congruent ellipses has the same density as the densest packing of circles, $\phi = \pi/\sqrt{12} \approx 0.9069$ [204]. This maximal density is realized by an affine (linear) transformation of the triangular lattice of circles. Such a transformation leaves the density unchanged. In three dimensions attempts at increasing the packing density yield some interesting structures, at least for needle-like ellipsoids. By inserting very elongated ellipsoids into cylindrical void channels passing through the ellipsoidal analogs of the densest ordered sphere packings (an affinely deformed face centered cubic or hexagonal close packed lattice), congruent ellipsoid packings have been constructed whose density exceeds 0.7405 and approaches 0.7585 in the limit of infinitely thin prolate spheroids (ellipsoids of revolution), i.e., when $\beta = 1$ and $\alpha \to \infty$ [203, 205].

However, there appears to be a widespread belief that for nearly spherical ellipsoids the highest packing fraction is realized by an affine transformation (stretch by α and β along two perpendicular axes) of the densest sphere packing, preserving the density at 0.7405. Mathematicians have often focused on *lattice packings*, where a single particle is replicated periodically on a lattice to obtain a crystal packing. For ellipsoids, a lattice packing is just an affine transformation of a sphere packing, and therefore a theorem due to Gauss [9] enables us to conclude that the densest lattice ellipsoid packing has $\phi \approx 0.7405$. The next level of generality involves nonlattice periodic packings (lattice packings with a multiparticle basis), where a unit cell consisting of several ellipsoids with at least two *inequivalent* orientations is periodically replicated on a lattice to fill Euclidean space. We will refer to these as *crystal packings*.

8.2 The Densest Known Ellipsoid Packing

In this chapter, we report on a family of crystal packings of ellipsoids that are denser than the densest Bravais lattice packing for a wide range of aspect ratios in the vicinity of the sphere point $\alpha = \beta = 1$, and for certain aspect ratios yields the densest known ellipsoid packings with $\phi \approx 0.7707$.

8.2.1 MD Simulations

Extensive experience with the LS algorithm for hard spheres has shown that for reasonably large packings, sufficiently slowing down the growth of the density, so that the hard-particle system remains close to the equilibrium solid branch of the equation of state, leads to packings near the FCC lattice [58, 56]. This is because the dense crystal is thermodynamically favored over disordered liquid and glassy states. This however requires impractically long simulation times for large ellipsoid packings. By running the simulation for very small unit cells, from 4 to 16 particles per unit cell, we were able to identify crystal packings significantly denser than the FCC lattice. The simulations were carried out by starting with a small number of particles inside a cubic cell and then growing the particles very slowly, at expansion rates $\gamma \approx 10^{-8} - 10^{-6}$, while allowing the unit cell to deform as described in Section 3.2.4.2. It is very important to allow for a deforming unit cell when considering small systems. If the lattice is kept constant, then the produced packings will at most be collectively jammed, and their density can strictly be increased to first order by deforming the lattice, as discussed in Chapter 5. If on the other hand the generated packings are strictly jammed, then their density cannot be increased by continuous motions of the particles, even if the unit cell deforms. For small systems, the lowering of the density due to boundary effects is substantial, and without using a deforming lattice we would not have found packings denser than the FCC sphere crystal.

The packings produced by the algorithm are not truly strictly jammed because of the difficulties of incorporating the deforming lattice in the MD algorithm, and also, there is no guarantee that even very slow runs will find the densest packing. By examining several of the densest computer-generated packings and subsequent analytical calculations suggested by the simulation results, we discovered ellipsoid packings with a remarkably high density of $\phi \approx 0.7707$. This result implies that among all possible choices of congruent ellipsoids, the maximum density attainable is bounded below by 0.7707.

8.2.2 Explicit Construction

We now describe the construction of a family of unusually dense crystal packings of ellipsoids. We start from the FCC lattice, viewed as a laminate of face-centered square planar layers of spheres, as illustrated in Fig. 8.1. We similarly construct layers from the ellipsoids by orienting the c semiaxis perpendicular to the layer, while orienting the a and b axes along the axes of the face-centered square lattice defining the layer, as shown in Fig. 8.1. In this process, we maintain the aspect ratio of the squares of side L of the face-centered square lattice defining the layer, i.e., we maintain

$$L = \frac{4\alpha}{\sqrt{1+\alpha^2}}, \tag{8.1}$$

which enables us to rotate the next layer by $\pi/2$ and fit it exactly in the holes formed by the first layer. This two-layer lamination is then continued *ad infinitum* to fill all space. This can be viewed as a family of crystal packings with a unit cell containing two ellipsoids.

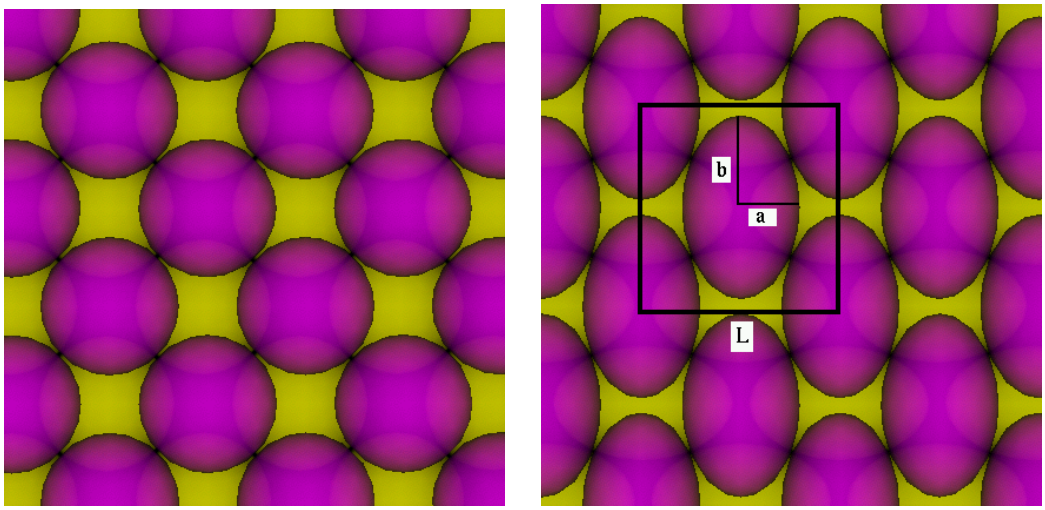

Figure 8.1: *(Left):* The face-centered cubic packing of spheres, viewed as a laminate of face-centered layers (in the (001) plane). The bottom layer is colored purple and the top layer yellow. *(Right)* A nonlattice layered packing of ellipsoids based on the FCC packing of spheres, but with a higher packing fraction.

We can calculate the minimal distance h between two successive layers (that preserves impenetrability) from the condition that each ellipsoid touches four other ellipsoids in each of the layers above and below it. This gives a simple system of equations (two quadratic equations and one quartic equation), the solution of which determines the density to be

$$\phi = \frac{16\pi\alpha}{3hL^2}. \tag{8.2}$$

Notice that the axis perpendicular to the layers can be scaled arbitrarily, without changing the density, and therefore it is sufficient to consider spheroids with $a = c = 1$. The density of this crystal packing as a function of the aspect ratio α is shown in Fig. 8.2, and is higher than the density of the FCC sphere packing for a wide range of aspect ratios around the sphere point $\alpha = 1$. The Taylor series expansion of Eq. (2) about the sphere point through second order in $\alpha - 1$ does not contain a linear term,

$$\phi = \frac{\pi}{\sqrt{18}}[1 + 6(\alpha-1)^2 + \mathcal{O}((\alpha-1)^3)].$$

Two sharp maxima with density of about 0.770732 are observed when the ellipsoids in the face-centered layers touch six rather than four in-plane neighbors, as shown in Fig. 8.3, i.e., when $L = 2\alpha$. This corresponds to an in-plane aspect ratio of $\sqrt{3}$. These two densest-known packings of spheroids are illustrated in the insets

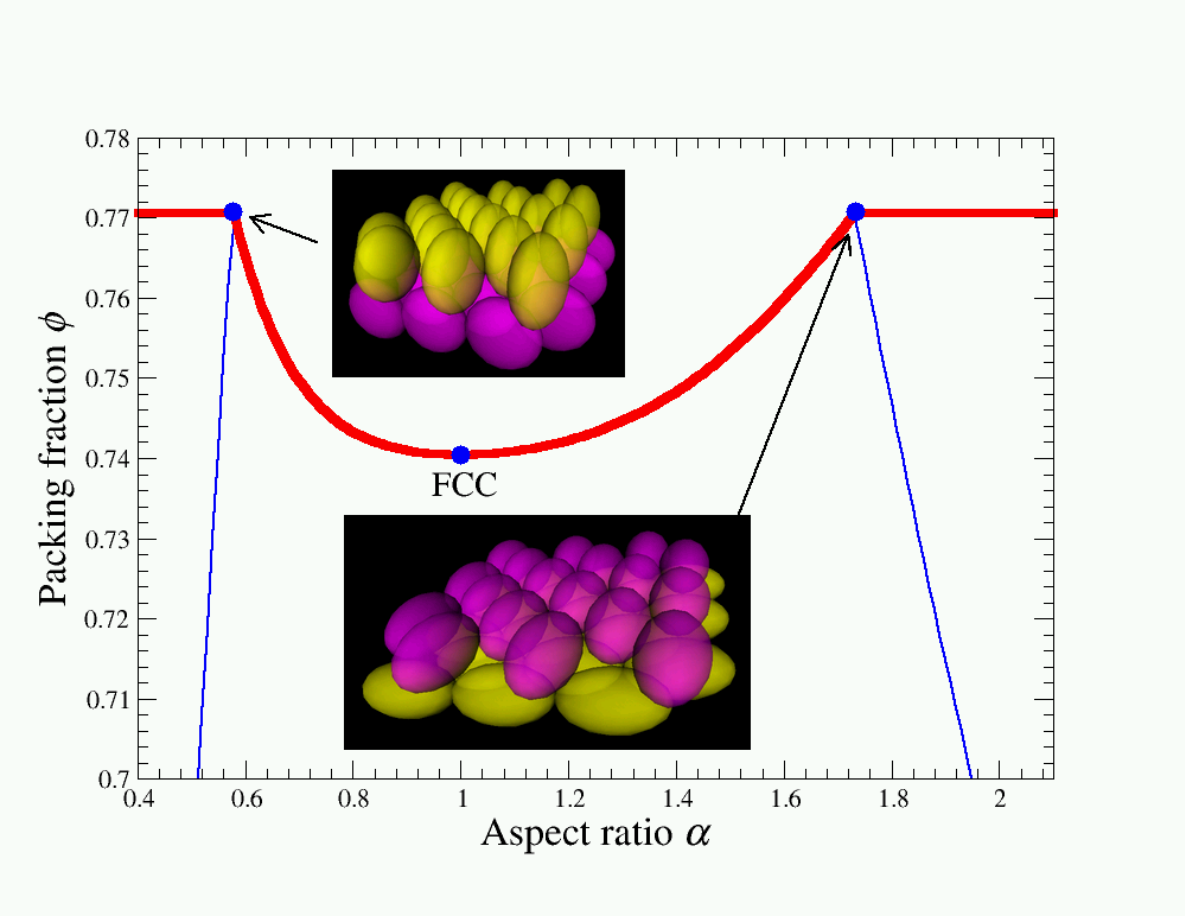

Figure 8.2: The density of the laminate crystal packing of ellipsoids (thin blue line) as a function of the ratio $\tilde{\alpha} = b/a$, which is essentially the aspect ratio but can also be less than unity [i.e., $\alpha = \max(\tilde{\alpha}, \tilde{\alpha}^{-1})$]. The third semiaxes can be fixed at $c = a$ for spheroids. The point $\tilde{\alpha} = 1$ corresponding to the FCC lattice sphere packing is shown, along with the two sharp maxima in the density for $\alpha = \sqrt{3}$, as illustrated in the insets for prolate ($\tilde{\alpha} = \sqrt{3}$) and oblate particles ($\tilde{\alpha}^{-1} = \sqrt{3}$). The presently maximal achievable density is also shown (thick red line), and is constant for $\alpha \geq \sqrt{3}$, as explained in the text. Note that unlike for random packings, the jamming density is a differentiable function of the ellipsoid axes.

in Fig. 8.2, and in these special packings each ellipsoid touches exactly 14 neighboring ellipsoids (compare this to 12 for the FCC lattice). As illustrated for two dimensions in Fig. 5.15, an affine deformation of the densest sphere packing gives an ellipsoid packing that is *not* strictly jammed [15, 72]. It is an interesting open question whether our denser laminated crystals are strictly jammed. Unfortunately, the algorithm to test for jamming described in Chapter 5 cannot be applied to these crystal packings directly because the packing is geometrically very degenerate and has a very high number of self stresses (and thus first order flexes), and therefore different stresses may block different unjamming motions. This is similar to the situation for the rectangular lattice of ellipses discussed in Section 5.5.4.

Figure 8.2 shows a rapid decrease in the packing fraction for large aspect ratios. However, it is a surprising fact that the maximal density of 0.770732 is also achievable whenever the aspect ratio α of the ellipsoids is greater than or equal to $\sqrt{3}$. The key observation is that the $x = y$ plane is a mirror symmetry plane in the above packings, so that an affine stretch by an arbitrary factor $s \geq 1$ along a direction in this plane will produce a packing of equal (stretched) ellipsoids, without changing the density. Stretching an ellipse with

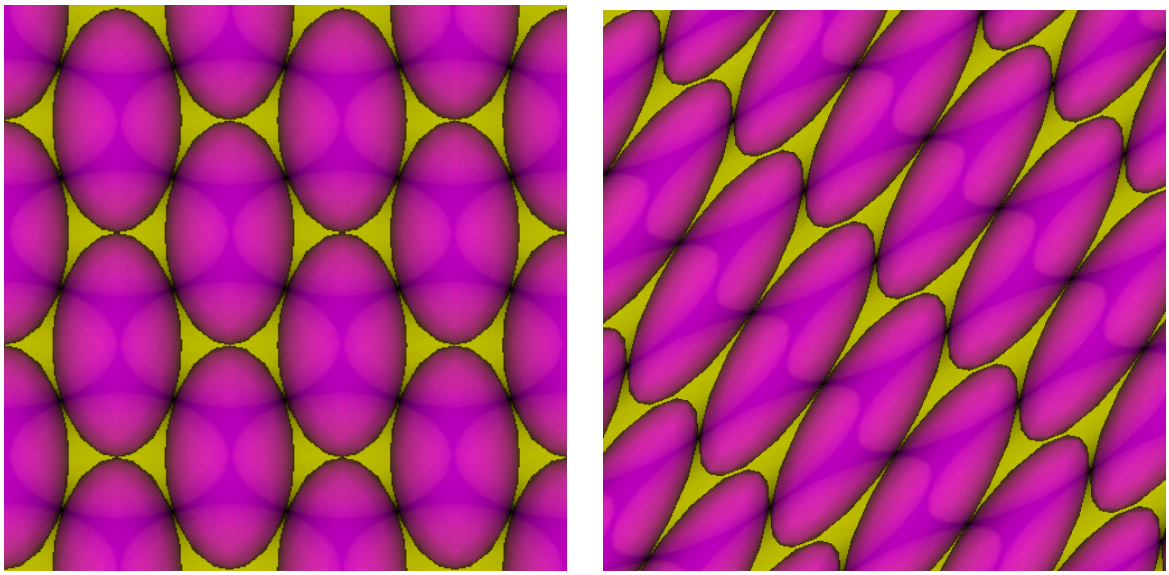

Figure 8.3: (*Left*) The layers of the densest known packing of ellipsoids with aspect ratio $\sqrt{3}$, as illustrated in Fig. 8.2. The same perpendicular view applies for both prolate and oblate particles. The layers can be viewed as either face-centered or triangular. (*Right*) The layers of the densest known packing of ellipsoids with aspect ratio $\alpha = 3$, as obtained by stretching the packing on the left along the $(\sqrt{2}/2, \sqrt{2}/2, 0)$ direction by a factor of 2.4842. The same perpendicular view applies for both prolate and oblate particles or general ellipsoids.

$\alpha = \sqrt{3}$ by a factor of s along the $x = y$ line produces an ellipse with aspect ratio

$$\alpha^2 = \frac{(2 + s^2 + 2s^4) + 2(1 + s^2)\sqrt{1 - s^2 + s^4}}{3s^2}, \tag{8.3}$$

which is always greater than $\sqrt{3}$ and achieves arbitrarily large values for sufficiently large s. Therefore, by stretching the packing in Fig. 8.3 along the $(\sqrt{2}/2, \sqrt{2}/2, 0)$ direction, we can obtain a packing with density 0.770732 for any aspect ratio $\alpha \geq \sqrt{3}$. The layers of such a packing of ellipsoids with $\alpha = 3$ are illustrated in Fig. 8.3. As before, scaling the axis perpendicular to the layers can be used to go between the prolate and oblate cases since the c semi-axes remains aligned with the z axes. Notice that the initial stretch can be along a direction with a nonzero z component, which produces alternative packings with the same density and $\alpha \geq \sqrt{3}$. The above stretch cannot be used to decrease the aspect ratio, so that for $\delta < \sqrt{3}$, our best results remain as shown in Fig. 8.2. In the limit of infinitely large stretch s (i.e., $\alpha \to \infty$), the particles *approach* perfect alignment that are either needle- or plate-like ellipsoids. However, the packings remain non-lattice arrangements with 14 contacts per particle and a density of 0.770732.

8.3 Thermodynamics

We use the collision-driven MD algorithm from Chapter 3 to look more closely at the thermodynamic behavior of prolate and oblate ellipsoids of aspect ratio $\alpha = \sqrt{3}$, for which the crystal packing is the densest known, $\phi \approx 0.77$. The first-order phase transition between the hard-sphere liquid (a more correct term is fluid) and crystal phases can be directly observed in molecular dynamics simulations, and the relevant dynamics (nucleation or relaxation) studied. In MD, one usually studies equilibrium properties by starting with a nonequilibrium system at a given density and then allowing it to equilibrate for a sufficiently long time. An alternative is to very slowly change the density in a quasi-equilibrium manner while tracking the relevant properties such as pressure or order-parameters. This kind of procedure allows one to directly observe the process of melting of the high-density crystal or the freezing of a liquid, and identify approximately

the transition points. The collision-driven MD algorithm we have presented is ideally suited for such a study. Namely, by imposing a very small rate of expansion/contraction γ of the particle extents, one can continuously change the density while keeping the system in quasi-equilibrium.

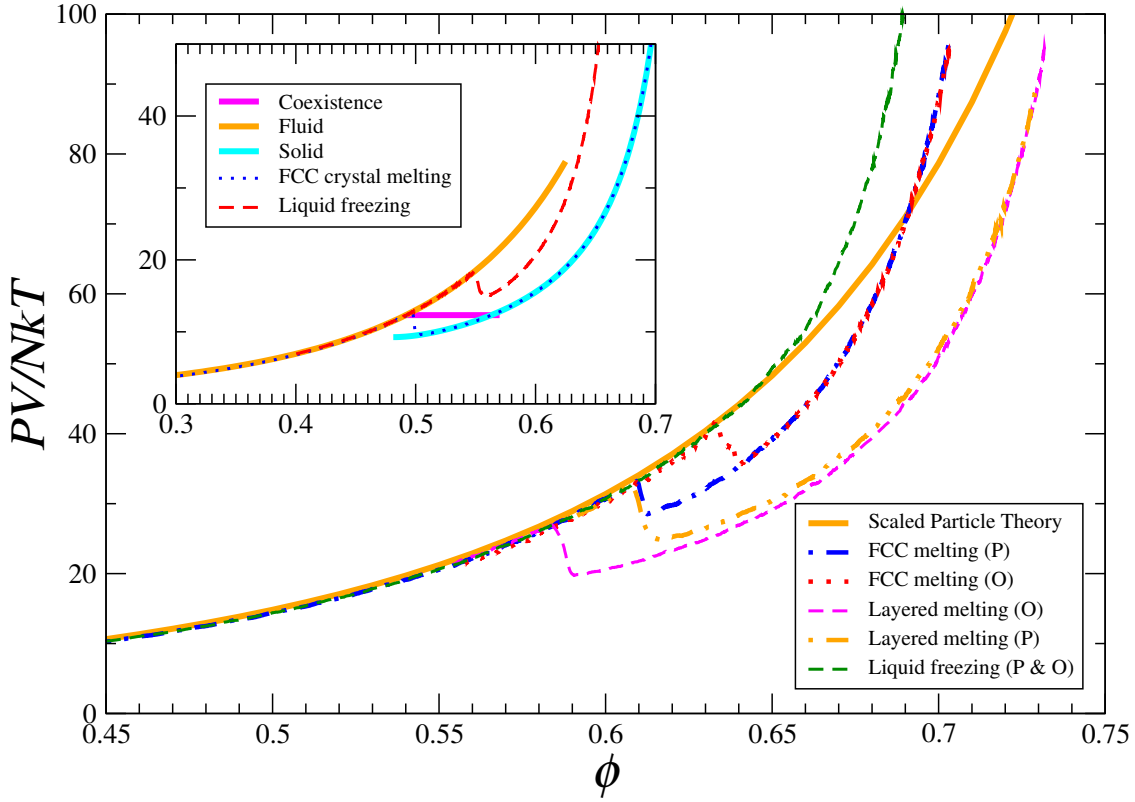

Figure 8.4: Equation-of-state (pressure-density) curves for prolate and oblate ellipsoids of aspect ratios $\sqrt{3}$, compared with spheres (inset), as obtained from quasi-equilibrium collision-driven MD with $N = 1024$ particles. The temperature kT is maintained at unity by frequent velocity rescaling and the "instantaneous" pressure and order parameters are averaged and recorded every 100 collisions per particle. The unit cell is kept fixed. Compare the inset with the more detailed results in Fig. 1.3. Our two-layered crystal [80] is melted from a density of $\phi = 0.73$, and an affinely deformed FCC crystal is melted starting from $\phi = 0.7$ (similarly for spheres in the inset), with $\gamma = -10^{-6}$. An attempt to freeze an isotropic liquid on the other hand fails and leads to a jammed metastable glassy configuration with $\phi \approx 0.72$ despite the slow expansion ($\gamma = 10^{-6}$), for both oblate and prolate ellipsoids. Investigating even slower expansion or larger systems is needed to obtain more quantitatively accurate results.

By starting with a low-density isotropic fluid and very slowly increasing the density, one can produce a superdense liquid and then observe a first-order freezing transition as soon as the metastable fluid becomes unstable, typically when the density approaches the maximal density of coexistence (as also observed in Ref. [53]). This freezing is a nucleation-activated (rare-event) process (the dynamics of which can be observed) and does not lead to perfect crystallization, but is clearly visible as a discontinuous pressure drop, as illustrated for hard-spheres in the inset in Fig. 8.4 and also in Fig. 1.3. One can reverse the process by starting with a perfect crystal, assuming that the stable high-density crystal structure is known, and slowly reduce the density until the crystal melts, typically as the density approaches the minimal density of coexistence, as illustrated in the inset in Fig. 8.4. In the figure, comparison is made to the semi-empirical results for the liquid, solid, and coexistence regions for hard spheres in the literature [53], and for ellipsoids comparison is made to scaled-particle theory for the isotropic fluid [206]. Unfortunately, direct coexistence is very difficult to observe in computer simulations, and requires the creation of an artificial interface between the two phases

[207]. Additionally, it is in principle important to include unit cell dynamics in order to allow for solid-solid transitions. However, it is difficult to do this in a stable manner across a range of densities and in this study we fix the unit cell.

In Fig. 8.4, we also show the melting of our newly-discovered two-layered ellipsoid crystal for an aspect ratio of $\sqrt{3}$ for prolate and oblate spheroids. The crystal melts into an isotropic fluid and no nematic phase is observed, as can be seen by monitoring the nematic order parameters, which rapidly goes to zero as the first-order transitions occur. For comparison, we also show the corresponding melting curves for the ellipsoid crystal obtained by affinely stretching or compressing an FCC sphere crystal along the $(0, 0, 1)$ direction. Additionally, we try to observe the freezing of the isotropic liquid by slow compression. However, it can be seen that despite the very slow compression the liquid does not freeze but rather jams in a metastable glass. This illustrates that systems of ellipsoids have a marked propensity toward (orientationally) disordered configurations and are very hard to crystallize. This is to be contrasted with the case of hard spheres where we easily observe freezing at the same expansion rate. A close agreement between the results for prolate and oblate spheroids is seen, although mode detailed studies with a wider range of system sizes and expansion rates are needed.

8.3.1 Crystal Nucleation and Glassiness

It is interesting to compare the densities of random ellipsoid packings, as reported in Chapter 5, against those of the densest-known crystal packing. In particular, the self-dual ellipsoids with $\beta = 1/2$ have the highest random packing densities, and we compare this density against the density of our crystal (which is independent of β) in Fig. 8.5. We see that for ollipsoids (recall that in Chapter 7 we termed ellipsoids with axes $1.25 : 1 : 0.8$ ollipsoids), the difference in density between the disordered ($\phi \approx 0.735$) and ordered ($\phi \approx 0.753$) packings is rather small, $\Delta\phi \approx 0.02$, especially when compared to the difference for spheres, $\Delta\phi \approx 0.1$. We expect that this will be reflected in a reduced tendency to crystallize, both through an increase in the freezing temperature from the sphere value $\phi_g \approx 0.495$, and also through enhanced tendency to form glasses upon compression. Indeed, recent calculations of the nucleation barrier and nucleation rate for crystallization of hard ellipsoid fluids have found that nucleation is very strongly suppressed [208], although it should be noted that this study did not use our crystal structure.

We do not carry out a more detailed investigation of the thermodynamics of hard ellipsoids in this work. For large aspect ratios it is well-known that ellipsoids form a *nematic* (*liquid-crystal*) phase [43], rather than a glassy phase. In the nematic phase the particle centroids are liquid-like, whereas the orientations are solid-like (aligned). For aspect ratios very close to unity, the liquid freeze into a *rotor* (*plastic crystal*) phase where the particle centroids are just like in a hard-sphere crystal, whereas the particle orientations are random (liquid-like) [208]. For aspect ratios in the range $1.2 < \alpha < 2$ we expect that there will be a first order phase transition from the isotropic liquid phase to a crystal that is like the crystal phase we described in Section 8.2, in which both translations and orientations show long-range periodic order. Our preliminary studies indicate that crystallization into such a crystal is kinetically suppressed and the liquid is thus more likely to form a glass.

It would be of interest to observe spontaneous crystallization of a hard-ellipsoid liquid. This would give hints about the densest packing of ellipsoids, since we know that the densest packing will be favored thermodynamically at sufficiently high densities. However, we have only been able to observe such spontaneous crystallization by starting with a translationally ordered state. We started with an FCC sphere crystal and transformed the spheres into ollipsoids by shrinking then along certain randomly selected axes. We then performed very long MD runs at the reduced density. We observed that the initial plastic crystal (rotor) phase spontaneously froze into an orientationally ordered crystal. Examination of the particle orientations showed that there was alignment along two perpendicular directions in alternating layers, just like in the densest crystal we identified (Section 8.2). This offers some hope that we may have indeed identified the densest packing, at least close to the sphere point. However, it is far from a proof. In the next section we will consider ways to rigorously identify the optimal packing of nearly spherical ellipsoids.

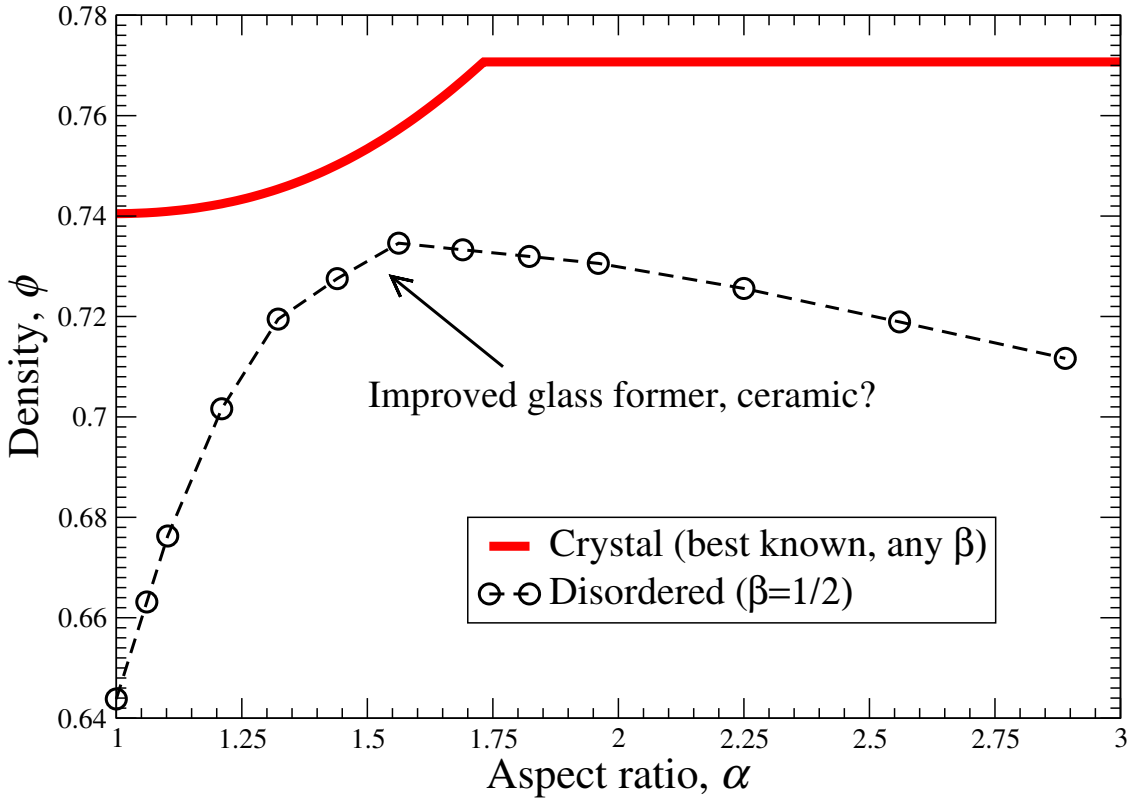

Figure 8.5: A comparison between the densities of random and crystal packings of ollipsoids, as determined in this work. We highlight the special point of ollipsoids, which is expected to give improved glass formers, better ceramic materials (highest densities and contact number), and may even be relevant to the shape of the eggs for certain species of fish.

8.4 Nearly Spherical Ellipsoids

There is nothing to suggest that the crystal packing we have presented here is indeed the densest for any aspect ratio other than the trivial case of spheres. We believe it is important to identify the densest periodic packings of ellipsoids with small numbers of ellipsoids per unit cell. This may be done using modern global optimization techniques, as has been done for various sphere and disk packing problems. However, this is a challenging project due to the complexity of the nonlinear impenetrability constraints between ellipsoids. The case of slightly aspherical ellipsoids is more tractable, since the best packing will be a perturbation of the FCC lattice with a broken orientational symmetry, and should thus be easier to identify. In Fig. 8.2 we see that the density of our crystal packing increases smoothly (quadratically) as asphericity is introduced, unlike for random packings, where a cusp-like increase is observed near $\alpha = 1$ [78]. Is there a crystal packing of ellipsoids which has to a linear increase in density for slightly aspherical ellipsoids?

Formally, the problem is similar to that of finding the densest packing of ellipsoids in the vicinity of a random sphere packing, as we studied in Section 5.9. Consider a (small) unit cell of N spheres, whose lattice is a sublattice of the FCC packing and density is ϕ_J^S. We assign orientations (quaternions) $\mathbf{W}$ to each of the spheres in this crystal packing of spheres. Then we perturb (shrink) the shape of the spheres to be slightly aspherical, with semiaxes $\mathbf{O} = O(\mathbf{I} + \varepsilon \boldsymbol{\epsilon}_O)$, where $s_\phi^\epsilon = \mathbf{e}^T \boldsymbol{\epsilon}_O$ is the relative change in volume of the particles due to the perturbation. Note that we can always take the largest entry in $\boldsymbol{\epsilon}_O$ to be unity and the smallest to be zero, $\boldsymbol{\epsilon}_O = \mathrm{Diag}\{0, \beta, 1\}$, and therefore the aspect ratio after the perturbation is $\alpha = 1 + \varepsilon$. Finally, we increase the density of the packing while keeping the orientations fixed, to find a jamming density

$$\phi_J(\mathbf{W}) = \phi_J^S \left[1 + s_\phi(\mathbf{W})\varepsilon + O(\varepsilon^2)\right].$$

The goal is to find the choice of orientations that maximize $s_\phi(\mathbf{W})$. No expansion can help us in this, since the orientational symmetry is broken when asphericity is added; the search must be over all sets of orientations (quaternions). Going back to the expansions of the impenetrability constraints between nearly spherical ellipsoids in Section 2.2.6, we see that $s_\phi(\mathbf{W})$ can be found from the solution to a linear programming problem,

$$\begin{aligned} s_\phi(\mathbf{W}) &= \max_{\mathbf{\Delta R}, s_\phi} s_\phi \\ \mathbf{A}^T \mathbf{\Delta R} - \frac{2}{d}\left(s_\phi - s_\phi^\epsilon\right)\mathbf{e} &\geq -2\mathbf{U}^T \mathcal{S}(\mathbf{W})\mathbf{U}, \end{aligned} \tag{8.4}$$

where the higher-order matrix (linear operator) $\mathcal{S}$ contains the information from the orientations, see Eq. (2.25). Here the lattice degrees of freedom need to be included in $\mathbf{R}$ and in the rigidity matrix, and we have set $\varepsilon = 1$. Note that the right hand side of the linear constraints in the LP (8.4) are sums of squares of quadratic polynomials of the quaternions $\mathbf{W} = (\mathbf{w}_i)$. The difficult problem is performing the (outer) maximization $\max_{\mathbf{W}} s_\phi(\mathbf{W})$.

We have tried to use some general-purpose global optimization software packages to solve for the optimal $\mathbf{W}$, however, we did not have success in obtaining a provably global optimum for reasonable numbers of particles in the unit cell N. Nevertheless, we give the details of the formulation of the global optimization problem to be solved below, in the hope that future work may be able to use it. If a global optimum of

$$\max_{\mathbf{W}} s_\phi(\mathbf{W}) = 0$$

could be proved, this would show that the density of the FCC hard-sphere crystal cannot be improved to first order as asphericity is added. Instead, one would have to look at second-order terms in the perturbation around the sphere point.

8.4.1 Global Optimization Approach

The problem $\max_{\mathbf{W}} s_\phi(\mathbf{W})$ together with the LP (8.4) is a clearly separated mixed LP-NLP polynomial global optimization problem, where the linear variables are $\mathbf{\Delta R}$, and the nonlinear variables are $\mathbf{W}$.

8.4.1.1 Formulation

We summarize the structure of the global optimization problem that gives the best packing of nearly spherical ellipsoids:

- *Variables*: $\Delta\mathbf{r}_i \in R^3$(displacements) and $\mathbf{w}_i = [s_i, \mathbf{p}_i] \in R^4$ (normalized quaternions), where $\mathbf{p}_i \in R^3$, for $i = 1, .., N$. There are also the degrees of freedom coming from the symmetric strain tensor (lattice deformation) $\hat{\boldsymbol{\epsilon}}$, and of course s_ϕ.
- *Objective*: Linear $f = s_\phi$.
- *Constraints*: For a sparse set of contacting pairs of particles $\{i, j\}$, with given *unit vectors* $\mathbf{u}_{ij}$ as parameters, and a given set of shape parameters $\boldsymbol{\epsilon}_O$, it must hold that:

$$\mathbf{u}_{ij}^T \Delta\mathbf{r}_{ij} + \sum_{k=1}^{3} \epsilon_k^O \left[g_k^2\left(\mathbf{u}_{ij}, \mathbf{w}_i\right) + g_k^2\left(\mathbf{u}_{ij}, \mathbf{w}_j\right)\right] \geq \frac{2}{3}\left(s_\phi - s_\phi^\epsilon\right)\mathbf{e},$$

  which are (nonconvex) polynomial constraints of degree 4, since the g_k's are *nonconvex quadratic* polynomials. Notice that the constraints are separable in the $\mathbf{w}$'s. There are also normalization conditions on the quaternions, $\|\mathbf{w}_i\| = 1$.

The structure of this hard mathematical programming problem is of the following form

$$\begin{aligned} &\min_{\mathbf{y},\mathbf{x}} \mathbf{c}^T\mathbf{x} \\ &\mathbf{A}\mathbf{x} \leq \mathbf{b}(\mathbf{y}) \end{aligned}$$

with $\mathbf{b}(\mathbf{y})$ being at most quartic polynomials of $\mathbf{y}$'s.

8.4.1.2 Hessian of Nonlinear Terms

The separable nonlinear terms in the constraints have the form

$$C = [\mathbf{g}(\mathbf{u}, \mathbf{w})]^T \boldsymbol{\epsilon} [\mathbf{g}(\mathbf{u}, \mathbf{w})],$$

where the elements of the vector $\mathbf{g}(\mathbf{u}, \mathbf{w}) \in R^3$ are quadratic polynomials in $\mathbf{w}$, and $\boldsymbol{\epsilon}$ is a diagonal matrix. The form of the quadratic polynomials is

$$\mathbf{g} = 2\left[(\mathbf{u}^T\mathbf{p})\mathbf{p} + s(\mathbf{u} \times \mathbf{p}) + \left(s^2 - \frac{1}{2}\right)\mathbf{u}\right],$$

where $\times$ denotes a cross product of two vectors. The Jacobian of $\mathbf{g}$ is given with

$$\mathbf{J} = \nabla_{\mathbf{w}}\mathbf{g} = 2\begin{bmatrix} \mathbf{u} \times \mathbf{p} + 2s\mathbf{u} \\ (\mathbf{u}^T\mathbf{p})\mathbf{I} + \mathbf{p}\mathbf{u}^T + s\mathbf{U} \end{bmatrix},$$

which gives the Hessian of the nonlinear term

$$\mathbf{H} = \nabla^2_{\mathbf{w}} C = 2\mathbf{J}^T\boldsymbol{\epsilon}\mathbf{J} + 4\widetilde{\mathbf{H}},$$

where the second term is

$$\widetilde{\mathbf{H}} = \begin{bmatrix} 2\mathbf{v}^T\mathbf{u} & (\mathbf{v} \times \mathbf{u})^T \\ \mathbf{v} \times \mathbf{u} & \mathbf{v}\mathbf{u}^T + \mathbf{u}\mathbf{v}^T \end{bmatrix}$$

with $\mathbf{v} = \boldsymbol{\epsilon}\mathbf{g}$.

Note that if we choose $\boldsymbol{\epsilon} \geq 0$ (which we can since the sign of the shape perturbation is not important), the first term in the Hessian, $\mathbf{J}^T\boldsymbol{\epsilon}\mathbf{J} \succeq 0$, and only the second piece $\tilde{\mathbf{H}}$ has negative eigenvalues (non-convexities). The eigenvalues of $\widetilde{\mathbf{H}}$ are two eigenvalues of multiplicity two,

$$\lambda = \mathbf{v}^T\mathbf{u} \pm \|\mathbf{v}\| = \|\mathbf{v}\|(\cos\theta \pm 1),$$

where θ is the angle between $\mathbf{v}$ and $\mathbf{u}$. The only *potentially* negative eigenvalue (non-convexity) therefore is $\|\mathbf{v}\|(\cos\theta - 1)$:

$$\lambda_{\min} = \|\mathbf{v}\|\left(\hat{\mathbf{v}}^T\mathbf{u} - 1\right),$$

where $\hat{\mathbf{v}}$ is a unit vector in the direction of $\mathbf{v}$.

8.5 Densification by Changing the Ellipsoid Shape

In Section 8.4 we considered the breaking or orientational symmetry in the hard-sphere crystal for nearly-spherical ellipsoids. This symmetry breaking is inherently a discontinuous change, since particle orientations suddenly start making a difference when asphericity is introduced. Once the orientational symmetry is broken however, and the structure of the ellipsoid crystal, and in particular, the particle orientations and the contact network, are determined, we expect that the packing structure and density will evolve smoothly as the asphericity $\varepsilon = \alpha - 1$ is increased. This smooth evolution of the packing structure could be followed, for example, by using an ODE solver to update the particle positions and orientations as the shape of the particle changes continuously with time. Of course, such a smooth evolution is only possible until some singular point is reached, such as the special aspect ratio $1 + \sqrt{3}$ for our crystal packing of ellipsoids, at which point the contact number jumped from 12 to 14.

In this Section we examine the change in the jamming density as one changes the shape of the particles in a nearly jammed packing, to first order. We allow for general changes in the particle shapes and sizes. In particular, for spheres, we examine how the jamming density changes as one changes the diameters of the spheres independently, that is, makes the packing polydisperse. Our goal is to find ways to generate very dense disordered polydisperse packings of disks and spheres. For ellipsoids, we focus on changing the shape of the ellipsoids while keeping the packing monodisperse. The goal is to design a procedure that can find the ellipsoid shape that produces denser packings.

8.5.1 Spheres

Consider a jammed polydisperse sphere packing with particle diameters $\mathbf{D}$ that is at a jamming gap δ (see Section 5.7), and the particle displacements from $\mathbf{R}_J$ are $\Delta\mathbf{R}$. The interparticle gaps are thus

$$\mathbf{h} = \mathbf{A}^T \Delta\mathbf{R} - \frac{\delta}{2}\widehat{\mathbf{A}}\mathbf{D},$$

where $\hat{\mathbf{A}}$ is the node-arc incidence matrix for the *undirected* contact network graph. The equilibrium interparticle forces, $\mathbf{Af} = \mathbf{0}$, are related to the jamming gap δ via [c.f. Eq. (5.19)]

$$\mathbf{f}^T\mathbf{h} = (\mathbf{Af})^T \Delta\mathbf{R} - \frac{\delta}{2}\left(\widehat{\mathbf{A}}\mathbf{D}\right)^T \mathbf{f} = -\frac{\delta}{2}\left(\widehat{\mathbf{A}}\mathbf{f}\right)^T \mathbf{D}. \tag{8.5}$$

Consider changing the particle diameters by $\Delta\mathbf{D}$. How does this affect the jamming density? The interparticle gaps change by $\Delta\mathbf{h} = \frac{1}{2}\widehat{\mathbf{A}}\Delta\mathbf{D}$ and the forces are not affected to first order. Therefore, according to Eq. (8.5), the change in the jamming gap is

$$\Delta\delta = -2\frac{\mathbf{f}^T\Delta\mathbf{h}}{\left(\widehat{\mathbf{A}}\mathbf{f}\right)^T\mathbf{D}} = -\frac{\left(\widehat{\mathbf{A}}\mathbf{f}\right)^T\Delta\mathbf{D}}{\left(\widehat{\mathbf{A}}\mathbf{f}\right)^T\mathbf{D}}.$$

At the same time, the change in the density is

$$\Delta\phi = \frac{S_d}{V}\left(\mathbf{D}^{d-1}\right)^T\Delta\mathbf{D},$$

where S_d is just the constant in front of D^{d-1} in the surface area of a sphere, so that the density is

$$\phi = \frac{S_d}{dV}\left(\mathbf{e}^T\mathbf{D}^d\right).$$

The change in the jamming density $\phi_J = (1 + d\delta)\,\phi$ to first order is

$$\Delta\phi_J \approx \Delta\phi + d\phi\Delta\delta,$$

which is proportional to

$$\Delta\phi_J \sim \left[\mathbf{D}^{d-1} - \frac{\mathbf{e}^T\mathbf{D}^d}{\left(\widehat{\mathbf{A}}\mathbf{f}\right)^T\mathbf{D}}\left(\widehat{\mathbf{A}}\mathbf{f}\right)\right]^T \Delta\mathbf{D}. \tag{8.6}$$

If the $\mathbf{D}-$equilibrium condition

$$\widehat{\mathbf{A}}\mathbf{f} = \frac{\left(\widehat{\mathbf{A}}\mathbf{f}\right)^T\mathbf{D}}{\mathbf{e}^T\mathbf{D}^d}\mathbf{D}^{d-1} = p_{\text{int}}\mathbf{D}^{d-1} \tag{8.7}$$

is satisfied, then no further growth of the jamming density is possible regardless of the specific $\Delta\mathbf{D}$. This condition has a very simple physical interpretation: The term $\widehat{\mathbf{A}}\mathbf{f}$ gives the total force pressing on the surface of the particle (without regard to direction), and therefore the constant p_{int} can be interpreted as the internal pressure inside the particles that resists the compressive forces. When this pressure is equal between all spheres, then the sphere diameters are in mutual equilibrium and changing them only exchanges volume between the spheres but cannot grow the overall volume (density).

As an example, consider perturbing a monodisperse sphere packing into a bidisperse one by randomly selecting a fraction x of the particles to increase their diameter by ΔD, $\Delta\mathbf{D} = \tilde{\mathbf{e}}$, where $\tilde{\mathbf{e}}$ is a vector with xN randomly positioned unit entries. Substitution of this and $\mathbf{D} = D\mathbf{e}$ into equation (8.6) produces

$$\Delta\phi_J \sim xN - \frac{N}{\left(\widehat{\mathbf{A}}\mathbf{f}\right)^T\mathbf{e}}\left[\left(\widehat{\mathbf{A}}\mathbf{f}\right)^T\tilde{\mathbf{e}}\right] = 0.$$

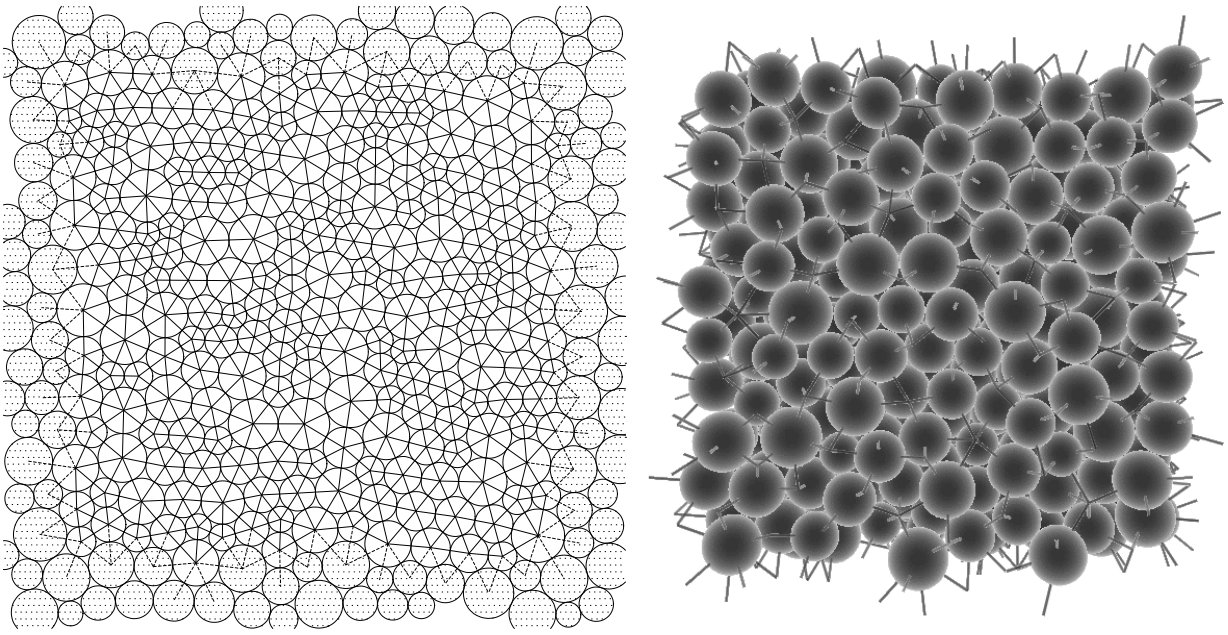

Figure 8.6: (Left) A very dense ($\phi \approx 0.91$) polydisperse packing of disks that is in $\mathbf{D}-$equilibrium, produced by an MD algorithm in which the particle diameters change independently. (Right) A very dense ($\phi \approx 0.68$) polydisperse packing of spheres that is approximately in $\mathbf{D}-$equilibrium.

As discussed in Chapter 13, we expect that in the bidisperse MRJ sphere packing the large and small disks will be fully demixed, that is to say, a randomly chosen subset of the particles in a monodisperse packing is chosen to be large disks and grown to the desired size. Our derivation suggests that polydispersity does not change the random packing density to first order in the size dispersity, consistent with computational observations in Ref. [69].

Equation (8.7) gives a very simply prescription for how to increase the jamming density of a packing that is not in $\mathbf{D}-$equilibrium: Just set the sphere growth rates proportional to

$$\boldsymbol{\gamma} = \frac{\left(\widehat{\mathbf{A}}\mathbf{f}\right)^T \mathbf{D}}{\mathbf{e}^T \mathbf{D}^d} \mathbf{D}^{d-1} - \widehat{\mathbf{A}}\mathbf{f}, \tag{8.8}$$

that is, if a sphere is being pressed on too much by its neighbors, shrink its diameter, otherwise, increase it. Since the forces $\mathbf{f}$ can be calculated using MD by averaging over collisions, if the growth rates are small enough this should lead to a slow increase in the jamming density in an actual implementation, ending in a configuration that is in $\mathbf{D}-$equilibrium. We have made a preliminary implementation of this procedure in our MD algorithm for hard spheres, and examples of packings that we have produced with it are shown in Fig. 8.6.

Starting from a disordered binary packing of hard disks, like the one in Fig. 4.8, we start the MD simulation at a small δ (i.e., high pressure), and let the particle diameters change independently according to Eq. (8.8) at constant (reduced) pressure, where the forces are estimated from the average collisional momentum exchange between particles. The final packing, that is in $\mathbf{D}-$equilibrium, is shown in Fig. 8.6. As expected, it has more contacts per particle due to the additional N degrees of freedom present (the diameters of the disks), more specifically, we expect that $\bar{Z} \geq 2(d+1) = 6$. In fact, the contact number is exactly six, $\bar{Z} = 6$, and the density is very close to that of the triangular crystal of monodisperse disks, $\phi \approx 0.91$. Starting with a disordered packing of monodisperse hard spheres, and using the same MD recipe, we obtained a polydisperse packing of hard spheres with density $\phi \approx 0.68$, also shown in Fig. 8.6. The algorithm slowed down substantially in three dimensions and was not able to find a packing in true $\mathbf{D}-$equilibrium, therefore, we expect the density can be increased further. The important observation is that these packings

are disordered according to any reasonable criterion, however, they are much denser than what would be considered "random" or MRJ packings. For example, if we use the same distribution of radii as the one for the polydisperse packing of spheres in Fig. 8.6 (at $\phi \approx 0.68$) and perform the usual LS algorithm of compressing a liquid into a jammed packing, we obtain a packing with significantly lower jamming density $\phi \approx 0.645$.

8.5.2 Ellipsoids

We now generalize the argument in the previous section to ellipsoids. In the case of spheres the matrix $\widehat{\mathbf{A}}$ has 1's in column $\{i, j\}$ in rows i and j. In the case of ellipsoids these dense blocks become vectors $\tilde{\mathbf{a}}$,

$$\tilde{\mathbf{a}} = \left(\mathbf{O}^{-1}\mathbf{Q}\widehat{\mathbf{n}}\right) \diamond \left(\mathbf{Q}\mathbf{r}_C\right),$$

as derived in Section 2.2.5. Note that $\tilde{\mathbf{a}}^T\mathbf{o} = \widehat{\mathbf{n}}^T\mathbf{r}_C$, so that in this case,

$$\mathbf{f}^T\mathbf{h} = -\frac{\delta}{2}\sum_i \left(\widehat{\mathbf{A}}\mathbf{f}\right)_i^T \mathbf{o}_i = -\frac{\delta}{2}\sum_i \sum_{j\in\mathcal{N}(i)} \mathbf{f}_{ij}^T\mathbf{r}_{iC}^j = -\delta\sum_{ij}\mathbf{f}_{ij}^T\mathbf{r}_{ij}. \tag{8.9}$$

If the axes of the ellipsoids change, then the change in the jamming gap is

$$\Delta\delta = -\frac{\sum_i \left(\widehat{\mathbf{A}}\mathbf{f}\right)_i^T \Delta\mathbf{o}_i}{\sum_{ij}\mathbf{f}_{ij}^T\mathbf{r}_{ij}}.$$

At the same time, the relative change in the volume of one ellipsoid is

$$\Delta V/V = \mathbf{e}^T\mathbf{O}^{-1}\Delta\mathbf{o}.$$

The total change in jamming density therefore is:

$$\Delta\phi_J \approx \Delta\phi + d\phi\Delta\delta \sim \sum_i \left\{ V_i\left(\mathbf{e}^T\mathbf{O}_i^{-1}\Delta\mathbf{o}_i\right) - \left(d\sum_j V_j\right)\frac{\sum_i \left(\widehat{\mathbf{A}}\mathbf{f}\right)_i^T \Delta\mathbf{o}_i}{\sum_{ij}\mathbf{f}_{ij}^T\mathbf{r}_{ij}} \right\}.$$

Of particular interest is the case when all the ellipsoids have the same size and they all change their shape in unison, $\Delta\mathbf{o} = \mathbf{O}\boldsymbol{\kappa}$, where $\boldsymbol{\kappa} \neq \mathbf{e}$ is some vector. In this case the above simplifies to

$$\Delta\phi_J \sim \left[\mathbf{e} - d\frac{\sum_i \left(\widehat{\mathbf{A}}\mathbf{f}\right)_i^T \mathbf{O}}{\sum_{ij}\mathbf{f}_{ij}^T\mathbf{r}_{ij}}\right]^T \boldsymbol{\kappa} = \left[\mathbf{e} - d\frac{\sum_i\sum_{j\in\mathcal{N}(i)} f_{ij}\left(\mathbf{O}\tilde{\mathbf{a}}_{ij}\right)}{\sum_{ij} f_{ij}\left(\widehat{\mathbf{n}}_{ij}^T\mathbf{r}_{ij}\right)}\right]^T \boldsymbol{\kappa}.$$

Since $\mathbf{O}\tilde{\mathbf{a}} = \left(\mathbf{Q}\widehat{\mathbf{n}}\right) \diamond \left(\mathbf{Q}\mathbf{r}_C\right)$, we finally get a condition for $\mathbf{O}$-equilibrium

$$\sum_i \sum_{j\in\mathcal{N}(i)} f_{ij}\left(\mathbf{Q}_i\widehat{\mathbf{n}}_{ij}\right) \diamond \left(\mathbf{Q}_i\mathbf{r}_{iC}^j\right) = \frac{\sum_{ij} f_{ij}\left(\widehat{\mathbf{n}}_{ij}^T\mathbf{r}_{ij}\right)}{d}\mathbf{e} = p_{\text{int}}\mathbf{e}.$$

We have not yet implemented an MD algorithm that would search for ellipsoid packings in $\mathbf{O}-$equilibrium, however, this is an interesting investigation for the future.

8.6 Conclusions

In this chapter we used molecular dynamics to discover ellipsoid packings that surpass the density of the FCC lattice packing even for nearly-spherical ellipsoids. The maximum density of $\varphi \approx 0.7707$ is achieved for aspect ratio of $\sqrt{3}$, and in this densest-know ellipsoid packing each ellipsoid has 14 touching neighbors. There is nothing to suggest that the crystal packing we have presented in this chapter is indeed the densest for any

aspect ratio other than the trivial case of spheres. We believe it is important to identify the densest periodic packings of ellipsoids with small numbers of ellipsoids per unit cell. We proposed how to do this by using modern global optimization techniques, as has been done for various sphere and disk packing problems. However, this is a challenging project due to the complexity of the nonlinear impenetrability constraints between ellipsoids. In particular, the case of slightly aspherical ellipsoids is very interesting, as the best packing will be a perturbation of the FCC lattice with a broken symmetry, and should thus be easier to identify. Once the breaking of orientational symmetry is resolved, aspect ratios close to unity could be handled by smooth continuation of the structure of nearly spherical packings. The results of such investigations could be used to formulate a Kepler-like conjecture for ellipsoids and understand the high-density phase behavior of the hard-ellipsoid system.

Chapter 9

Pair Correlation Function of Jammed Sphere Packings

In this chapter we study the approach to jamming in hard-sphere packings, and, in particular, the pair correlation function $g_2(r)$ around contact, both theoretically and computationally [75]. Our computational data unambiguously separates the narrowing delta-function contribution to g_2 due to emerging interparticle contacts from the background contribution due to near contacts. The data also shows with unprecedented accuracy that disordered hard-sphere packings are strictly isostatic, i.e., the number of exact contacts in the jamming limit is exactly equal to the number of degrees of freedom, once rattlers are removed. For such isostatic packings, we derive a theoretical connection between the probability distribution of interparticle forces $P_f(f)$, which we measure computationally, and the contact contribution to g_2. We verify this relation for computationally-generated isostatic packings that are representative of the maximally random jammed state. We clearly observe a maximum in P_f and a nonzero probability of zero force, shedding light on long-standing questions in the granular-media literature. We computationally observe an unusual power-law divergence in the near-contact contribution to g_2, persistent even in the jamming limit, with exponent -0.4 clearly distinguishable from previously proposed inverse square root divergence. Additionally, we present high-quality numerical data on the two discontinuities in the split-second peak of g_2, and use a shared-neighbor analysis of the graph representing the contact network to study the local particle clusters responsible for the peculiar features. Finally, we present the first computational data on the contact-contribution to g_2 for vacancy-diluted FCC crystal packings and also investigate partially crystallized packings along the transition from maximally disordered to fully ordered packings. We find that the contact network remains isostatic even when ordering is present. Unlike previous studies, we find that ordering has a significant impact on the shape of P_f for small forces.

9.1 Introduction

Jamming in hard-sphere packings has been studied intensely in past years (see [4, 6] and references therein). In this Chapter we investigate the pair correlation function $g_2(r)$ of the classical three-dimensional hard-sphere system near a jamming point for both disordered (amorphous, often called random) as well as ordered (crystal) jammed packings. The basic approach follows that of Ref. [52], developed further for crystal packings of rods, disks and spheres in Ref. [145]. We focus on *finite* sphere packings that are almost *collectively jammed* [15, 71], in the sense that the configuration point is trapped in a very small region of configuration space around the point representing the jammed *ideal* packing [71]. Difficulties with extending the results to infinite packings will be discussed in what follows. In the ideal jammed packing particle contacts necessary to ensure jamming are exact, and the particles cannot at all displace, even via collective motions. Such ideal jammed (or rigid) packings have long been the subject of mathematical inquiry [138]; however, they are not really attainable in numerical simulations where produced packings invariably have some interparticle gaps (even taking into account the unavoidable roundoff errors). It is therefore instructive to better understand the approach to this ideal jammed state computationally and theoretically, which is

our primary objective.

We choose as our main tool of exploration the shape of the venerable orientationally-averaged pair correlation function $g_2(r)$ around contact. This is because this function is a simple yet powerful encoding of the distribution of interparticle gaps. In the jamming limit, it consists of a delta function due to particle contacts and a background part due to particles not in contact. As the jamming limit is approached, it is expected that the delta-function contribution will become more localized around contact. We derive the first exact theoretical model for this narrowing for isostatic packings (defined below), connecting g_2 to the probability distribution of interparticle forces P_f, and verify the relation numerically. In this work, we present computational data with unprecedented proximity to the jamming limit, for the first time clearly separating the narrowing delta-function contribution from the apparently persistent diverging background contribution. The data show that our disordered packings have an exactly isostatic contact network in the jamming limit, but with an unusual multitude of nearly closed contacts, manifested as a power-law divergence in the near-contact contribution to g_2, persistent even in the jamming limit, with exponent -0.4 clearly distinguishable from previously proposed inverse square root divergence. We study the properties of the contact network and find, contrary to previous studies, no traces of polytetrahedral packing, but rather a complex local geometry, indicating that the geometric frustration due to the constraints of global jamming on the local geometry is nontrivial. Additionally, we study the evolution of the salient features of $g_2(r)$ along the transition from maximally disordered to fully ordered packings by inducing partial crystallization in the packings. We find that both $g_2(r)$ and P_f are significantly affected by crystallization; however, the contact network remains isostatic. We thus demonstrate by example that isostaticity is not synonymous with randomness.

In this Chapter we focus on short-ranged pair correlations, i.e., on the features of the pair correlation function close to contact. In Chapter 10 we will continue our study of pair correlations in disordered jammed hard-sphere packings by focusing instead on the long-ranged tail of g_2, or equivalently, on the behavior of the structure factor $S(k)$ for small k. In Chapter 11 we will examine random sphere packings in dimensions higher than three, and we will find a very similar behavior to the one we study in detail here for $d = 3$.

9.2 Theoretical Considerations

As discussed in Sections 4.3.1.1 and 5.7.1, for jammed sphere packings, in the jamming limit $\delta \to 0$ the set of displacements that are accessible to the packing $\mathcal{J}_{\Delta\mathbf{R}}$ approaches a convex *polytope* $\mathcal{P}_{\Delta\mathbf{R}}$ [52, 145], which is a simplex for isostatic packings. A troublesome aspect, discussed in Ref. [52], is that infinite packings can never be jammed in the above sense unless $\delta = 0$, due to the appearance of unjamming mechanisms involving collective density fluctuations. Nevertheless, computational studies indicate that macroscopic properties derived using this polytope-based approach do not depend on N, even as $N \to \infty$. Looking at Eq. (5.24), we see that in the jamming limit the reduced pressure is asymptotically given by the free-volume equation of state [52],

$$p = \frac{PV}{NkT} = \frac{1}{\delta} = \frac{d}{(1 - \phi/\phi_J)}. \tag{9.1}$$

Relation (9.1) is remarkable, since it enables one to accurately determine the true jamming density of a given packing even if the actual jamming point has not yet been reached, just by measuring the pressure. We later numerically confirm the validity of Eq. (9.1) in the vicinity of the jamming point.

9.2.1 Interparticle Force Networks

During a molecular dynamics simulation, the particles travel around $\mathbf{R}_J$ and the configuration explores $\mathcal{P}_{\boldsymbol{\Delta}\mathbf{R}}$. One can average the exchange of momentum between any two pairs of particles which share a contact in the jammed limit (i.e., whose contact forms a face of $\mathcal{P}_{\Delta\mathbf{R}}$), hereafter referred to as *first neighbors*, to obtain an average interparticle *force* (momentum transfer per unit time or impulse [117]). This kind of method for measuring interparticle forces has previously been used in work on dense granular flows [117]. The vector of collisional forces $\mathbf{f}$ compares directly to the inter-grain force networks which have been the subject of intense experimental and theoretical study in the field of granular materials [3, 209, 16, 210]. As discussed in Chapter 5, the existence of a self-stress $\mathbf{f}$ is a necessary condition for jamming, specifically, these forces

must be in equilibrium,

$$\mathbf{A}\mathbf{f} = \mathbf{0}, \tag{9.2}$$

they must be nonnegative, $\mathbf{f} \geq 0$, and here we normalize them to have a unit average, $\bar{f} = \mathbf{e}^T\mathbf{f}/M = 1$, in the tradition of the granular media literature. Our numerical investigations indicate that indeed the set of time-averaged collisional forces approach local equilibrium as the time horizon T of the averaging increases, in a inverse-power law manner, $\|\mathbf{A}\mathbf{f}\| \sim T^{-1}$. We can therefore obtain interparticle forces relatively accurately given sufficiently long molecular dynamics runs. While Eq. (9.2) will have a unique solution if and only if the contact network of the packing is isostatic, even for hyperstatic packings, such as the FCC packing, the equilibrium set of forces should be unique. In fact, one can prove that the force between two particles will be proportional to the surface area of the face of $\mathcal{P}_{\mathbf{\Delta R}}$ formed by the contact in question.

It is interesting to observe that if one has an *arbitrary* point $\mathbf{\Delta R} \in \mathcal{P}_{\mathbf{\Delta R}}$, the interparticle gaps due to nonzero jamming gap will be $\mathbf{h} \approx \mathbf{A}^T\mathbf{\Delta R} + \delta D\mathbf{e}$, so that

$$\mathbf{f}^T\mathbf{h} \approx (\mathbf{A}\mathbf{f})^T\mathbf{\Delta R} + MD\delta = MD\delta. \tag{9.3}$$

Eq. (9.3) enables one to determine how far from the jamming density a packing is without actually reaching the jamming point. This can be a useful alternative to using Eq. (9.1) when the hard-sphere pressure is not available, but interparticle forces are, such as, for example, with packings generated by algorithms using stiff "soft" spheres [31].

As discussed in Chapter 5, hypostatic packings cannot be jammed. However, it is possible for a hypostatic packing to be *locally maximally dense*, in the sense that no continuous motion of the particles can increase the density to first order. In other words, the particles must first move and unjam (which must be possible for a hypostatic sphere packing) before the density can increase. In particular, a packing of contacting particles for which a set of interparticle forces $\mathbf{f}$ in equilibrium exists, is locally maximally dense. In a sense, the interparticle forces resist further increase of the density. As we discuss later, our packing generation algorithm sometimes terminates with such packings since it tries to continually increase the density.

9.2.2 Pair Correlation Function Around Contact

We now turn to the central subject of this work: The shape of the (orientation-averaged) pair correlation function $g_2(r)$ for small jamming gaps. In particular, we will focus on interparticle distances r that are very close to D. We express $g_2(l)$ in terms of the nonnegative *interparticle gaps* $l = r - D$. Note that so far we have denoted gaps with h, however, in this section, we will use h to denote heights in a simplex. Only the M first-neighbor particle pairs will contribute to the shape of $g_2(l)$ right near contact, i.e., for gaps up to $l_{\max}$, where $l_{\max}$ is the largest distance from the centroid of $\mathcal{P}_{\mathbf{\Delta R}}$ to one of its faces. This contribution will become a delta function in the jamming limit. Particle pairs not in contact will not contribute to $g_2(l)$ until gaps larger than the minimal further-neighbor gap l_{FN}, and for now we will implicitly assume that $l_{FN} \gg l_{\max}$, so that there is a well-defined *delta-function region* $g_2^{(\delta)}(l) \equiv g_2(l \ll l_{FN})$. This delta function region has previously been investigated theoretically for crystal packings, primarily [145]. In this work, we derive exact theoretical expressions for this region for isostatic packings, as well as numerically study vacancy-diluted FCC crystals and partially crystallized packings.

9.2.2.1 Isostatic Packings

We first focus on the probability distribution for observing an interparticle gap l, $P_l(l)$, which is related to $g_2^{(\delta)}(l)$ via a simple normalization factor. The contribution $\tilde{P}(l)$ from a specific contact is determined from the area $\tilde{S}(l)$ of the cross section of $\mathcal{P}_{\mathbf{x}}$ with a plane parallel to the face corresponding to the contact and at a distance l from the face, $\tilde{P}(l) \sim \tilde{S}(l)$ [145]. The critical observation we make is that for an isostatic contact network, $\mathcal{P}_{\mathbf{\Delta R}}$ is a *simplex* and thus immediately we get $\tilde{S}(l) \sim [(h-l)/h]^M$, where h is the height of the simplex corresponding to this particular face, $h = M\left|\mathcal{P}_{\mathbf{\Delta R}}\right|/S$, $S = \tilde{S}(0)$. After normalization of $\tilde{P}(l)$ and averaging over all interparticle contacts, we obtain that

$$P_l(l) = \int_{h=l}^{\infty} \frac{M}{h}\left[1 - \frac{l}{h}\right]^M P_h(h)dh,$$

which shows that if we know the distribution P_h of heights for the simplex, or equivalently, the distribution of surface areas S of the faces of the polytope $P_S(S)$, we would know P_l and thus $g_2^{(\delta)}$.

Since the interparticle force $f \sim S$, we see immediately that the distribution of face areas is equivalent to the distribution of interparticle forces $P_f(f)$, as we derived in Section 6.2.2.1,

$$h = \frac{M}{f}\Delta D,$$

where $\Delta D = \delta D$. This gives in the limit $M \to \infty$

$$\begin{aligned} P_l(l) = \int_{f=0}^{M/l} \frac{f}{\Delta D}\left[1 - \frac{lf}{M\Delta D}\right]^M P_f(f)df \quad &\approx \\ \frac{1}{\Delta D}\int_0^\infty fP_f(f)\exp(-fl/\Delta D)df \quad &= \quad \frac{1}{\Delta D}\mathcal{L}_{l/\Delta D}\left[fP_f(f)\right], \end{aligned}$$

where $\mathcal{L}_s$ denotes the Laplace transform with respect to the variable s. We have the normalization condition $\int_0^\infty P_l(l)dl = 1$ and additionally

$$DP_l(0) = \frac{D}{\Delta D}\int_0^\infty fP_f(f)df = \frac{D}{\Delta D} = p.$$

If we now relate $P_l(l)$ to $g_2^{(\delta)}(l)$,

$$g_2^{(\delta)}(l) = \frac{2MV}{4\pi D^2 N^2}P_l(l) = \frac{\bar{Z}D}{24\phi}P_l(l),$$

where $\bar{Z} = 2M/N = 2d = 6$ is the *mean coordination number*, we obtain the central theoretical result

$$g_2^{(\delta)}(l) = \frac{p}{4\phi}\mathcal{L}_{l/\Delta D}\left[fP_f(f)\right]. \tag{9.4}$$

9.3 Computational Results

We use event-driven molecular dynamics [11] as the primary computational tool for our investigations. This enables us to perform exact molecular dynamics on hard-particle packings very close to the jamming point, which is not possible with traditional time-driven molecular dynamics algorithms, as discussed in Chapter 3. The algorithm monitors a variety of properties during the computational run, including the "instantaneous" pressure, as calculated from the total exchanged momentum in all interparticle collisions during a certain short time period Δt. By allowing the shape of the particles to change with time, for example, by having the sphere diameter grow (shrink) uniformly at a certain (possibly negative) expansion rate $dD/dt = 2\gamma$, one can change the packing density. If the change is sufficiently slow, the system will be in approximate (metastable) equilibrium during the densification, and one can rather effectively gather quasi-equilibrium data as a function of density.

Event-driven molecular dynamics (see Ref. [11] and references therein) in which the particles (quickly) grow in size in addition to their thermal motion at a certain expansion rate, starting from a random (Poisson) distribution of points, produces a jammed state with a diverging collision rate. This is the well-known Lubachevsky-Stillinger (LS) packing algorithm [12, 13], which we have used and modified [11] to generate all the disordered hard-sphere packings for this study. During the initial stages, the expansion has to be fast to suppress crystallization and maximize disorder [56], and delaying further discussion to later sections, we will assume that the disordered packings used in this study are representative of the MRJ state. It is important to note that the algorithm typically produces packings that have *rattling* particles, i.e., particles that do not have true contacts with particles in the jammed *backbone*[1] of the packing, and can be removed without affecting the jamming category of the final packing. We will discuss procedures for identification of such rattlers in what follows.

[1] The backbone is formed by the collection of particles that participate in the jamming force network [71].

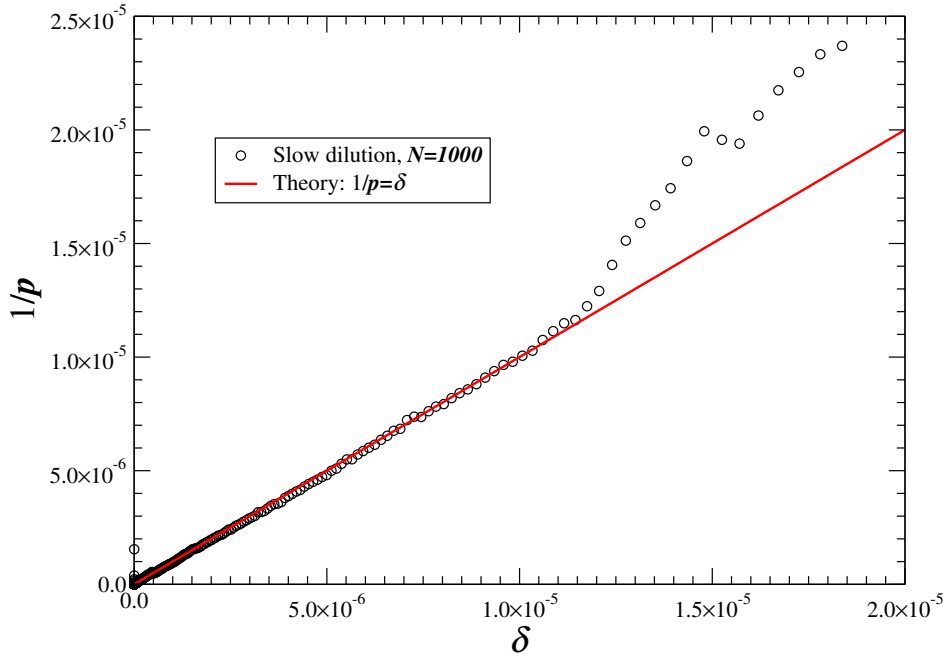

Figure 9.1: The inverse of the "instantaneous" (averaged over several hundred collisions per particle) pressure of a nearly jammed (isostatic) packing of 1000 particles, as it is slowly diluted (using a negative expansion rate for the particles in the molecular dynamics algorithm $\gamma = -10^{-5}$) from $\phi_J \approx 0.627$ until an unjamming particle rearrangement occurs. Up to this occurrence, the free-volume theoretical relation $p = \delta^{-1}$ is satisfied to very high accuracy. There is a short transient region during the initial equilibration of the packing. Rattlers have been removed from the packing.

To our knowledge, no verification of the exactness of Eq. (9.1) for disordered packings exists in the literature. The perfect FCC crystal is stable until rather low densities, and the pressure seems to be rather accurately predicted by the free-volume approximation in a wide range of densities around close packing. This has been observed in the literature and a suitable corrective term was determined [54]. However, for disordered packings, previous studies have identified a coefficient smaller than 3 in the numerator, namely 2.67 [211, 53]. In Fig. 9.1, we numerically confirm the validity of Eq. (9.1) with very high accuracy for disordered packings. In Fig. 9.2, we show the change of the coefficient (the configurational heat capacity in units of Nk) $C = (1-\phi/\phi_J)p$ with density. Agreement with the theoretical $C = d = 3$ is observed sufficiently close to the jamming point, but with rapid lowering of the coefficient from 3 away from the jamming point. This is because for sufficiently large jamming gaps, contacts other than the M true contacts start contributing to the collisions, and the polytope-based picture we presented so far does not apply exactly. We demonstrate this in Fig. 9.2 by showing the number of contacts which participate in collisions (*active contacts*) as the jamming point is approached. Our investigations indicate that previous studies did not examine at the range of densities appropriate for the theory presented above and did not properly account for the rattlers.

9.3.1 Disordered Packings

We have verified in previous publications that LS packings are typically collectively jammed [72] using a testing procedure based on linear programming [71]. Unfortunately, the linear programming library used in the implementation cannot really achieve the kind of numerical accuracy that we require in this work, specifically that for packings which are jammed almost to within full numerical precision ($\delta = 10^{-15} - 10^{-12}$). Additionally, it cannot handle three-dimensional packings of more than about a thousand particles. Another test for jamming, which we have found to be reliable for the purposes of this work, is to take the final packing produced by the LS procedure and then run standard event-driven molecular dynamics on it for long periods

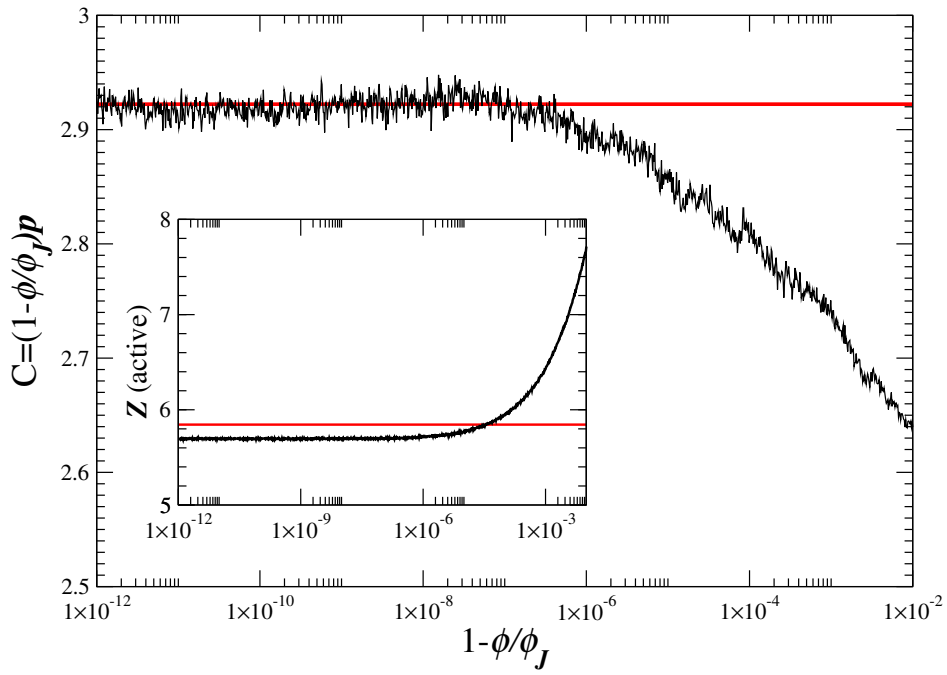

Figure 9.2: The coefficient C during a typical slow densification (expansion rate is 10^{-4}) of a 10,000-particle system, starting from an equilibrated liquid at $\phi = 0.5$ up to jamming. The final packing has 259 rattlers, so the expected coefficient is $3 \cdot 0.9741 \approx 2.92$, a value which is shown with a red line. It is clear that close to the jamming point Eq. (9.1) is very accurate, but a marked lowering from a coefficient of 3 is seen for pressures lower than about 10^6, likely explaining the coefficient 2.67 reported in works of Speedy [211, 53]. The inset shows the estimated "collisional" coordination, defined as the average number of different particles that a particle has collided with during a time interval of about 100 collisions per particle, during the same densification. The expected number $6 \cdot 0.9741 \approx 5.85$ is shown (this number is not asymptotically reached exactly since some of the M contacts do not participate in collisions frequently enough to be registered during the time interval used), and we see that as many as 8 contacts per particle are active at sufficiently large jamming gaps.

of time (on the order of thousands to hundreds of thousands of collisions per particle) and monitor the "instantaneous" pressure. We called this the shrink-and-bump test for jamming in Section 4.3.3.1. If the packing is jammed, this pressure will be stable at its initial value. However, if the packing is not truly jammed, we have observed that the pressure slowly decays with time, the slower the "pressure leak" the more "jammed" the initial packing is, as illustrated in Fig. 9.3. Similar observations are made in Ref. [211]. In addition, we track the average particle displacement (from the initial configuration) and check to see if there is a systematic drift with time away from the initial configuration. The two tests always agreed: A pressure leak always corresponds to a systematic drift away from the initial configuration. Note that this kind of test would not work for hypostatic jammed packings of ellipsoids, since there the pressure exhibits large oscillations, as discussed in Section 5.7.3.

We have observed that LS packings densified to within numerical capability only pass this rigorous jamming test of having no pressure leak if during the final stages of the LS densification the expansion rate is very small compared to the average thermal velocity (maintained constant via a velocity rescaling thermostat [11]) of the particles (about five orders of magnitude or less). Similar observations are made in Ref. [211]. If the expansion rate is too fast, we have found that the packings jam in slightly hypostatic configurations, where there are not enough particle contacts to ensure jamming. In particular, some particles have 2 or 3 contacts (and of course rattlers are present). In order for a set of balanced forces to exist (which as we discussed is a necessary condition for a packing to be locally maximally dense) when a particle has less than

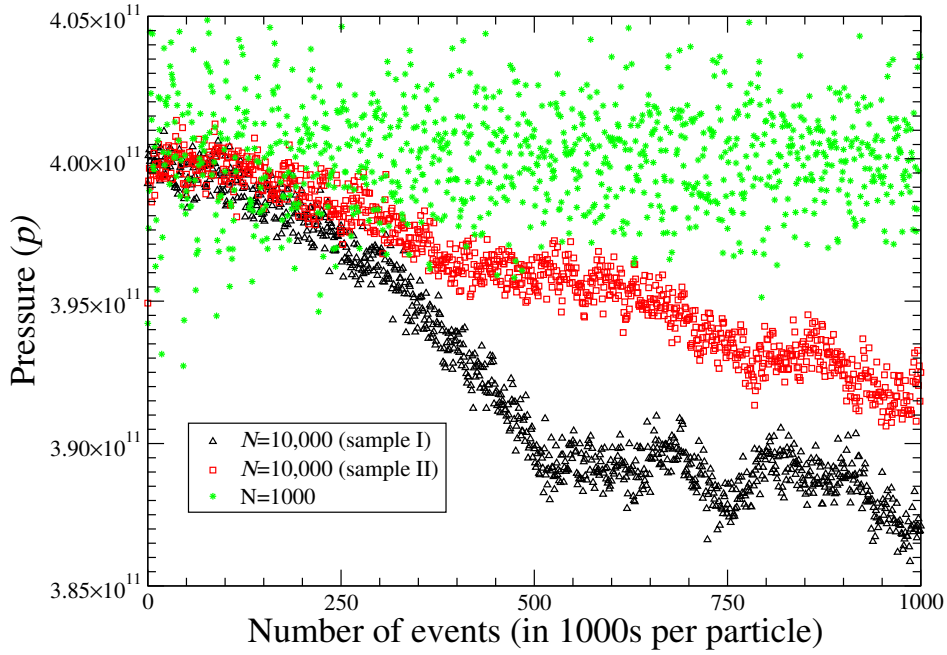

Figure 9.3: The short-term ("instantaneous") pressure versus number of events (mostly binary collisions) processed by the molecular dynamics algorithm [11], corresponding to a total run of about half a million collisions per particle. For the 1000-particle packing the pressure is stable, but for the larger packings a systematic pressure leak is observed.

4 contacts, these contacts must be in a degenerate geometric configuration, namely 3 coplanar or 2 collinear contacts. We have indeed verified that this is what happens in the hypostatic packings produced by the LS algorithm. The number of such geometric peculiarities increases with increasing expansion rate, and also for more ordered packings, as we discuss later.

We illustrate the progress of the densification during the final stages of the algorithm in Fig. 9.4. The figure shows, for several snapshots of the packing during the densification, the *cumulative coordination number*

$$Z(l) = \frac{N}{V}\int_{r=D}^{D+l} 4\pi r^2 g_2(r)dr = 24\phi\int_{r=D}^{D+l}\left(\frac{r}{D}\right)^2 g_2(r)\frac{dr}{D},$$

i.e., the average number of particles within a gap l from a given particle. We we will often use this quantity instead of $g_2(l)$. For the first time in the vast literature on random hard-sphere packings, a clear separation is seen between the delta-function contribution $Z^{(\delta)}(l)$, which becomes more localized around contact, and the background increase in the mean coordination from the isostatic contact value of $\bar{Z} = 6$, which remains relatively unaffected by the densification. For small packings ($N = 1000$), the value of $Z(l)$ is fixed at 6 for a remarkably wide range of gaps, as much as 9 orders of magnitude for the final packings. Fast densification is seen to lead to hypostatic packings in Fig. 9.4, leaving a certain fraction of the contacts "open". Stopping the expansion invariably leads to a decay of the macroscopic pressure for such hypostatic packings.

By using heuristic strategies, we were able to find (slow) densification schemes which produced packings which are indeed ideally jammed within almost full numerical precision, at least for packings of $N = 1000$ particles or less. In fact, the plateau in $Z(l)$ was at exactly (up to a single contact!) an isostatic number of contacts, $M = 3N - 2$, for all the packings produced via a carefully guided LS algorithm. It is essential that here N is the number of particles in the jammed backbone of the packing [71], i.e., rattlers [13] with fewer than 2 contacts have been removed from the packing. It seems that the algorithm produces packings with about 2.2% rattlers, and so the density of the disordered packings we look at is typically $\phi \approx 0.625 - 0.630$, rather than the widely known $\phi \approx 0.64$. Despite a concentrated effort and lots of expended CPU time, we

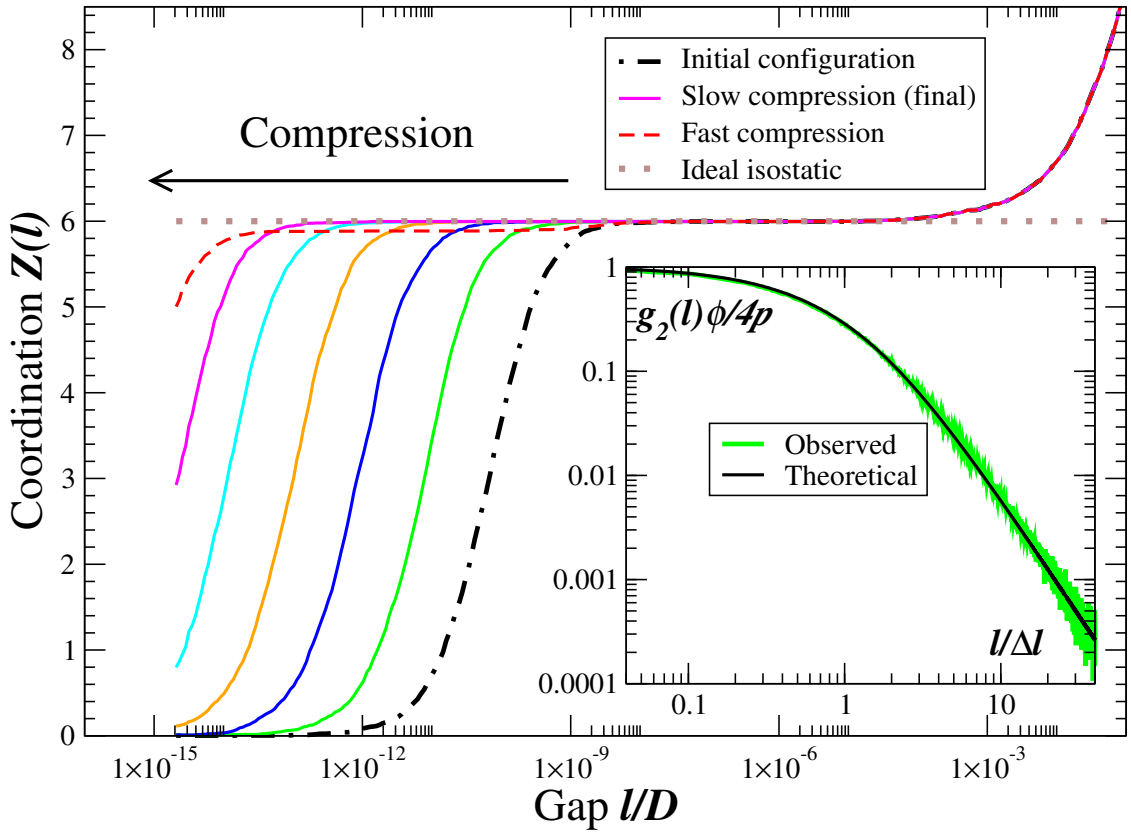

Figure 9.4: The cumulative coordination $Z(l)$ (i.e., the integral of $g_2(l)$) as a function of the gap tolerance l, for a sequence of snapshots of a 1000-particle packing during the final compression stages of the LS algorithm. Each snapshot is shown with a separate solid curve and only the last one is labeled in the figure legend. For a sufficiently slow expansion (expansion rate is 10^{-5} times the average thermal velocity), the packing is clearly seen to jam in an isostatic configuration. A hypostatic configuration is found for fast expansion (expansion rate is comparable to the thermal velocity). The inset shows the properly normalized derivative of $Z(l)$, right around contact, along with a comparison to our semi-theoretical prediction for $g_2^{(\delta)}(l)$, for a packing with $\delta = 2.5 \cdot 10^{-12}$.

have been unable to achieve true isostaticity for $10,000$-particle packings[2]. This is illustrated in Fig. 9.3, where it is clearly seen that the pressure in the large packings does not remain constant over long periods of time (about a million collisions per particle). It is therefore not strictly justified to consider these packings within the framework of ideal jammed packings that we have adopted here. However, it is readily observed that over finite and not too long time intervals (for example several thousands of collisions per particle), the large packings conform to the predictions of the theory developed here. In particular, the collisional forces form a balanced force network with essentially the same $P_f(f)$ as the truly jammed smaller packings, and the pressure is given by Eq. (9.1) with very high accuracy, where δ can be determined, for example, via Eq. (9.3). We have observed no systematic differences in any of the correlation functions or distributions between the jammed isostatic packings with 1000 particles and the ones with 10000 particles, other than the better binning resolution of the larger packings and larger statistical variability among the small packings. Results given in subsequent sections will indicate that the fact we could not achieve true jamming for very large packings is an inherent property of the kinds of packings we consider, rather than a failure of the simulation method. We therefore believe it is justified to use the larger packings for certain analysis where better statistics are needed.

[2]When carefully densified, the packings typically lacked only a few contacts to achieve isostaticity.

The main goal of this work is to explore and explain Fig. 9.4, and in particular, to investigate both the "delta-function", or *contact*, contribution $g_2^{(\delta)}$, which should integrate to produce the isostatic average coordination $\bar{Z} = 2M/N = 6$, and the "background" or *near-contact* $g_2^{(b)}$, for gaps from about $100\delta D$ to $10^{-1}D$. This latter one has already been observed in an experimental study of hard spheres [199], and in computational studies of stiff "soft" spheres [32, 31], including work on soft spheres that followed ours and is closely related [212]. These various studies find a nearly square-root divergence, $g_2^{(b)}(l) \sim 1/\sqrt{l}$, and Ref. [32] observes that this is an integrable divergence and thus clearly separate from the delta function. Our results, shown in Fig. 9.4, are the first unambiguous and precise separation of the two pieces of the pair correlation function around contact near jamming for hard spheres. Our numerical data has precision ($\delta < 10^{-13}$) not previously attained, since such proximity to ideal jammed hard-sphere packings can only be achieved in a true hard-sphere algorithm, and at present only event-driven molecular dynamics seems to provide the required numerical robustness. It is rather interesting that although graphs showing the hard-sphere $g_2(l)$ in the literature have clearly demonstrated a divergence in $g_2(l)$ near contact for at least three decades [20], this seems to never have been clearly documented or investigated. We are led to believe that researchers were under the false impression this divergence is a signature of the delta-function contribution, and thus expected it to further narrow and disappear at true jamming.

9.3.1.1 Delta-Function (Contact) Contribution

We first verify that our theory correctly predicts the shape of $g_2^{(\delta)}(l)$. In order to verify relation (9.4) numerically, a form for $P_f(f)$ is needed. Force networks in particle packings have been the subject of intense theoretical and experimental interest [3, 210, 169, 16, 213], and it has been established that P_f decays exponentially at large forces for a variety of models. The behavior of P_f for small forces has not been agreed upon, the central question being whether the infinite-system-limit $P_f(0)$ is nonzero. No theoretical model has been offered yet that truly answers this question. We note that a recent model reproduces all of the major characteristics of P_f that we observe, including a positive $P_f(0)$, even though it is presently restricted to two dimensions [214]. Part of the difficulty is that the answer likely depends not only on the system in question, but also on the definition of f. In a true ideal collectively jammed isostatic packing, which is necessarily finite, all interparticle forces, *must* be strictly positive, and in fact are determined uniquely through Eq. (9.5),

$$\mathbf{f} = \begin{bmatrix} \mathbf{A} \\ \mathbf{e}^T \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{0} \\ 1 \end{bmatrix}, \tag{9.5}$$

without any mention of interparticle potentials or influence of external fields or loads like gravity, or thermal dynamics. The limiting probability distribution of these interparticle forces as the packing becomes larger, if it exists, can be positive at the origin, indicating that finite but large packings have limiting polytopes with a few extremely small faces, or equivalently, are very elongated along certain directions. We have numerically studied the form of $P_f(f)$ for almost jammed random packings of $N = 1000$ and $N = 10,000$ spheres by using molecular dynamics to observe the collisional forces between first neighbors, and also by directly using Eq. (9.5) for the smaller packings[3] (this offers better accuracy for small forces). The results are shown in Fig. 9.5. We clearly see a peak in $P(f)$ for small forces, as observed in the literature for jammed packings of soft particles [169], and it appears that there is a finite positive probability of observing zero interparticle force. We will return to this point later.

The observed $P_f(f)$ can be well fitted for medium and large forces by $P_f(f) = (Af^2 + B)e^{-Cf}$, with a small correction needed to fit the small-force behavior, as used in Fig. 9.5. This small correction has a negligible impact on the Laplace transform of $fP_f(f)$, and in fact a very good approximation to $g_2^{(\delta)}(l)$ in Eq. (9.4) is provided by just using

$$\mathcal{L}_x\left[fP_f(f)\right] = \frac{6A}{(x+C)^4} + \frac{B}{(x+C)^2}. \tag{9.6}$$

In the inset in Fig. 9.4, we show a comparison between the $g_2^{(\delta)}(l)$ we observe computationally and the one given by Eqs. (9.4) and (9.6) and the empirical fit to $P_f(f)$ in Fig. 9.5. An essentially perfect agreement is

[3]Efficiently inverting the rigidity matrix for very large three dimensional packings is a rather challenging numerical task which we have not yet tackled.

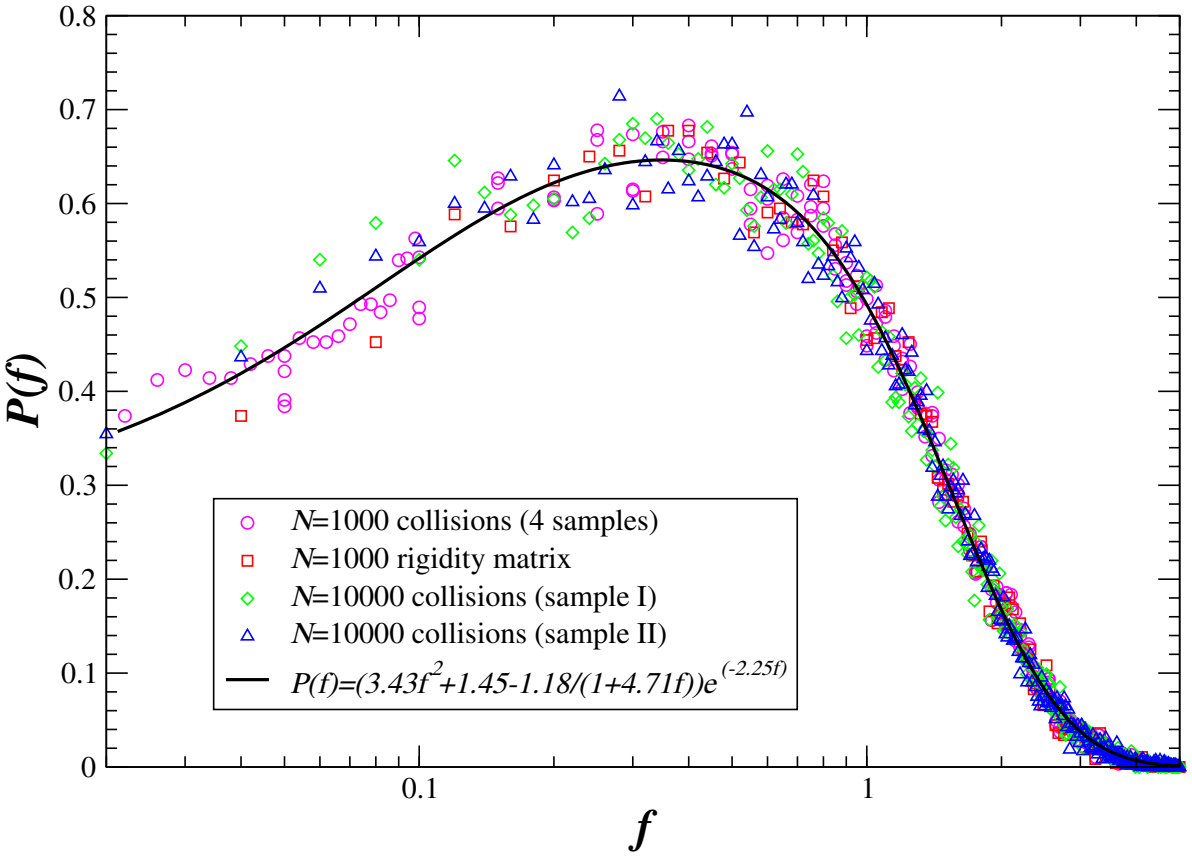

Figure 9.5: Computational data on the interparticle force distribution along with the best fit we could achieve. Packings of both 1,000 and 10,000 particles, using either molecular dynamics to average the collisional forces, or inversion of the rigidity matrix, were used, consistently producing the same probability distribution. Comparison to other data in the granular-media literature is beyond the scope of this work.

observed. Our focus here is on small forces, however we do wish to note that our data cannot confidently rule out a Gaussian component to P_f for large forces and that a slight quadratic component does seem to be visible when $P_f(f)$ is plotted on a log-log plot.

9.3.1.2 Near-Contact Contribution

In Fig. 9.6 we investigate the near-contact contribution to $g_2(l)$. We have found that $Z^{(b)}(l)$ has a power law behavior over a surprisingly wide range of gaps, up to the first minimum of g_2 at $l \approx 0.25D$, $Z^{(b)}(l) \approx 11(l/D)^{0.6}$, as shown in the figure. Note that this range is too wide for

$$g_2^{(b)}(x) = \frac{1}{24\phi(1+x)^2}\frac{dZ^{(b)}(x)}{dx}$$

to be a perfect power law, where $x = l/D$, as used to fit numerical data in other studies (which have not investigated nearly as wide a range of gaps as we do here) [32, 199]. The observed exponent is clearly distinguishable from an inverse square root divergence in $g_2^{(b)}(l)$, as proposed in the literature [32], and it is consistent with the experimental exponents reported in Ref. [199]. Our study has higher statistical accuracy than previously realized, however, it is not clear if there aren't also systematic effects due to the different protocols used to prepare the packings in studies such as Ref. [32].

We do not have a theoretical explanation for this functional behavior of $Z^{(b)}(l)$, however, the remarkable quality of the fit in Fig. 9.6 hints at the possibility of a (simple) scaling argument. Some simple observations

can be made by assuming that

$$Z(x) = \bar{Z} + ax^{1-\alpha} \text{ for } 0 < x \leq \beta, \tag{9.7}$$

where α is an exponent $0 \leq \alpha \leq 1$, and $\beta < 1$ determines the extent of this power-law dependence. The corresponding pair correlation function of course exhibits an inverse power-divergence with exponent α, except when $\alpha = 1$, when it is identically zero[4]. The exponent α clearly will depend on the amount of order present in the packing, i.e., the position of the packing in the density-order diagram of Fig. 1.4. We expect that it will increase with increasing order, since $\alpha \to 0$ would indicate a constant $g_2(l)$ near contact, a signature of the ideal gas, while $\alpha \to 1$ would indicate a clear distinction between the first and second shell of neighbors (i.e., a wide range of gaps with very few contacts) typical of crystal packings. Under the assumption that a power-law divergence in g_2 is appropriate, an intermediate value of α between 0 and 1, as we find numerically, is therefore expected. Some bounds on the range of possible α can be obtained from bounds on $Z(x)$ derived from geometric constraints (for example, $Z(x) < 13$ for a certain range of x since the sphere kissing number is 12 in three-dimensions), but the exact value is not simple to predict[5].

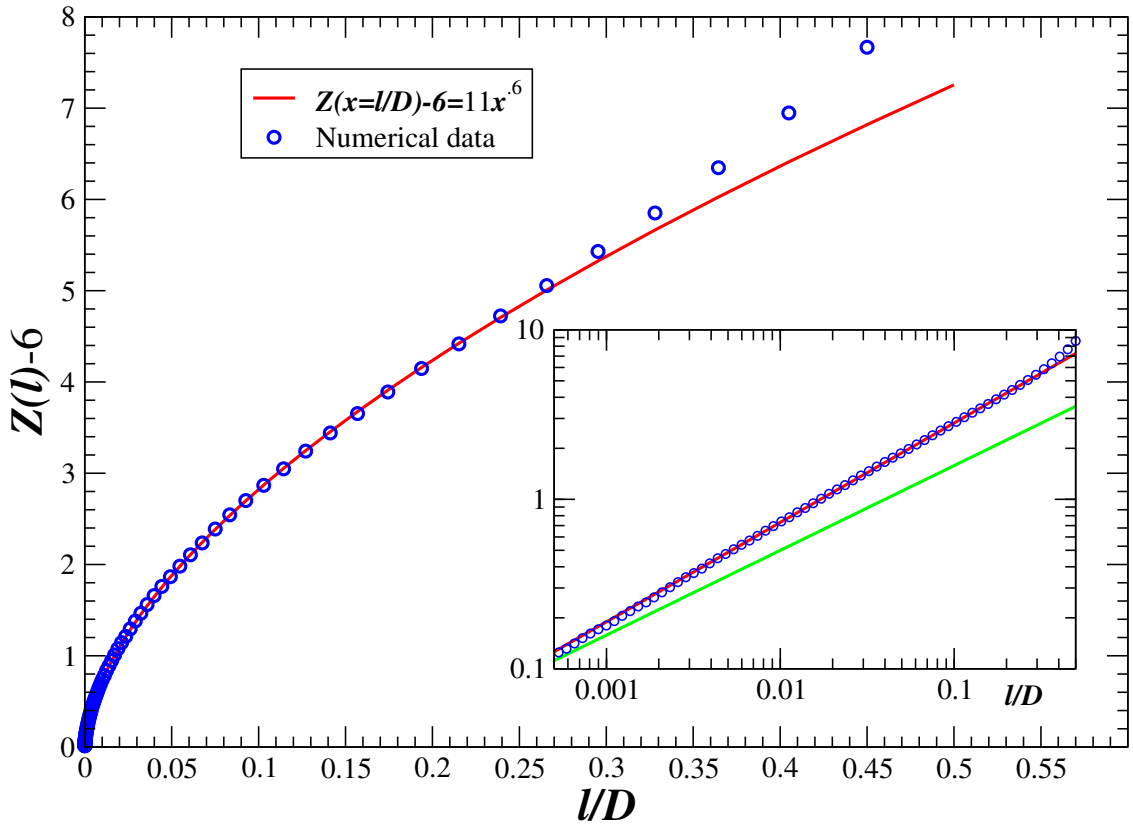

Figure 9.6: The near-contact $Z^{(b)}(l)$ for a nearly-jammed 10,000-particle packing, along with a power law fit for small gaps, shown in both a linear-linear scale and a log-log scale (inset). In this inset we also show a line with slope 0.5 (i.e., a square root dependence), which is clearly inconsistent with the numerical data. The rattlers have been removed from the packing.

It is important to emphasize that we have removed the rattlers from the packing before computing $Z(x)$. In Ref. [212] it is argued that this is unphysical, and that rattlers should be kept in the packing, which is reported to lead to a power exponent $\alpha \approx 0.50$. Indeed, if rattlers are kept and $\bar{Z}$ adjusted accordingly, we

[4]Note that $g_2^{(b)}(x)$ cannot have a simple-pole divergence since this would lead to a logarithmic divergence in $Z^{(b)}(x)$, which must be finite for all finite x.

[5]The three parameters α, β and a are thus not independent of one another. For example, requiring that $g_2^{(b)}(x) > 1$ and $Z(x) < 12$ for $0 < x \leq \beta$ gives the weak constraints $a(1-\alpha) > 24\phi\beta^2(\beta-1)^\alpha$ and $a(\beta-1)^{1-\alpha} < 12 - \bar{Z}$.

find a power-law divergence in $Z(x)$ with exponent between 0.5 and 0.6, closer to 0.5. However, we wish to emphasize that the quality of the power law fit is much worse than that shown in Fig. 9.6. Additionally, the exact placement of the rattlers inside their cage of neighbors affects the observed $Z(x)$. For example, an algorithm that tries to keep particles in contact, such as the Zinchenko algorithm [28], might produce a packing in which a rattler touches d of its caging particles. This would increase the observed contact number. On the other hand, algorithms such as the LS algorithm will randomly sample over the different rattler placements inside their cage. Some algorithms, for example, based on energy minimization, might tend to push rattlers as far from their caging neighbors as possible, leading to less contacts per particle. For these reasons, we have analyzed the jammed backbones of the packings. As we will see in Chapter 11, we observe the same goodness of fit and exponent of 0.6 for $Z(x)$ for jammed sphere packings in dimension higher than 3, suggesting a possible geometrically universal feature underlying the observed divergence in the number of near contacts.

9.3.1.3 Away from Contact: Split Second Peak

Although the primary focus of this work is on the behavior of $g_2(r)$ around contact, it is instructive to also look at the split second peak of the pair correlation function, shown for a sample of packings of $10,000$ particles in Fig. 9.7. Only two clear discontinuities are seen, one at exactly $r = \sqrt{3}D$, and one at $r = 2D$. The latter is very clearly asymmetrical, with a sharp decrease in g_2 at $r = 2D^+$. Although the first discontinuity is less pronounced and statistics are not good enough to unambiguously determine its shape, it appears that it also has the same shape as the second discontinuity, only of smaller magnitude. The split second peak is of great importance because it is a clear signature of the strong local order in the first two coordination shells of the packing, and in fact observations have been made that along with the appearance of a peak in $P_f(f)$ for small forces, the splitting of the second peak of g_2 is a signature of jamming [169]. It is therefore important to try to understand the local geometrical patterns responsible for the occurrence of these structures in g_2. Analysis similar to ours is performed for packings of soft spheres near the jamming threshold in [212], and find similar results.

9.3.1.4 Contact-Network Statistics

The exact geometry of the jammed configuration $\mathbf{R}_J$ is determined (not necessarily uniquely) from its contact network, which as we have demonstrated is the network of first-neighbor interactions and can easily be separated from further-neighbor interactions. Fig. 9.8 shows the histogram of local coordination numbers as a function of the first-neighbor cutoff τ, i.e., the histogram of the number of particles within distance $(1 + \tau)D$ from a given particle. It is seen that for sufficiently small τ ($\tau < 10^{-5}$) the histograms are independent of the exact cutoff used (this is true down to $\tau \approx 100\delta$ or so, which can be as small as 10^{-12} in some of our packings). It is interesting to observe that the contact-number probability distributions conform very well to a Gaussian shape, at discrete points between four and ten, for all of the cutoffs shown in the figure. A number of particles having less than 2 contacts are seen, and these are clearly *rattlers* and we have removed them from consideration from all of the final packings we analyze here. We observe that such particles remain with fewer than 2 contacts for a very wide range of τ and are easy to identify. In some cases, however, we cannot unambiguously identify a handful of the particles as rattlers or non-rattlers. This is typical for packings which are not sufficiently close to their jamming point, packings which have been produced using fast expansion in the LS algorithm, or packings which are very large. It is safest to not remove such particles as rattlers.

This work is the first time a clear look has been provided at the exact contact network of disordered hard-sphere packings. Previous studies have either used soft atoms, in which case the definition of a contact is not clear cut unless one carefully takes limits of a stiff interaction potential [210], and therefore in such studies τ has been typically set to correspond to the location of the first minimum in $g_2(r)$, or have used Voronoi tessellations to define neighbors. Even studies which have actually used hard particles have resorted to such definitions unsuitable to investigating the jamming limit, mostly because the numerical precision required to separate the true contacts from the near contacts has not been achieved up to now [215]. Such investigations, the literature of which is too vast to cite, have found a plethora of local coordination patterns typical of *polytetrahedral* packing, including icosahedral order [215].

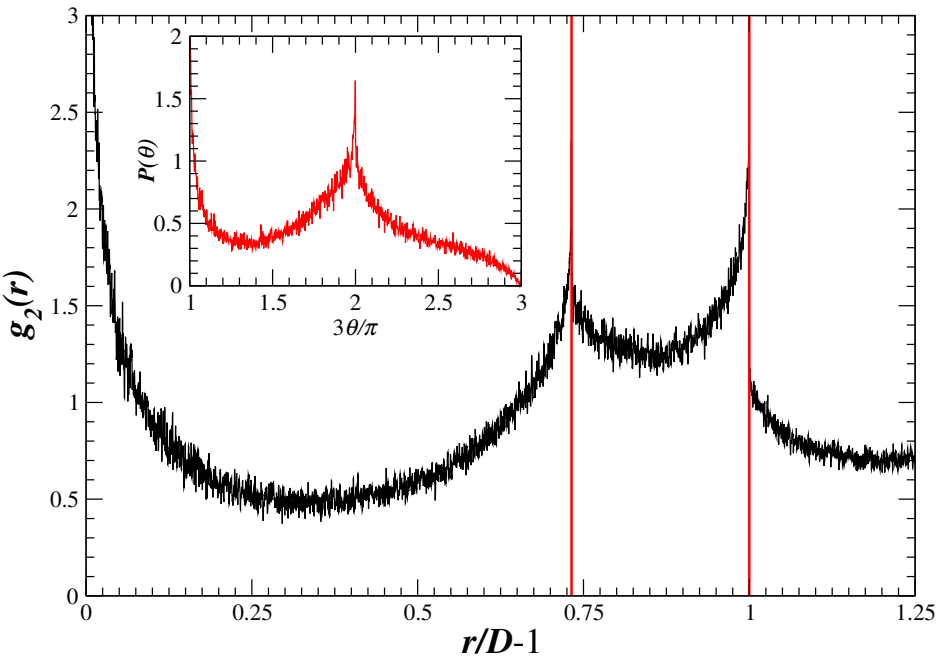

Figure 9.7: Computational data on the split second peak of $g_2(r)$ averaged over 5 packings of 10,000 particles. The values $r = \sqrt{3}D$ and $r = 2D$ are highlighted, and match the two observed discontinuities. Also visible is the divergence near contact. The inset shows the probability distribution $P_\theta(\theta)$ of bond-pair angles in the contact network of the packings, also revealing two divergences at $\theta = \pi/3$ and at $\theta = 2\pi/3$. No peaks are observed at $r = \sqrt{2}D$ or $r = \sqrt{5}D$, which are typical of crystal packings, indicating that there is no detectable crystal ordering in the packing.

We therefore attempted to do a similar *shared-neighbor* [215] analysis for the contact networks of our disordered packings, and look for local clusters reminiscent of polytetrahedral packing. Our procedure, based on looking at the contact network as an undirected graph, was as follows. For each particle, we extracted the subgraph corresponding to the first neighbor shell of the particle (this includes contacts between the neighbors), extracted its connected components, and counted the number of occurrences of a given subgraph (using graph algorithms that can test for graph isomorphism to form equivalence classes). The results were surprising. By far the most prominent patterns were a central particle contacting a *chain* of 1, 2, 3, 4 or 5 contacting particles. The chains were almost never closed, other than for chains of length 3 (which together with the original particle form a contacting tetrahedron), and this was itself rare. The probability of finding a chain of length n seems to decay exponentially, $P(n) \sim \exp(-1.2n)$. This study found very few tetrahedra, and so polytetrahedral local ordering is certainly not apparent in the contact networks. We also performed the same analysis for a range of τ's, all the way up to $\tau = 0.1D$ (which raises the average coordination significantly above 6), but still found the open linear chains to be the dominant pattern. We further attempted to include second neighbors in the analysis, however including all second neighbors led to very large subgraphs of a very broad variety, so classification was not possible. We further restricted our attention only to second neighbors which are very close to the given particle (within $0.1D$, for example), and this also found very few tetrahedra.

One of our goals was to determine if certain simple local coordination patterns are responsible for each of the three features of $g_2(r)$ we previously documented: the power-law divergence near contact, and the discontinuous, if not diverging peaks at $r = \sqrt{3}D$ and $r = 2D$. We had little success in accounting for the first one by restricting attention to only the first two neighbor shells in the true contact network. In particular, we looked at all the near contacts (for example, with gaps less than $0.01D$) and whether the almost contacting particles were in fact second neighbors in the contact network. Indeed most were, however the majority only shared one particle as a first neighbor, or two or three first neighbors which were not themselves

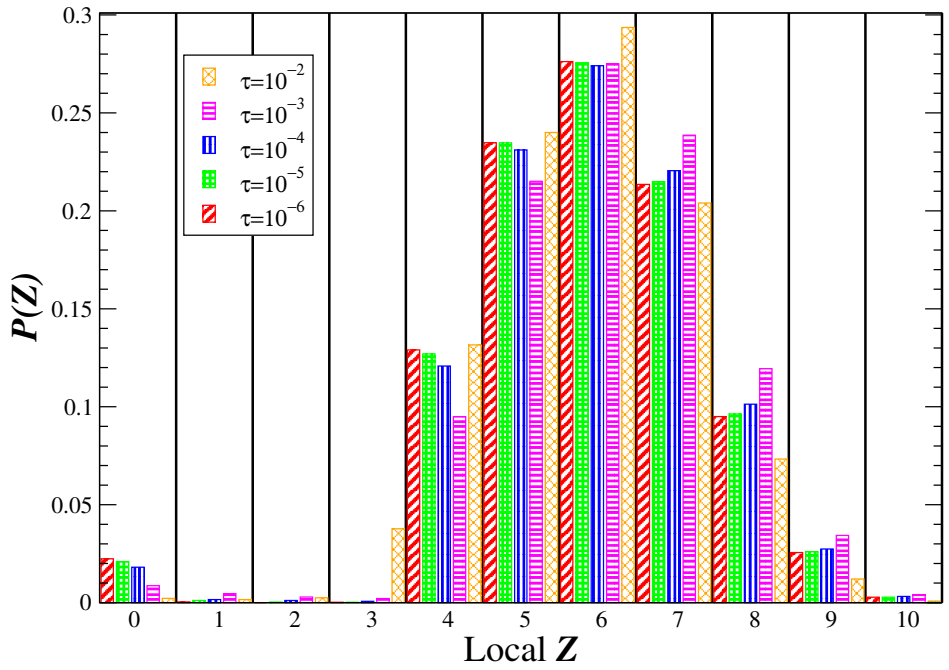

Figure 9.8: The probability distribution of local contact numbers as the cutoff used in defining neighbors is increased. Rattlers are clearly seen, and a relative maximum at $Z = 6$ is seen. Note that only one particle with 11 neighbors is observed, and very few have as many as 10 neighbors. No particle with 12 contacting neighbors has been observed in any of our packings, indicating a lack of crystallinity.

first neighbors. It was therefore not possible to isolate one particular local geometry as responsible for the multitude of near contacts. An interesting quantity we measured is the probability distribution $P_\theta(\theta)$ of bond-pair angles θ in the contact network, meaning the angles between two contact bonds of a given particle. This distribution is shown in the inset in Fig. 9.7, and shows divergences at $\theta = \pi/3$ and $\theta = 2\pi/3$, which correspond to distances $r = 2D\sin(\theta/2)$ of $r = D$ and $r = \sqrt{3}D$. Although there is no divergence at $\theta = \pi$, the corresponding distribution of distances does show a divergence at $r = 2D$.

We had more success with a shared-neighbors analysis for the split second peak. This was because we could increase τ and thus progressively relax the definition of first neighbor. We found that with increasing τ, an increasing majority of particle pairs at a distance close to $\sqrt{3}D$ were second neighbors, and that an increasing majority of them shared two neighbors which were themselves neighbors. This corresponds to two edge-sharing *approximately* equilateral coplanar triangles, a configuration which has been suggested as being responsible for the first part of the split second peak [20]. Note however that we do not observe any discontinuity in g_2 at $r = 1.633D$, which corresponds to two face-sharing tetrahedra, which is another configuration often mentioned in connection with the split second peak [20]. A similar analysis for the peak at $2D$ indicated that the majority of particle pairs at this distance share one neighbor between them, which represents an approximately linear chain of three particles, a configuration which has long been known to be responsible for the second part of the split second peak of g_2.

9.3.2 Ordered Packings

In this work we have focused on disordered hard-sphere packings, and have found a multitude of unexpected singular features, such as a long power-law tail in $g_2^{(\delta)}(l)$, a nonzero $P_f(f = 0)$, and a power-law divergence in $g_2^{(b)}(l)$. It is important to realize that the properties we observe are not universal and will change as one changes the amount of ordering of the packings. In particular, dense ordered packings like the FCC crystal are not isostatic, and we have no theory that can predict the shape of $g_2^{(\delta)}$. We therefore resort to a computational investigation of ordered and partially ordered sphere packings.

9.3.2.1 Vacancy-Diluted FCC Crystal Packings

It was the behavior of crystal packings around the jamming point that was the subject of Refs. [52] and [145], and these works inspired this investigation. For crystal packings, there is no ambiguity in defining first neighbors, and the FCC packing has $Z = 12$ contacts per particle, which is twice the isostatic value. Therefore, the limiting polytope $\mathcal{P}_{\mathbf{x}}$ is not a simplex and, as argued in Ref. [145], it is expected that for an FCC packing $g_2^{(\delta)}(l)$ will have a single peak for small gaps. We indeed observe this computationally, for the first time, as shown in Fig. 9.9.

Furthermore, we have prepared vacancy-diluted FCC packings by removing a fraction p of the spheres from a perfect crystal, $0 \leq p \leq 4$ (here $p = 0$ corresponds to the perfect crystal). The FCC lattice is composed of 4 interpenetrating cubic lattices. We obtain the vacancy-diluted crystal with the lowest density by removing one of these 4 cubic lattices (i.e., $p = 1/4$), as shown in the inset in Fig. 9.9. This gives a packing with density of about $\phi \approx 0.56$ and mean coordination $\bar{Z} = 8$ and is still collectively jammed. In fact, it is likely that more spheres can be removed with a more elaborate procedure [56]. We can add back a randomly chosen fraction $q = 1/4 - p$ of the previously removed quarter of the spheres, to obtain $0 < p < 1/4$. The delta-function contributions to g_2 for several p's are shown in Fig. 9.9. It is rather surprising that the pair correlation function for the $p = 1/4$ packing no longer shows a peak, but is monotonically decaying. In fact, by changing p one can obtain packings with $g_2^{(\delta)}(l)$ that has zero slope at the origin.

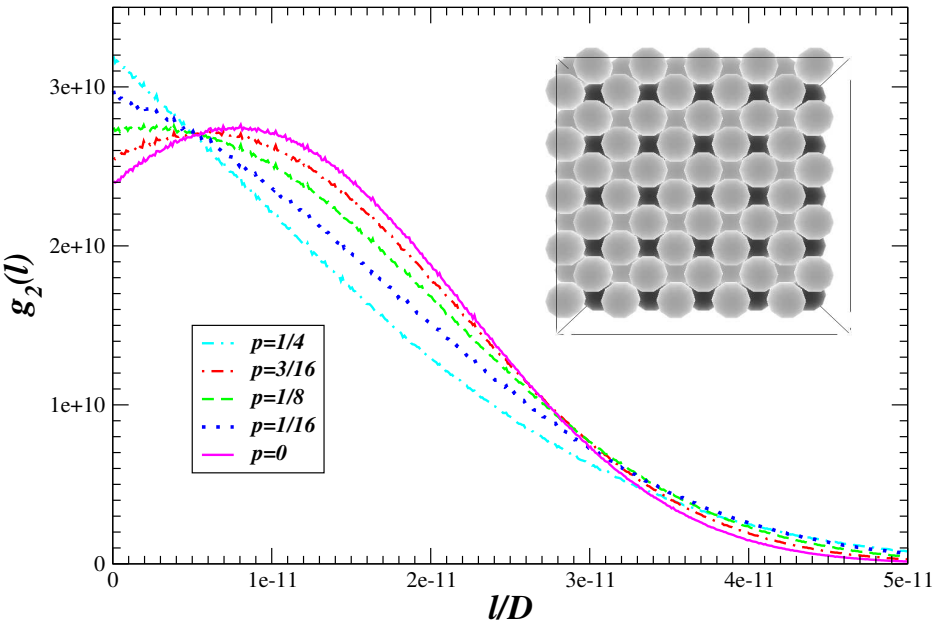

Figure 9.9: The first-shell $g_2^{(\delta)}(l)$ for a collection of FCC crystal packings with a fraction p of the spheres removed, starting with $N = 13,500$ particles. The inset shows the packing with most vacancies, where every 4^{th} sphere is removed to form a cubic sublattice of vacancies (colored dark). Intermediate p's are achieved by randomly adding back some of the spheres to the sublattice. The density has been reduced by $\delta = \sqrt{2} \cdot 10^{-11}$ from close packing.

It is interesting to note that for the (vacancy-diluted) FCC packings $g_2^{(\delta)}(l)$ decays in a Gaussian manner, and in fact is perfectly fitted by a modified Gaussian, $g_2^{(\delta)}(l) = (Al^2 + Bl + C) \exp\left[(l - D)^2\right]$, as suggested by an approximate theory [see Eq. (6) in Ref. [216]]. This fast decay is to be compared to the slow power-law decay for the disordered packings [c.f. Eq.(9.6)], hinting at possible connection to the stability of the crystal packings versus the metastability of the glass packings [217]. Additionally, we show the force distribution $P_f(f)$ for these ordered packings in Fig. 9.10, illustrating that, in contrast with the disordered packings, very small forces are not observed. It would be interesting to know if the perfect FCC crystal can be vacancy-diluted to an isostatic packing and still be collectively or strictly jammed, and what the corresponding force

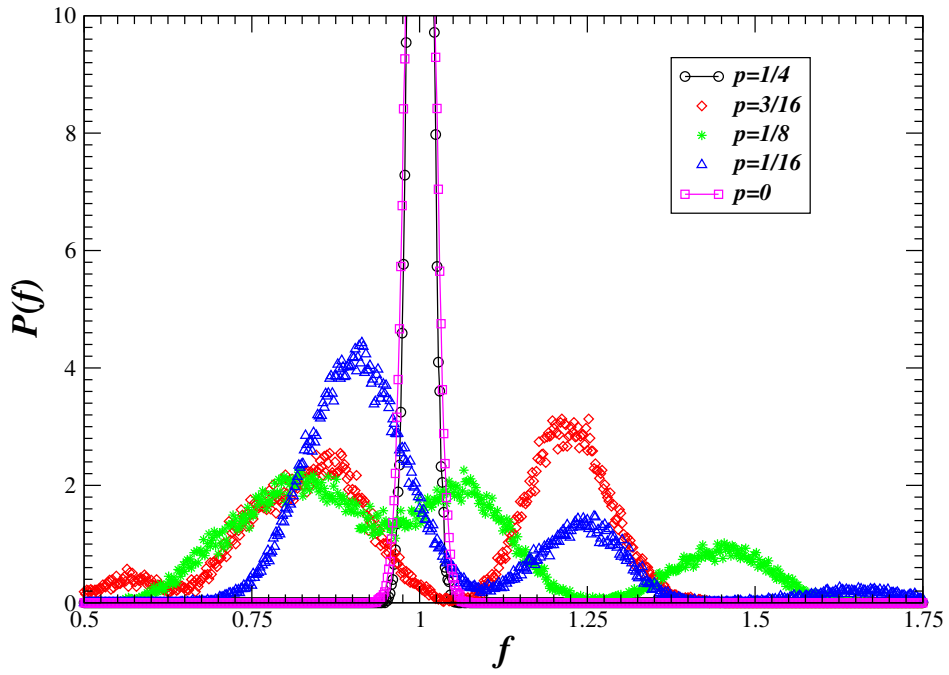

Figure 9.10: The force probability distribution for the collection of FCC crystal packings from Fig. 9.9. For the pure crystal and the crystal with most vacancies, all of the particle pairs are identical and therefore the probability distribution would be a delta function if forces are averaged over an infinite time horizon. For the intermediate p's, multiple relatively broad peaks are observed. In contrast with the disordered case, very small forces are *not* observed.

distribution would be.

9.3.3 Partially Crystallized Packings

As previously explained, the Lubachevsky-Stillinger algorithm can produce partially crystalline sphere packings if a sufficiently small expansion rate is used and nucleation of crystallites occurs during the densification. This is demonstrated in Fig. 9.11, where we show the evolution of the pressure during the densification of an initially liquid sample (i.e., a state on the stable equilibrium liquid branch), for a range of expansion rates γ. The slower the expansion is, the more crystalline the final packings become, as can be seen from the fact that the final density increases and from the evolution of the peaks in $g_2(r)$, as shown in Fig. 9.12. Additionally, the structure factor $S(\mathbf{k})$ shows more anisotropy and localized peaks. More detailed studies of crystallization using hard-sphere molecular dynamics have been performed by other researchers [218, 217]. Here we are merely interested in how crystallization affects the properties we have studied in detail for the disordered packings.

The packings shown in Fig. 9.11 clearly have nucleated crystals, and so one may anticipate that there is a qualitative distinction between them and the "random" packings produced by suppressing crystallization. However, as demonstrated in Fig. 9.13, slower densification leads to larger densities and more ordered packings even if crystallization is suppressed and no visible nucleation occurs. This indicates that there is a continuum of packings from most disordered to perfectly ordered [58] packings, so that one needs to be careful in interpreting results obtained from packings produced by just one, possibly non-trivially biased, algorithm. For example, Ref. [169] relates the occurrence of a peak in $P_f(f)$ to jamming. However, as we show next, jammed packings do not necessarily exhibit this peak if they are sufficiently ordered.

For the sake of brevity, we will only briefly discuss some interesting features of g_2 for the partially crystallized packings. Since the perfect FCC/HCP crystals have $\bar{Z} = 12$, one expects that, as partial crystallization occurs, somehow the number of first neighbors per particle should increase from the isostatic

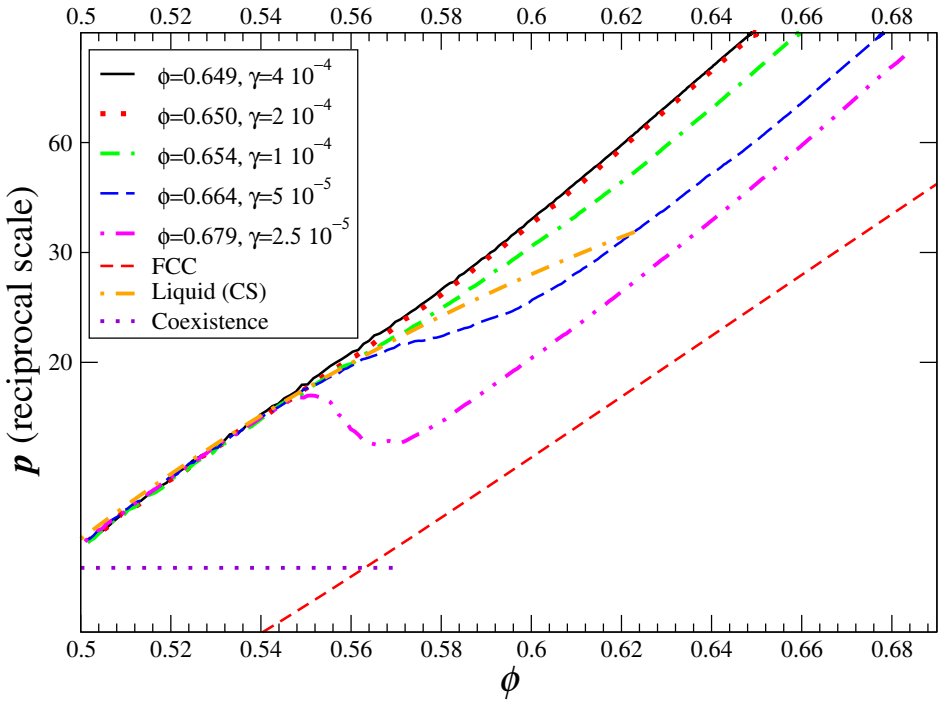

Figure 9.11: Compression of an initially liquid system with $\phi = 0.5$ to jamming with several different expansion rates γ (the mean thermal velocity is 1, in comparison). The pressure is plotted on a reciprocal scale (the tick marks being equally spaced in equal increments of p^{-1}, increasing in the usual direction), to highlight the expected linear relation (9.1) near jamming. The pressure-density curves for the perfect FCC crystal [54], the accepted fluid-solid coexistence region, and the widely known Carnahan-Starling equation of state for the fluid branch, are also shown for comparison. Sufficiently fast compression suppresses crystallization and leads to densities around $0.64 - 0.65$, and slower compression allows for partial crystallization, typically occurring around $\phi \approx 0.55$, which is the end of the coexistence region (i.e., the density where the crystal necessarily becomes thermodynamically favored). This produces denser packings which exhibit more crystal ordering the denser they are.

value of $\bar{Z} = 6$. However, this is not really so if one *properly* defines first neighbors via true contacts in the final jammed packing. In fact, if one plots $Z(l)$ for partially crystallized packings (we omit this plot), a qualitatively similar curve to that shown in Fig. 9.4 is seen, with $\bar{Z}$ clearly close to the isostatic value of 6. However, the background $Z^{(b)}(l)$ shows a faster rise the more crystalline the packing is [consistent with a larger exponent α as defined in Eq. (9.7)], so that indeed an increase of the cumulative coordination is seen for sufficiently large gaps. Additionally, we observe that nearly crystalline packings easily jam in noticeably hypostatic configurations, with a higher probability of observing particles with only 2 or 3 contacts and a less flat plateau in $Z(l)$.

All of these findings are readily explained. The basic premise, used widely in the granular media literature, is that random perturbations to either the particle-size distribution or to the boundary conditions will break some of the contacts in an otherwise perfect crystal down to the isostatic value. This is because additional contacts in excess of $\bar{Z} = 6$ imply special correlations between the positions of the particles, which one expects to destroy with random perturbations. Such random perturbations are provided in the case of partially crystallized packings by the fact that the crystallites need to jam against a partially amorphous surroundings, and this induces complex strains that break some of the perfect-crystal contacts.[6] However, the geometric peculiarities of the underlying crystal remain; for example, there is a multitude of nearly

[6]We mention in passing that we have observed similar results by starting with a perfect FCC crystal, applying a small (but not too small) random strain, and then jamming the packings. This typically yields almost perfectly crystal packings which are nonetheless clearly frustrated by the random strain to have $\bar{Z} \approx 6$.

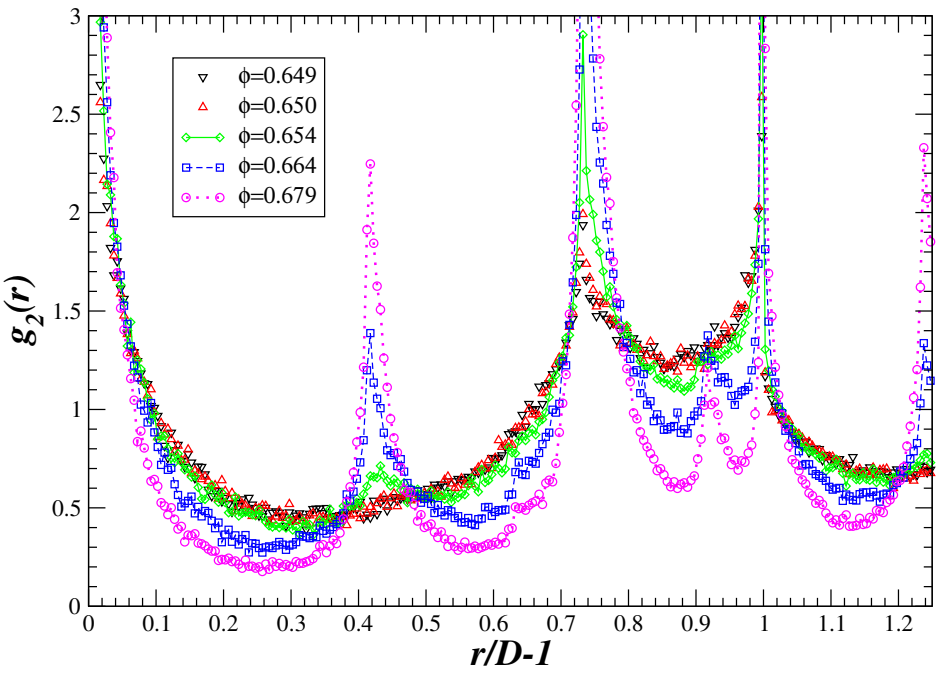

Figure 9.12: The evolution of the peaks in $g_2(r)$ as crystalline order is increased, for the packings from Fig. 9.11. The formation of peaks at distances typical of the FCC lattice, such as $r = \sqrt{2}$, is clearly seen. It is interesting to note that a peak is observed at $\sqrt{11/3} \approx 1.91$, which is a fifth-neighbor distance in the HCP (but not the FCC) lattice (a similar HCP peak at $\sqrt{8/3} \approx 1.63$ is barely visible). This is in agreement with numerous previous theoretical and numerical investigations of crystallization [218].

collinear (in fact lines of aligned particles) or coplanar contacts, which leads to the occurrence of much more pronounced *force chains* (chains of large forces propagating along a nearly straight line) and a sharp increase in the probability of occurrence of small forces. We indeed observe this in Fig. 9.14, where we show that for sufficiently ordered packings there is no longer a peak in $P_f(f)$ for small forces, but rather a monotonic decrease of $P_f(f)$, apparently exponential for sufficiently large forces. This is in contrast to previous studies of the effect of order on force distributions in granular piles [196, 219], which did not register a significant impact of the ordering. However, these studies examine the distribution of forces in granular piles and a direct comparison is beyond the scope of this work.

9.4 Conclusions

The results presented in this work settle some long-standing questions and confusions in the literature. For the first time, we showed both theoretically and computationally how the delta-function portion of $g_2(r)$ is formed as jamming is approached, for a true hard-sphere packing. Our investigation focused on maximally disordered (MRJ) sphere packings with a packing fraction $\phi \approx 0.64 - 0.65$. We presented the first true hard-sphere computational data on the power-law divergence in the near-contact portion of g_2, in agreement with previous observations in the literature for stiff soft spheres, but with a distinguishably different exponent of -0.4. We confirmed that this divergence persists even in the true jamming limit for hard particles. We presented high-quality data on the probability distribution of interparticle forces $P_f(f)$, especially focusing on small forces, demonstrating a maximum at small forces and a nonzero intercept at $f = 0$. A local analysis of the topology of the contact network found few traces of tetrahedra and an overwhelmingly complex local connectivity, and was successful in accounting for the structures responsible for the split second peak of $g_2(r)$. A computational study of the delta-function contribution to $g_2(r)$ for vacancy-diluted FCC crystals showed a faster than exponential decay, unlike the slow power-law decay for the disordered isostatic packings. Finally,

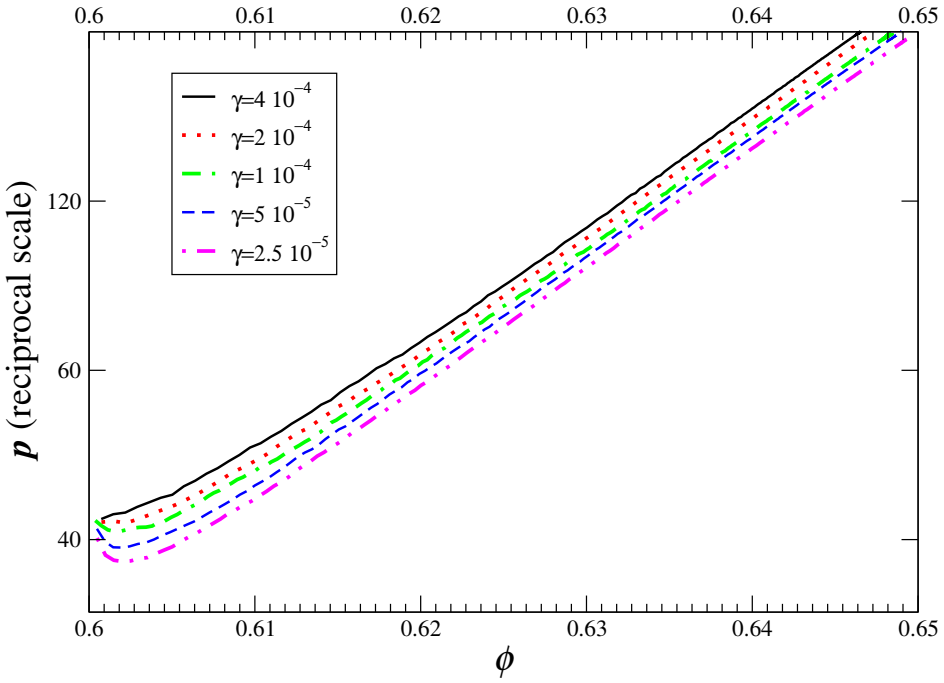

Figure 9.13: Compression of an initially (metastable) liquid system with $\phi = 0.6$ to jamming with several different expansion rates, as in Fig. 9.11. For this range of expansion rates, crystallization is suppressed due to the large initial density and all final packings are apparently disordered and would be ordinarily identified as random, however, it is clear that slower compression leads to higher densities, and thus the final packings are not all identical, but rather, some are more ordered than others, as can be verified by the slight increase in bond-orientational order metric Q_6 [58], for example.

we investigated packings on the transition from maximally disordered to maximally ordered, and found that partially crystallized packings produced by our algorithm are still nearly isostatic despite having a higher density, and that $P_f(f)$ loses the peak for sufficiently ordered packings.

This work has raised several important questions. The computational observations undermine the very applicability of the ideal jammed packing model to large (maximally) disordered packings of spheres, as produced by most algorithms in use today. First, a very unusual power law divergence in $g_2(l)$ is observed near contact, leading to a multitude of particle pairs just away from contact. Similarly, a power-law decay is seen in the contact part of $g_2(l)$. As the packings become larger, one can expect the tails of the two power laws to start overlapping by an observable number of contacts, blurring the distinction between true contacts and almost contacts. Even more troubling is the observation that there appears to be a positive probability of observing a zero force in the contact network of the packings, indicating the presence of geometric degeneracies in the contact network. The above observations may explain why we have had trouble generating truly jammed packings of $N = 10,000$ particles. However, we do not see a reason why very large but finite packings collectively jammed ideal packings could not be constructed. The question of what algorithm can produce disordered (and thus likely isostatic) packings which are jammed and devoid of some or all of the above peculiarities, as is the FCC crystal packing[7], for example, remains open.

[7]Note that the observations we list as troubling are separate from the rather general objections due to the inapplicability of the concept of ideal jamming to infinite packings, which apply to crystal packings as well [52].

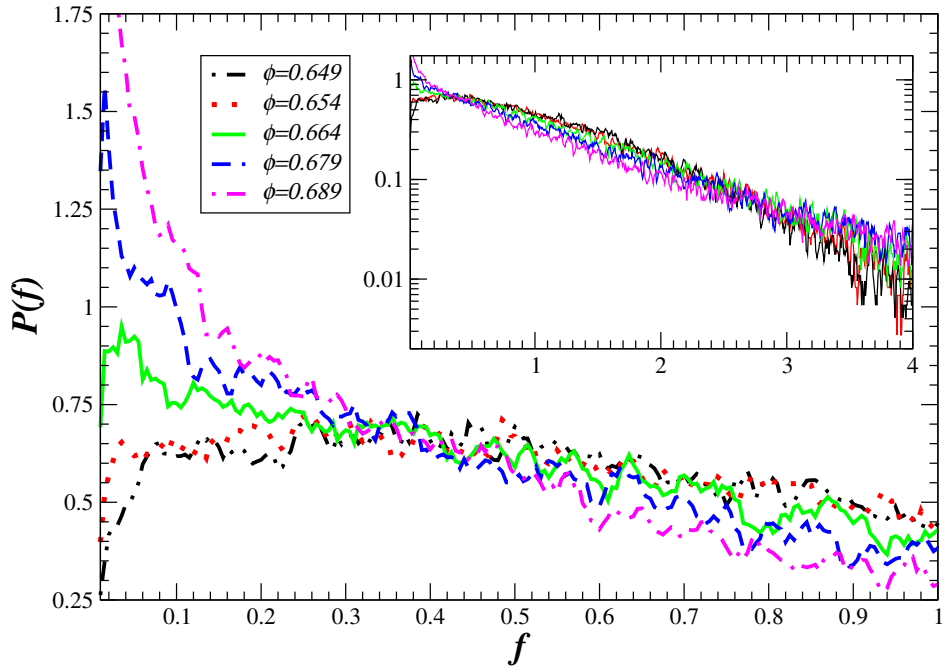

Figure 9.14: The evolution of $P_f(f)$ as crystalline order is increased, showing the disappearance of the peak at small forces. The inset shows a log-log view of the plot, and is consistent with exponential decay for large forces.

Chapter 10

Density Fluctuations in Jammed Sphere Packings

Continuing on previous theoretical investigations of local density fluctuations in atomic systems in Ref. [202], we computationally study jammed disordered hard-sphere packings of as many as one-million particles [76]. We show that local density fluctuations are suppressed from volume to surface ones, i.e., the generated packings are hyperuniform, with structure factor which vanishes at the origin to within 10^{-3}. The numerical results point to a strange non-analytic linear dependence of the structure factor around the origin, $S(k) \sim |k|$. In addition to exponentially damped oscillations seen in liquids, this implies a weak power-law tail in the total correlation function, $h(r) \sim -r^{-4}$. The direct pair correlation function shows a significant long-range tail outside the core, unlike the stable liquid. Finally, we demonstrate that the generated packings are saturated and all voids in the mechanically rigid backbone of the packing are filled with rattlers. Our results illuminate many open questions about density fluctuations in glassy atomic systems.

10.1 Introduction

The characterization of local density fluctuations in many-particle systems is a problem of great fundamental interest in the study of condensed matter, including atomic, molecular and granular materials. Previous work in Ref. [202] was concerned with the quantitative characterization of density fluctuations in point patterns, and in particular, those in which infinite wavelength fluctuations are completely suppressed, i.e., the structure factor $S(\mathbf{k})$ vanishes at the origin. In these so-called *hyperuniform* (or *superhomogeneous* [220]) systems, the variance in the number of points inside a large window grows slower than the volume of the window, typically like the window surface area. Most known examples of hyperuniform systems are either ordered lattices or quasi-crystals [202, 220]. An important open problem is the identification of statistically homogeneous and isotropic atomic systems (e.g., glasses) that are hyperuniform.

For equilibrium liquids and crystals $S(k = 0) = \kappa_T/\kappa_T^0$, where κ_T^0 is the ideal gas compressibility, and is thus positive. *Strictly jammed* sphere packings are rigorously incompressible (and non-shearable) [111]; however, they are also nonequilibrium systems. If the compressibility equation remains valid even in certain nonequilibrium cases, it would predict that for hard-sphere systems in the jamming limit $S(0) \approx 3/p^2$, where the reduced $p = PV/NkT$ diverges in the jamming limit according to Eq. (9.1). In Ref. [202] it was conjectured that all *saturated*[1] strictly jammed packings are hyperuniform. In this chapter we demonstrate that MRJ packings are indeed hyperuniform and saturated. Moreover, we observe an unusual non-analytic structure factor $S(k) \sim |k|$, or equivalently, a quasi-long ranged negative tail of the total pair correlation function $h(r) \sim -r^{-4}$.

We prepare jammed packings of hard spheres under periodic boundary conditions using a modified Lubachevsky-Stillinger (LS) packing algorithm [12], as detailed in Chapter 9. The generated disordered sphere packings typically have volume fractions in the range $\phi = 0.64 - 0.65$, and to a good approximation the packings should be representative of the maximally random jammed (MRJ) state [58]. The MRJ state

[1] A saturated packing is one in which no additional particles can be added.

for hard-particle packings is intimately related to jamming and dynamic arrest in both granular and glassy materials. For this study, we have generated a dozen packings of $N = 10^5$ and $N = 10^6$ particles jammed up to a reduced pressure of 10^{12} using an expansion rate of 10^{-3} [75] with $\phi \approx 0.644$. Generating such unprecedented one-million-particle packings was necessary in order to study large-scale density fluctuations without relying on dubious extrapolations.

10.2 MRJ Packings Are Saturated

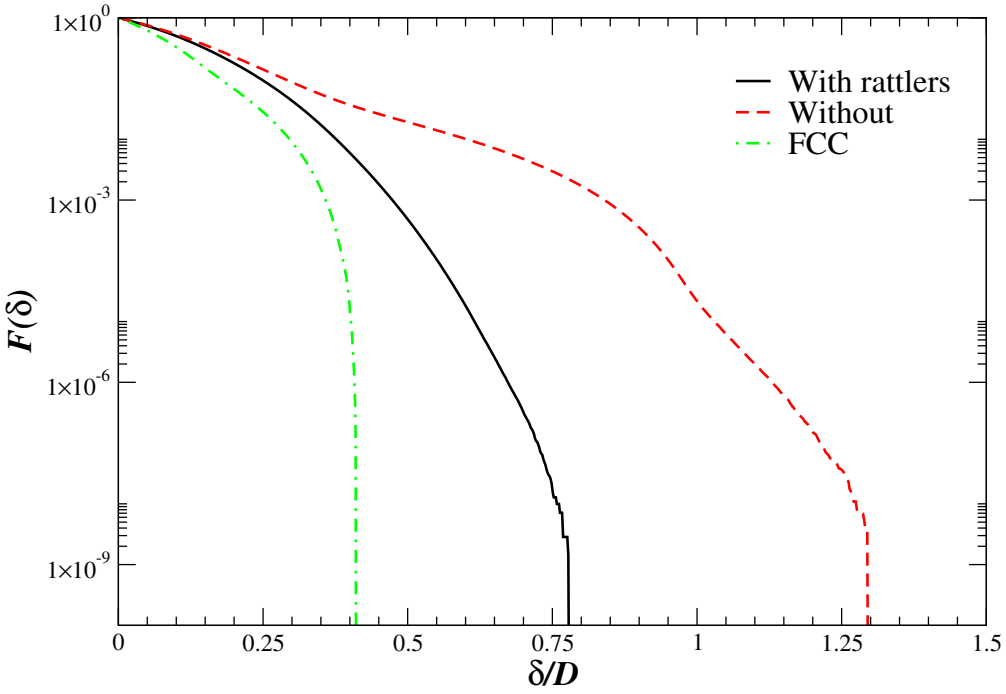

Figure 10.1: The pore-size function $F(\delta)$ for a packing with $N = 10^5$ particles, with and without the rattlers. The method of trial spheres with $2 \cdot 10^9$ trials was used [6]. A very similar $F(\delta)$ with cutoff around $\delta_{\max} \approx 0.8D$ is observed when $N = 10^6$, when rattlers are present. The cutoff is however not as sharply defined as it is for the FCC crystal, shown for comparison.

The packings generated by the LS and other algorithms have a significant fraction ($\sim 2.5\%$) of *rattling* particles which are not truly jammed but can rattle inside a small cage formed by their jammed neighbors [75]. These rattlers make a negligible contribution to the mechanical properties of the system, including the pressure, and can be removed sufficiently close to the jamming point. However, they are important when considering density fluctuations. Removing the rattlers will produce small but observable long-wavelength density fluctuations. Assuming that the rattlers are more or less randomly distributed among all particles, a hyperuniform packing from which the rattlers are then removed would have $S(0) \approx 0.025 > 0$. Similarly, the hyperuniformity could be destroyed by randomly filling large-enough voids with additional rattlers. It is therefore important to verify that the jammed packings are saturated, i.e., that there are no voids large enough to insert additional rattlers. Figure 10.1 shows the complementary cumulative pore-size distribution $F(\delta)$ [6], which gives the probability that a sphere of diameter δ could be inserted into the void space, with and without the rattlers. Clearly there is no room to insert any additional rattlers; the largest observed voids are around $\delta_{\max} \approx 0.8D$. The algorithm used to produce the packings appears to fill all void cages with particles.

10.3 Structure Factor Near $k = 0$

Long-wavelength density fluctuations in the packing are determined by the large-distance pair correlations, or equivalently, by the structure factor for small wavevectors. It is very difficult to study long-wavelength density fluctuations accurately in 3D computer simulations. When periodic boundary conditions apply with a periodic box of length L, particle correlations can only be studied up to a distance $L/2$, and there are large finite-size corrections for distances comparable to L. Additionally, as we show later, strong statistical fluctuations appear due to finite system size, making it necessary to use even larger systems to measure pair correlations at large distances. In reciprocal space, $S(k)$ can only be measured for $k \geq 2\pi/L$, with large discretization errors for the smallest wavevectors. To overcome these finite-size effects, it was necessary to generate a packing of one million particles.

Consider a large isotropic three-dimensional packing of N hard spheres of diameter D, with average number density $\rho = N/V$ and average volume fraction $\phi = \pi\rho D^3/6$. We employ the usual pair correlation function $g_2(x = r/D)$ or the total correlation function $h(x) = g_2(x) - 1$ in real space, or the equivalent Fourier representation given by the structure factor

$$S(K = kD) = 1 + 24\phi \int_0^\infty \frac{\sin(Kx)}{Kx} x^2 h(x) dx.$$

Of particular interest are the moments of $h(x)$, $\langle x^n \rangle = \int_0^\infty x^n h(x) dx$. Computer-generated packings are always finite, and thus binning techniques must be used to obtain probability densities like h. Accordingly, we prefer to use the more readily measurable *excess coordination*

$$\Delta Z(x) = 1 + 24\phi \int_0^x w^2 h(w) dw.$$

This is the average excess number of points inside a spherical window of radius xD centered at a particle, compared to the ideal-gas expectation $8\phi x^3$. Any integral containing $h(x)$ can easily be represented in terms of $\Delta Z(x)$ using integration by parts. For the structure factor we get $S(K) = \lim_{R\to\infty} S(K, R)$, where

$$S(K, R) = \Delta Z(R) \frac{\sin(KR)}{KR} - \int_0^R \Delta Z(x) \frac{d}{dx} \frac{\sin(Kx)}{Kx} dx. \tag{10.1}$$

This has quadratic behavior near $k = 0$ when expanded in a Taylor series,

$$S(K) \approx S(0) + \frac{K^2}{3} \int_0^\infty x \left[\Delta Z(x) - S(0)\right] dx, \tag{10.2}$$

where $S(0) = \Delta Z(x \to \infty)$ vanishes for a hyperuniform system. For large x, an *explicit* finite-size correction of order $1/N$ needs to be applied to the infinite-system excess coordination, $\Delta Z(x) \approx S(0)\left[1 - 8\phi x^3/N\right]$ [221], as it is clear that the excess coordination must vanish for windows as large as the whole system.

Figure 10.2 shows $S(k)$ for the simulated packings as obtained via a direct Fourier transform (DFT) of the particle positions,

$$S(\mathbf{k}) = N^{-1} \left| \sum_{j=1}^{N} \exp\left(i\mathbf{k}^T \mathbf{r}_j\right) \right|^2,$$

where $\mathbf{k}$ is a reciprocal lattice vector for the periodic unit cell. This calculation potentially involves many reciprocal lattice points and can only be done for a limited range of wavevectors near the origin due to memory and CPU constraints, and cannot easily and accurately be made faster using Fast Fourier Transforms (FFT). Recent non-uniform FFT algorithms can help obtain the DFT very efficiently [222, 223]. To obtain an approximation to the radially symmetric infinite-system $S(k)$, we average over the reciprocal lattice vectors inside a spherical shell of thickness $2\pi/L$. Using Eq. (10.1) together with a numerical (truncated) $\Delta Z(x)$ quickly gives $S(k)$ over a wide range of wavelengths. However, the behavior near the origin is not reliable since it depends on the analytic extension for the tail of $\Delta Z(x)$. The results of the DFT calculations are shown in Fig. 10.2, and they closely match the one obtained from $\Delta Z(x)$ for wavelengths smaller than about 20 diameters.

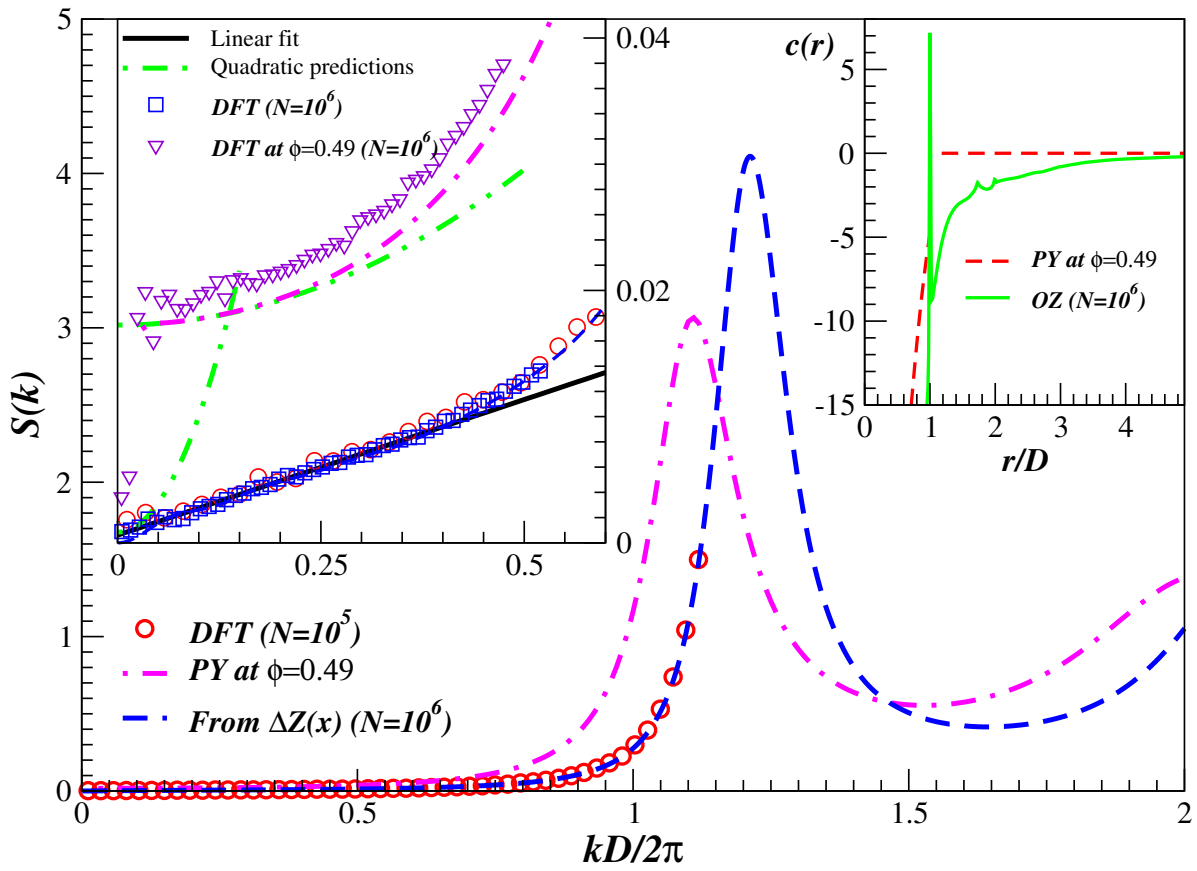

Figure 10.2: Structure factor for a 10^6-particle packing and for a hard-sphere liquid near the freezing point ($\phi = 0.49$), as obtained via two alternative numerical methods and also from the Percus-Yevick (PY) theory for the liquid [55]. DFT results are also shown over a larger range of K for a 10^5-particle packing. The left inset focuses on the range close to the origin, showing that while a parabola matches the liquid data reasonably well [$S(K) \approx 0.02 + 4 \cdot 10^{-3} K^2$ according to PY theory], it does not appear appropriate for the jammed packing for large-to-intermediate wavelengths. The very linear behavior of the DFT data for the jammed packing in the range up to $K/2\pi < 0.4$ is remarkable. The right inset shows $c(r)$ convolved (smudged) with a narrow Gaussian [due to numerical truncation of $S(k)$]. The peak at $r = D$ is thus in fact (almost) a δ-function.

Figure 10.2 demonstrates that the saturated packing is indeed hyperuniform[2] to within $S(0) < 10^{-3}$, as conjectured in Ref. [202]. The behavior of $S(k)$ near the origin is very surprising. The observed $S(k)$ follows closely a *non-analytic linear* relationship[3] well-fitted by $S(K) \approx 6.1 \cdot 10^{-4} + 3.4 \cdot 10^{-3} K$ over the whole range $K/2\pi < 0.4$. By contrast, analytic quadratic behavior is observed for a liquid sample at $\phi = 0.49$, as shown in the figure. Theoretical finite-size corrections to the small-k behavior of $S(k)$ have only been considered for relatively low-density liquid systems with relatively small N [221, 224, 225], and do not appear useful for our purposes. Although estimating the corrections to the DFT data analytically is certainly desirable, such corrections appear to be rather small at least for the well-understood liquid at $\phi = 0.49$. Comparison among the different $N = 10^6$ samples shows that statistical fluctuations in $S(k)$ near the origin are very small. It is important to again point out that it is essential to keep the rattlers when calculating the structure factor, since the removal of the rattlers introduces long-wavelength density fluctuations, as illustrated in Fig. 10.3.

[2]We believe that this small residual non-hyperuniformity is an artifact of the packing algorithm. For example, by reducing the particle expansion rate one observes an increase in the hyperuniformity (and also in density [75]).

[3]Since $S(k)$ is an even function, its derivative must vanish at the origin for it to be analytic.

The exact positioning of the rattlers inside their cages affects $S(k)$ and may also be contributing to the small but non-zero value we observe.

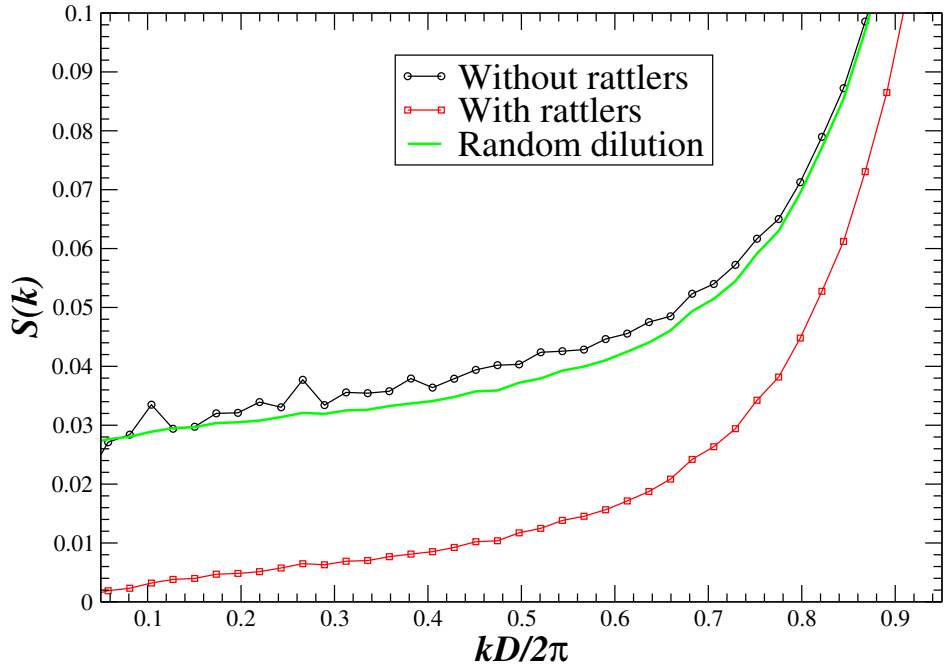

Figure 10.3: Structure factor obtained via DFT as in Fig. 10.2, with and without the rattlers, and with a random 2.5% of the particles removed from the packing. It is seen that removing the rattlers is equivalent to removing the same number of randomly-chosen particles, raising the structure factor to $S(0) \approx 0.025 > 0$.

Equation (10.2) shows that if h is truly short-ranged, the structure factor must be analytic (i.e., an even power of k, usually quadratic), near the origin. Our numerical observations point strongly to a linear $S(K)$ for small K. It is interesting to note that such non-analytic behavior is assumed in the so-called Harrison-Zeldovich power spectrum of the density fluctuations in the early Universe [220], and has been verified experimentally with high accuracy. If this observation $S(K) \sim |K|$ survives simulations of even larger jammed hard sphere systems, using a variety of packing algorithms, it would imply a negative algebraic power-law tail $h(x) \sim -x^{-4}$ uncharacteristic of liquid states. Such a quasi-long range correlated decrease in the density around a particle is typically only seen in systems with long-range interactions.

10.3.1 Direct Correlation Function

A long-ranged tail must appear in the *direct* correlation function $c(r)$ for a strictly hyperuniform system due to the divergence of $\tilde{c}(0)$, in a kind of "inverted critical phenomenon" [202]. Diverging length scales have been identified at a jamming transition for systems of soft spheres [188]. Such a tail is uncharacteristic of liquids where the range of $c(r)$ is substantially limited to the range of the interaction potential. The direct correlation function can numerically be obtained from its Fourier transform via the Ornstein-Zernike (OZ) equation,

$$\tilde{c}(k) = (\pi/6\phi)\,\frac{S(k)-1}{S(k)},$$

and we have shown it in the inset in Fig. 10.2, along with the corresponding PY anzatz for $c(r)$ at $\phi = 0.49$ which is assumed to vanish outside the core. Two unusual features relative to the liquid are observed for our jammed packing. First, there is a positive δ-function at contact corresponding to the $\bar{Z} = 6$ average touching neighbours around each jammed particle [75]. Second, there is a relatively long tail outside the core, the

exact form of which depends on the behavior of $S(k)$ around the origin[4].

10.4 Number Fluctuations

The numerical coefficient in the power-law tail in $h(x)$ suggested by the linear behavior of the structure factor near the origin is very small and cannot be directly observed, as we will show shortly. It is however possible to observe its effect on large-scale density fluctuations. For monodisperse hard sphere systems it suffices to focus only on the positions of the sphere centers and consider density fluctuations in *point patterns*. Following Ref. [202], consider moving a spherical window of radius $R = XD$ through a point pattern and recording the number of points inside the window $N(X)$. The number variance is exactly [202],

$$\begin{aligned} \sigma^2(X) &= \left\langle N^2(X) \right\rangle - \left\langle N(X) \right\rangle^2 \\ &= \frac{3\phi}{2} \left[(2X)^2 \Delta Z_0(2X) - \Delta Z_2(2X) \right] \end{aligned}$$

where $\Delta Z_n(x) = \int_0^x w^n \Delta Z(w) dw$ denotes a running moment of ΔZ.

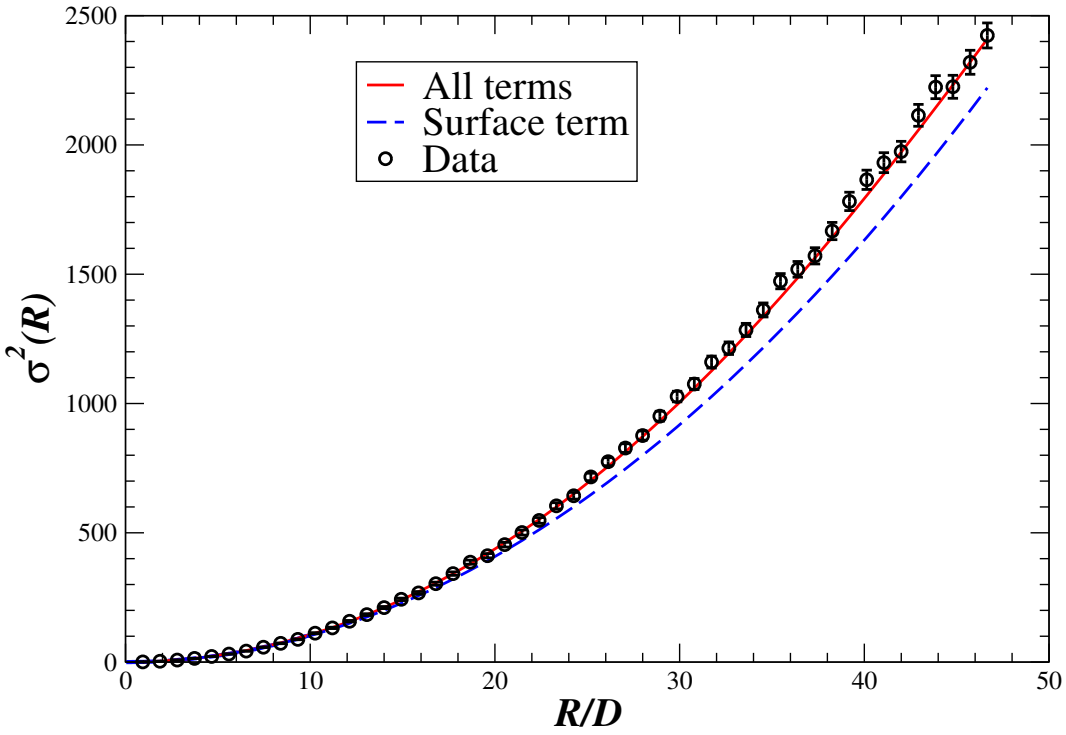

Figure 10.4: The variance σ^2 as a function of the window radius for a 10^6-particle packing. The uncertainty in the variance, as shown with error bars, is estimated to be of the order of $\sigma^2/\sqrt{M}$, where $M = 10^4$ is the number of windows used for a given window. Also shown is the theoretically predicted dependence of the form $AX^3 + CX^2 \ln X + B_0 X^2$, along with just the surface term $B_0 X^2$, which dominates the density fluctuations.

Asymptotically, for large windows, in an infinite system with analytic $S(k)$,

$$\sigma^2(X) \approx AX^3 + BX^2,$$

where

$$A = 8\phi \left(1 + 24\phi \left\langle x^2 \right\rangle\right) = 8\phi S(0)$$

[4]We used the linear fit to $S(k)$ when producing the figure, implying $c(r) \sim r^{-2}$.

is the *volume* fluctuation coefficient, and

$$B = -144\phi^2 \left\langle x^3 \right\rangle = 6\phi \Delta Z_0 \, (x \to \infty)$$

is the *surface* fluctuation coefficient. When a non-integrable power-law tail exists in $\Delta Z(x)$, asymptotically the "surface" fluctuation coefficient contains an additional logarithmic term in the moment $\left\langle x^3 \right\rangle$, giving a correction $B(X) = B_0 + C \ln X$ and a variance

$$\sigma^2(X) \approx AX^3 + CX^2 \ln X + B_0 X^2.$$

Such a logarithmic correction does not appear for any of the examples studied in Ref. [202]. Explicit finite-size effects for non-hyperuniform systems yield an additional correction $A(X) = 8\phi S(0) \left[1 - 8\phi X^3/N\right]$ [226]. Figure 10.4 shows numerical results for the number variance as a function of window size, along with the predicted asymptotic dependence, including both the logarithmic and N^{-1} corrections[5]. Both corrections need to be included in order to observe this close a match between the data and theory. The constants $S(0)$ and C were obtained from the linear fit to $S(k)$, while $B_0 \approx 1.02$ was obtained by numerically integrating $\Delta Z(x)$, as explained shortly[6].

10.5 Tail of the Total Correlation Function

We now turn our attention to real space to observe directly the large-distance behavior of h or equivalently ΔZ. For equilibrium liquids with short-ranged potentials it is expected that the asymptotic behavior of $h(x)$ is exponentially damped oscillatory [228, 229], of the form

$$h(x) \sim \frac{C}{x} \exp\left(-x/\xi\right) \cos\left[K_0(x - x_0)\right]. \tag{10.3}$$

However, it is not clear whether the decay is still exponential for glass-like nonequilibrium jammed systems. Previous studies have looked at much smaller systems, where explicit finite-size effects dominate, and also focused on the liquid phase [224, 225]. Figure 10.5 shows the numerical $\Delta Z(x)$ along with a relatively good exponentially damped oscillatory fit

$$\Delta Z(x) \approx 5.47x \exp(-x/1.83) \cos(7.56x - 2.86)$$

over the range $5 < x < 15$. Compare this fit to $\xi \approx 1.4$ and $K_0 \sim 7.9$ reported in Ref. [36]. It would be desirable to look at larger x and, in particular, directly observe the long-range inverse power tail predicted from the linear behavior of $S(k)$.

The use of cubic periodic boundary conditions imply that only pair distances up to $x_{\max} = \sqrt[3]{\pi N/24\phi}$ can be studied, and additionally one expects to see implicit finite-size effects due to the use of periodic boundaries in addition to explicit finite size effects. However, the most important for our study is the uncertainty $\delta Z(x)$ coming from statistical variations among finite systems. An estimate of the magnitude of the uncertainty $\delta Z(x)$ can be derived as follows. Ignoring the fact that the windows used in obtaining $\Delta Z(x)$ are centered on the particles, the expected uncertainty in the number of points per window is $\sigma(x)$. Since we are averaging over N windows, the expected uncertainty is $\delta Z(x) \approx \sigma(x)/\sqrt{N}$. Our numerical experience has demonstrated that this is a good estimate despite the assumptions in its derivation. If N_s different packings were averaged over, for a hyperuniform system, the uncertainty is of the order of

$$\delta Z(x) = x\sqrt{\frac{B}{N \cdot N_s}}.$$

The maximal uncertainty is thus of the order of $\delta Z(x_{\max}) \sim 1/\sqrt[6]{N}$, which decreases rather slowly with increasing system size, and in practice the range of x for which sufficient accuracy can be obtained is rather

[5] Additional *implicit* finite size effects due to the periodicity of the system have been considered for hard disks in Ref. [227], and they have been shown to be significantly smaller. While finite-size effects for hyperuniform systems have not been studied theoretically, Fig. 10.4 suggests they are also small.

[6] The surface coefficient B_0 cannot be determined from just the linear part of $S(k)$ near the origin.

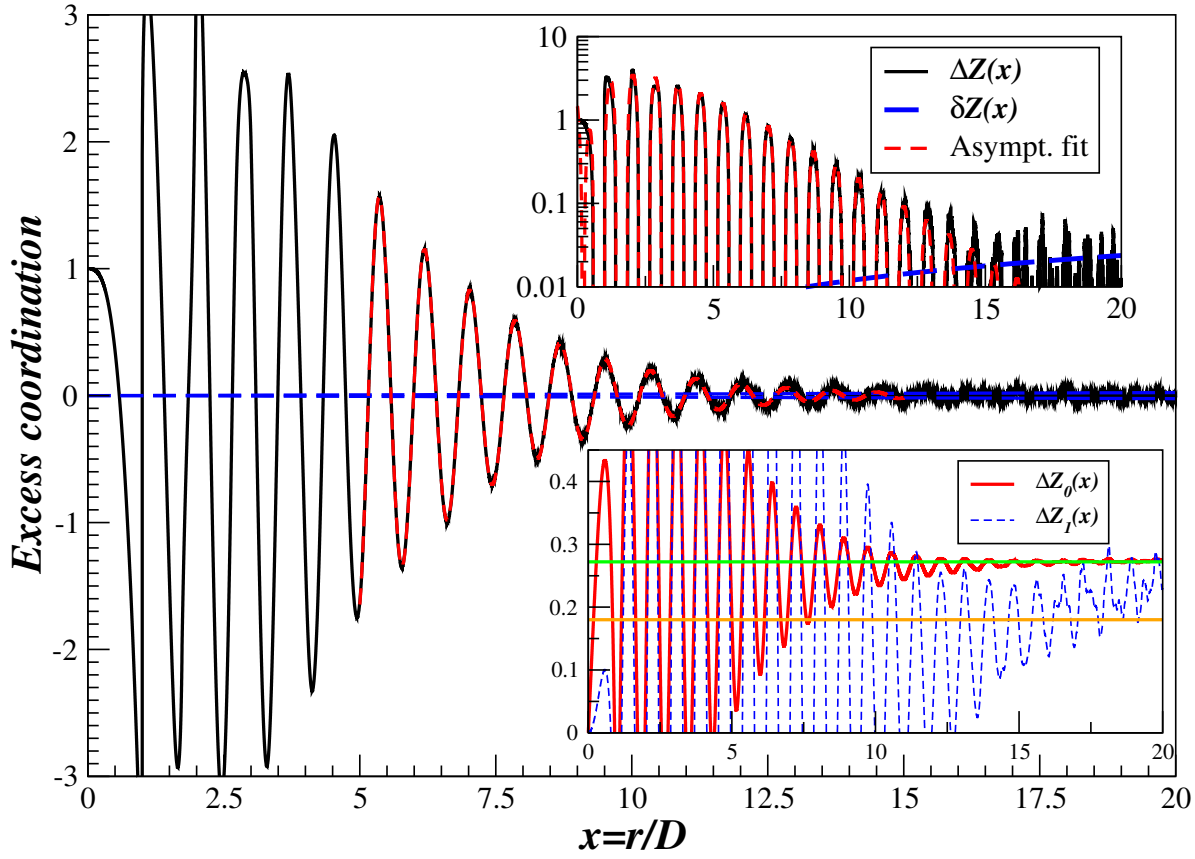

Figure 10.5: The excess coordination for a 10^6-particle packing, along with the best fit of the form (10.3) for the tail, and the estimated uncertainty. Statistical fluctuations overcome the actual correlations after $x \approx 15$. Averaging over nine samples only shrinks the magnitude of the fluctuations by three, without revealing additional information. The inset on the top uses a logarithmic scale, and the inset on the bottom shows the zeroth and first running moments along with their asymptotic values as estimated from the tail fit. Note that for the range of x shown explicit finite-size corrections are small (less than 5%).

limited. We have systematically observed that $\Delta Z(x)$ exhibits unusually strong *correlated* fluctuations at the tail of the range of certainty, most likely due to amplification of the underlying small geometrical correlations by statistical noise. These strong oscillations in $\Delta Z(x)$ for large x are illustrated in Fig. 10.6. It is therefore essential to restrict the observations to the region of validity as determined from $\delta Z(x)$, and as seen in Fig. 10.5 this limits the range of validity of the numerical $\Delta Z(x)$ to at most $x < 15$.

Within the range of validity of the observed $\Delta Z(x)$ the damped oscillatory fit is perfectly appropriate. We smoothly combined the actual numerical data for $x < 10$ with the fitted decaying tail for $x > 10$, and numerical integration of this smoothed $\Delta Z(x)$ gives $B_0 \approx 1.02 \pm 0.02$, as used in producing Fig. 10.4. This smoothed $\Delta Z(x)$ was used to obtain $S(k)$ via Eq. (10.1) when producing Fig. 10.2. The estimated power-law tail due to the non-analytic behavior of $S(k)$, $\Delta Z(x) \approx 4.4 \cdot 10^{-3} x^{-1}$, is too small to be observed directly, as it is completely overwhelmed by the statistical fluctuations in $\Delta Z(x)$.

10.6 Conclusions

We have given computational results for million-particle jammed disordered hard sphere packings and demonstrated that they are saturated and hyperuniform. We found that $S(k)$ is nonanalytic at the origin in striking contrast to liquid behavior. There are many open fascinating questions. Can a geometrical significance be

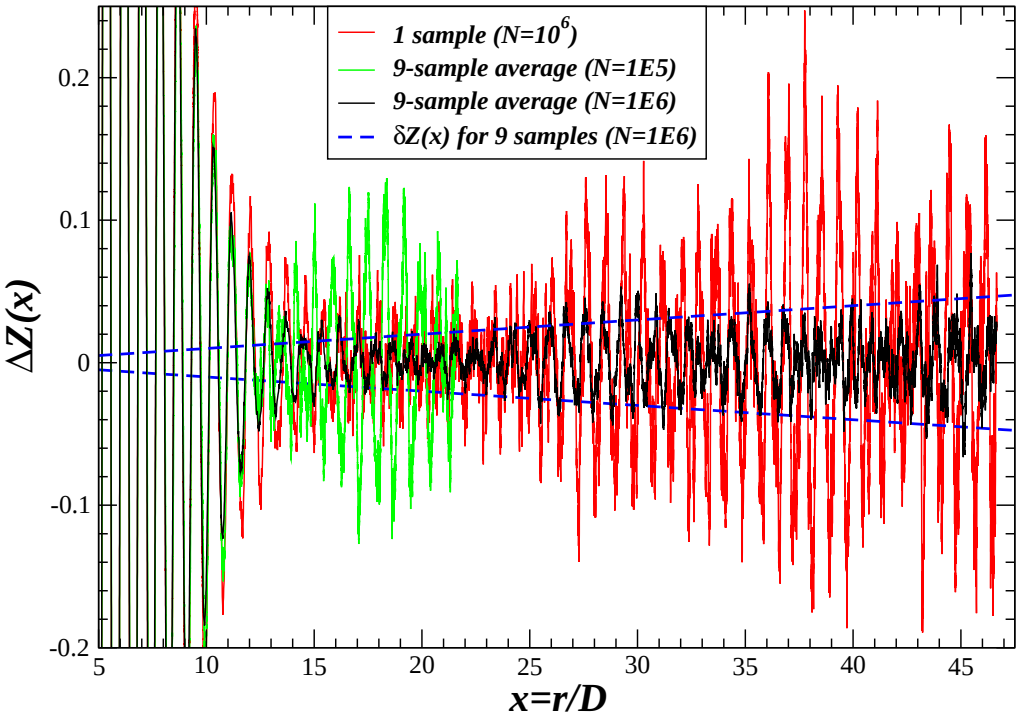

Figure 10.6: The excess coordination as in Fig. 10.5 but shown over the full range of distances up to half the simulation box length. Statistical fluctuations overcome the actual correlations after $x \approx 15$ and they are correlated over short distances, leading to large oscillations that may falsely be perceived as a signature of persistent long-range correlations.

attached to the period of oscillations K_0 in the jamming limit, or to the cutoff of $F(\delta)$? We believe that the strict jamming and saturation conditions demand hyperuniformity of our packings. We argue that the observed nonanalytic behavior of $S(k) \sim k$ is a direct consequence of the condition of maximal disorder on the jammed packing. The exponent p appears to increase with increasing order: It approaches infinity for ordered lattices, is two for perturbed lattices, and is one for MRJ. In this sense, the MRJ packings are markedly more disordered: they have macroscopic density fluctuations which are much larger than crystalline packings. Quantitative understanding of this aspect of disorder and its relation to density fluctuations remains a fascinating open problem.

Chapter 11

Packing Hyperspheres in High Dimensions

In this chapter we present the first study of disordered jammed hard-sphere packings in four-, five- and six-dimensional Euclidean spaces [77]. Using a modified Lubachevsky-Stillinger packing generation algorithm, we obtain the first estimates for the packing fractions of the maximally-random jammed (MRJ) states for space dimensions $d = 4$, 5 and 6 to be $\phi_{MRJ} \simeq 0.46$, 0.31 and 0.20, respectively. Calculations of the pair-correlation function $g_2(r)$ and structure factor $S(k)$ for these states show that short-range ordering appreciably decreases with increasing dimension, consistent with a recently proposed "decorrelation principle" stating that unconstrained correlations diminish as the dimension increases and vanish entirely in the limit $d \to \infty$. As in three dimensions (where $\phi_{MRJ} \simeq 0.64$), the packings show no signs of crystallization, are isostatic, and have a power-law divergence in $g_2(r)$ at contact with power-law exponent $\simeq 0.4$. Across dimensions, the cumulative number of neighbors equals the kissing number of the conjectured densest packing close to where $g_2(r)$ has its first minimum. Additionally, we obtain estimates for the freezing and melting packing fractions for the equilibrium hard-sphere fluid-solid transition, $\phi_F \simeq 0.32$ and $\phi_M \simeq 0.39$, respectively, for $d = 4$, and $\phi_F \simeq 0.19$ and $\phi_M \simeq 0.24$, respectively, for $d = 5$. Although our results indicate the stable phase at high density is a crystalline solid, nucleation appears to be strongly suppressed with increasing dimension.

11.1 Introduction

In this chapter we extend the studies of hard sphere systems in three dimensions presented in previous chapters, and in particular Chapters 9 and 10, to four, five and six dimensions. Of particular interest to us are (nonequilibrium) disordered jammed packings of hard spheres and their statistical and mechanical properties, such as the maximally random jammed (MRJ) state [58, 56], pair correlations [75], isostaticity [75], and density fluctuations [76].

A hard-sphere packing in d-dimensional Euclidean space $\mathcal{R}^d$ is an arrangement of congruent spheres, no two of which overlap. As in a variety of interacting many-body systems [230], we expect studies of hard-sphere packings in high dimensions to yield great insight into the corresponding phenomena in lower dimensions. Analytical investigations of hard-spheres can be readily extended into arbitrary spatial dimension [231, 232, 233, 234, 235, 236, 237, 238, 239, 6, 202, 240, 241] and high dimensions can therefore be used as a stringent testing ground for such theories. Mean-field-like models that are only approximate in lower dimensions become more accurate in higher dimensions, enabling exact theoretical calculations in the limit $d \to \infty$. Along these lines and of particular interest to this paper, predictions have been made about long-wavelength density fluctuations [202] and decorrelation [240, 241] in disordered hard-sphere packings in high dimensions. Additionally, the optimal packing of hard spheres in high dimensions is also of interest in error-correcting codes in communications theory [9].

Our focus in this paper will be the study of hard-sphere packings in four, five and six dimensions. We consider equilibrium hard-sphere systems for $d = 4$ and $d = 5$. We report estimates of the freezing, melting, and kinetic glass transition densities, and find a decreased tendency to crystallize with increasing dimension.

We also numerically study MRJ packings of hard spheres for $d = 4$, 5 and 6 that are at least collectively jammed and report the first estimates of the packing fractions of the MRJ states [58] in these dimensions to be $\phi_{MRJ} \simeq 0.46$, 0.31 and 0.20, respectively. We find that short-range ordering exhibited by $g_2(r)$ and $S(k)$ appreciably diminishes with increasing dimension, consistent with a recently proposed "decorrelation principle" stating that unconstrained spatial correlations diminish as the dimension increases and vanish entirely in the limit $d \to \infty$ [240].

11.1.1 Previous Work

Equilibrium thermodynamic properties of hard-sphere packings for $d = 4$ and $d = 5$ have been studied both computationally and with approximate theories. For the low-density fluid, lower-order virial coefficients, B_2, B_3, and B_4, are known exactly for arbitrary dimensionality [231, 232]. Higher-order virial coefficients have been calculated by Monte Carlo simulation, B_5, B_6 and B_7 for both $d = 4$ and $d = 5$ and B_8 for $d = 4$ [234], or analytically [233]. The pair correlation function for equilibrium fluids has been studied and a decrease in ordering with increasing dimension was readily apparent [242]. Hard-sphere systems have been shown to undergo a (first-order) fluid-solid phase transition by numerical simulations for $3 \leq d \leq 5$ [243] and with approximate theories for d as high as 50 [238]. The freezing points for $d = 4$ and $d = 5$ were estimated numerically to occur at packing fractions (fraction of space covered by spheres) $\phi_F \sim 0.5\phi_{max} = 0.31$ and $\phi_F \sim 0.4\phi_{max} = 0.19$, respectively, and it was conjectured that freezing occurs at lower packing fractions relative to close packing as the dimension increases [243].

At sufficiently large densities, the packing of spheres with the highest jamming density has the greatest entropy because the free-volume entropy dominates over the degeneracy entropy. Therefore, the high-density equilibrium phase corresponds to the optimal packing, i.e., maximal density. The densest packing for $d = 3$ was recently proved by Hales [39] to be attained by the FCC lattice with packing fraction $\phi_{max} = \pi/\sqrt{18} = 0.7404\ldots$. The kissing number Z, the number of spheres in contact with any given sphere, for the FCC lattice corresponds to the maximal kissing number $Z_{max} = 12$ for $d = 3$. One of the generalizations of the FCC lattice to higher dimensions is the D_d checkerboard lattice, defined by taking a cubic lattice and placing spheres on every site at which the sum of the lattice indices is even (i.e., every other site). The densest packing for $d = 4$ is conjectured to be the D_4 lattice, with packing fraction $\phi_{max} = \pi^2/16 = 0.6168\ldots$ and kissing number $Z = Z_{max} = 24$, which is also the maximal kissing number in $d = 4$ [9]. For $d = 5$, the densest packing is conjectured to be the D_5 lattice, with packing fraction $\phi_{max} = 2\pi^2/(30\sqrt{2}) = 0.4652\ldots$ and kissing number $Z = 40$ [9]. For $d = 6$, the densest packing is conjectured to be the "root" lattice E_6, with density $\phi_{max} = 3\pi^3/(144\sqrt{3}) = 0.3729\ldots$ and kissing number $Z = 72$ [9]. The maximal kissing numbers Z_{max} for $d = 5$ and $d = 6$ are not known, but have the following bounds: $40 \leq Z_{max} \leq 46$ for $d = 5$ and $72 \leq Z_{max} \leq 82$ for $d = 6$ [9]. In very high dimensions, it has been suggested that random packings of spheres might have a higher density than ordered packings, enabling the intriguing possibility of disordered ground states and hence thermodynamic glass transitions [240]; see also Ref. [241].

11.1.2 Simulation Procedure

We use event-driven molecular dynamics (MD) and a modified Lubachevsky-Stillinger (LS) algorithm [12], as in Chapter 9 (see also Section 1.2), to produce collectively-jammed hard-sphere packings. As discussed in Chapter 3, our algorithm uses periodic boundary conditions applied to a hypercubic cell, in which a fundamental cell containing N spheres is periodically replicated to fill all of Euclidean space. We also use the cell method, in which the computational domain is divided into cubic cells and only neighboring cells are checked when predicting collisions for a given sphere. Since the number of neighboring cells, as well as the number of spheres per cell, increases considerably with increasing dimension, working in high dimensions is computationally intensive. Additionally, eliminating excessive boundary effects requires on the order of ten sphere diameters per simulation box length, i.e., on the order of $N = 10^d$ spheres. Due to the increasing computational load with increasing dimension, we cannot yet study $d > 6$. Implementing the near-neighbor list (NNL) techniques from Section 3.3.2, as well as parallelization, are necessary in order to study higher dimensions. Dimension-independent C++ codes used to generate the data in this paper can be downloaded at `http://cherrypit.princeton.edu/Packing/C++`.

11.2 Thermodynamic Properties

The temperature in equilibrium systems of hard spheres is a trivial variable; it does not affect equilibrium configurational correlations, leaving only one independent thermodynamic state variable, which can be taken to be either the reduced pressure $p = PV/Nk_BT$ or the density ϕ, related through the equation of state (EOS). Hard-sphere systems undergo a (first-order) fluid-solid phase transition, characterized by a melting point, which is the density at which the crystal thermodynamically begins to melt, and a freezing point, which is the density at which the fluid thermodynamically begins to freeze. Equilibrium properties, such as the melting and freezing points, are studied here using small expansion rates ($\gamma = 10^{-5} - 10^{-9}$) and periodic rescaling of the average sphere velocity to one, such that the total change in kinetic energy of the system, due to the collisions between growing spheres, was kept small. In this section we only consider four and five dimensions due to (presently) prohibitive computational costs for higher dimensions. Section 1.2 and in particular Fig. 1.3 contains results closely related to the results presented here for three dimensions.

11.2.1 Equilibrium Properties

The low density phase for hard spheres is an isotropic fluid (which is formally a gas, but at high densities has the microstructure of a liquid) in all dimensions. It is expected that this fluid will undergo a first-order phase transition into a non-isotropic solid phase with long-range translational order (i.e., a crystal), at least in lower dimensions [243, 238]. In this section we use MD to estimate the density ranges for the equilibrium fluid, solid, and coexistence regions, in four and five dimensions.

11.2.1.1 Freezing Point

Figure 11.1 shows the reduced pressure p as a function of density ϕ for (a) simulations of $d = 4$ systems of spheres placed in a D_4 lattice with negative expansion rate $\gamma = -10^{-6}$ and (b) simulations of $d = 5$ systems of spheres placed in a D_5 lattice with negative expansion rate $\gamma = -10^{-5}$. The pressure initially follows the (lower) crystal branch, until the system becomes mechanically unstable and jumps onto the (higher) fluid branch. Also plotted is the theoretical prediction of Luban and Michels (LM) for the equation of state [235], which agrees well with our numerical results for the fluid branch for $d = 4$, but less so for $d = 5$. It is a computational observation that crystals become mechanically unstable, giving rise to a sudden jump in pressure, at a density close to the freezing point [244, 245]. Such "superheating" (undercompression) is most likely due to the difficulty of achieving coexistence in finite systems, although we are not aware of a theoretical analysis. From the results in Fig. 11.1, we estimate the freezing points for $d = 4$ and $d = 5$ to be $\phi_F \simeq 0.31 - 0.32$ and $\phi_F \simeq 0.19 - 0.20$, respectively.

11.2.1.2 Melting Point

The melting points for $d = 4$ and $d = 5$ can also be estimated from the data in Fig. 11.1. Since throughout the coexistence region the fluid and solid have the same absolute pressure P, the melting density can be estimated as the density on the crystal branch with the same absolute pressure P as that at the freezing point. The coexistence region is plotted in Fig. 11.1 and the melting packing fractions for $d = 4$ and $d = 5$ are estimated to be $\phi_M \simeq 0.38 - 0.40$ and $\phi_M \simeq 0.24 - 0.26$, respectively. We also observe that the reduced pressure at the freezing point is $p_F \simeq 12$ in both $d = 4$ and $d = 5$, which agrees with the reduced pressure at the freezing point for $d = 3$, $p_F \simeq 12.3$, obtained from free energy calculations [116].

The melting point was also estimated for $d = 4$ (higher dimensions are presently too computationally demanding) by slowly densifying a system of spheres, initially a fluid, and looking for the onset of partial crystallization, again by monitoring the reduced pressure p as a function of density ϕ. Due to the difficulty of observing coexistence in finite systems and the relatively high activation barrier, simulated hard-sphere systems become "supercooled" (over-compressed) and nucleation does not occur until the melting density is surpassed. Consequently, the density at which partial crystallization appears for sufficiently slow expansion provides a reasonable estimate for the melting density, as illustrated for $d = 3$ in Fig. 1.3. As in Fig. 1.3, we plot the estimated jamming density $\tilde{\phi}_J(\phi)$ [c.f. Eq. (1.3)] instead of the pressure $p(\phi)$, as shown in Fig. 11.2 for a system of 648 spheres in $d = 4$. The onset of partial crystallization causes a dramatic jump in $\tilde{\phi}_J(\phi)$, as the jamming density of the crystal is much higher than the jamming density of a disordered packing. The

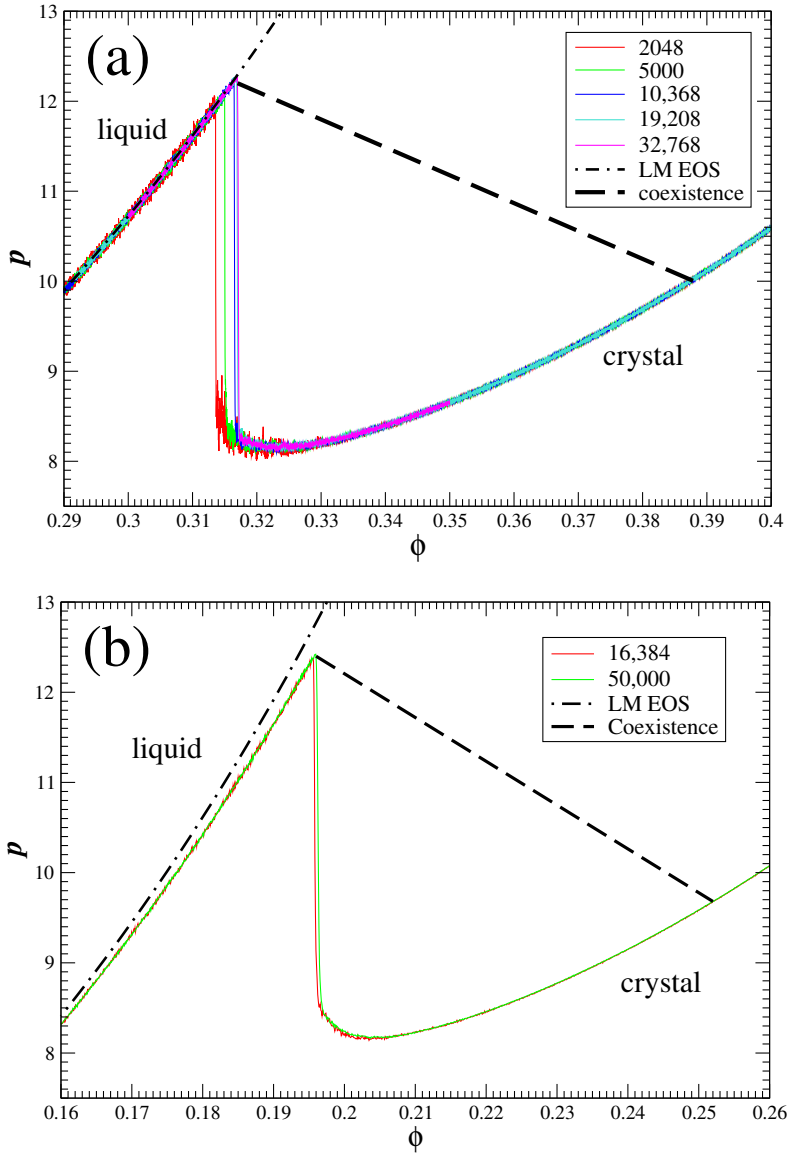

Figure 11.1: Reduced pressure p as a function of density ϕ, for a range of system sizes (see legend), for (a) $d = 4$ systems of spheres, initially in a D_4 lattice, and negative expansion rate $\gamma = -10^{-6}$ and (b) $d = 5$ systems of spheres, initially in a D_5 lattice, and negative expansion rate $\gamma = -10^{-5}$. The number of particles N was chosen to make a perfect D_d lattice with periodic boundary conditions, i.e., $N = (2n)^d/2$ for $n\epsilon\mathbb{Z}$. Also plotted is the theoretical prediction of Luban and Michels (LM) for the equation of state [235].

intersection of the curves with the line $\tilde{\phi}_J(\phi) = \phi$ gives the final jamming density. Sufficiently fast expansion suppresses crystallization and leads to packing fractions around $0.45 - 0.47$. Slower expansion allows for partial crystallization, typically around $\phi_M \simeq 0.38 - 0.39$, which is our rough estimate of the melting point, in agreement with our estimate from the results in Fig. 11.1. More accurate estimates can only be obtained using free-energy calculations. Since crystallization is a nucleated process, it is not surprising that the same expansion rates γ can crystallize at different packing fractions and onto different crystal branches, for example, $\gamma = 10^{-8}$ (a) and (b) in Fig. 11.2.

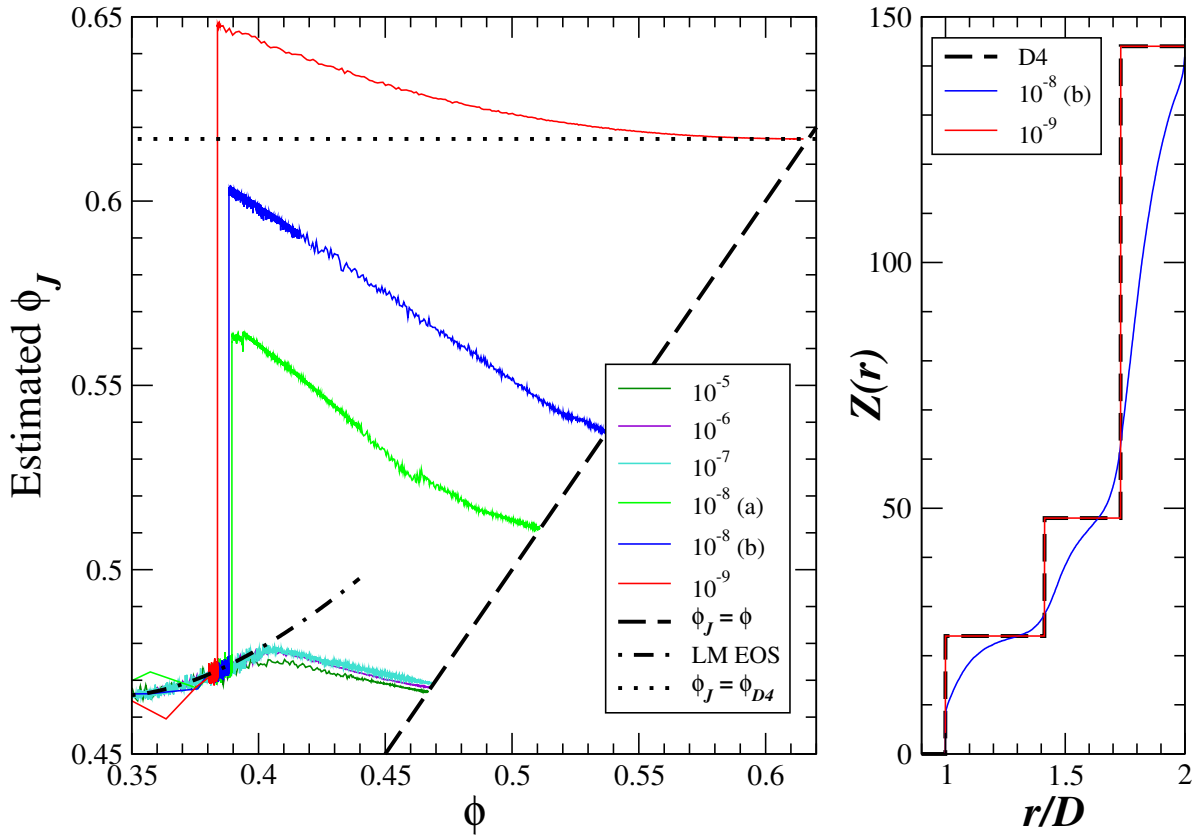

Figure 11.2: (*Left*): The estimated jamming packing fraction $\tilde{\phi}_J$ as a function of density ϕ for systems of 648 spheres for $d = 4$ with various expansion rates (see legend and note that there are two samples labeled (a) and (b) for $\gamma = 10^{-8}$). (*Right*): The cumulative coordination $Z(r)$ (i.e., the number of contacts) for the perfect D_4 lattice and for the partially crystallized packings at $p > 10^{12}$ obtained for expansion rates $\gamma = 10^{-8}$ and $\gamma = 10^{-9}$. The jamming packing fraction for the $\gamma = 10^{-8}$ packing is $\phi = 0.511$, and the jamming packing fraction for the $\gamma = 10^{-9}$ packing agreed up to 12 significant figures with the density of the D_4 lattice, $\phi = \pi^2/16 \simeq 0.617$.

To determine whether the crystallized packings were spontaneously forming a D_4 lattice, the conjectured densest packing in four dimensions, we computed the average cumulative coordination number $Z(r)$, which is the average number of sphere centers within a distance r from a given sphere center. The inset to Fig. 11.2 shows $Z(r)$ for a perfect D_4 lattice and for the crystallized packings with $\gamma = 10^{-8}$ and $\gamma = 10^{-9}$ (corresponding colors represent the same packing). The sharp plateaus for the D_4 lattice correspond to the coordination shells and the number of spheres in the first shell is the kissing number $Z_{max} = 24$. The packing with $\gamma = 10^{-9}$ formed a perfect D_4 lattice. The packing with $\gamma = 10^{-8}$ partially crystallized with a final density of $\phi \simeq 0.511$.

11.2.2 Kinetic Glass Transition

Figure 11.3 shows the estimated jamming packing fraction $\tilde{\phi}_J(\phi)$, as in Fig. 11.2, but for a system of $10,000$ spheres, instead of 648 spheres, in four dimensions. In contrast to the 648 sphere system, there is no sign of partial crystallization for the $10,000$ sphere system. In fact, molecular dynamics was performed at packing fractions of $\phi \simeq 0.38 - 0.42$ for 10 million collisions per sphere and there was no significant drop in pressure indicative of partial crystallization. The curves in Figs. 11.2 and 11.3 exhibit a bump around $\phi_G \simeq 0.41$, suggesting a kinetic transition from the fluid branch to a glassy branch.

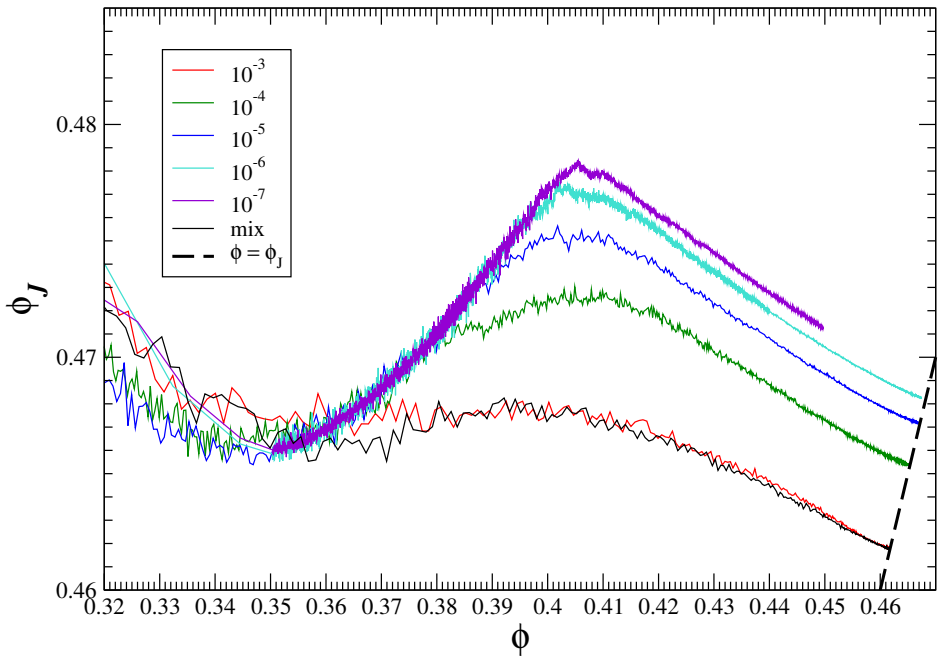

Figure 11.3: The estimated jamming packing fraction $\tilde{\phi}_J$ as a function of density ϕ for a system of $10,000$ spheres for $d = 4$ with various expansion rates. The curve labeled "mix" corresponds to the following sequence of expansion rates: $\gamma = 10^{-2}$ until $p = 10$, $\gamma = 10^{-3}$ until $p = 10^4$, $\gamma = 10^{-4}$ until $p = 10^6$, and $\gamma = 10^{-5}$ until $p = 10^{12}$.

Figure 1.3 showed the estimated jamming packing fraction $\tilde{\phi}_J$ for systems of spheres for $d = 3$ with various positive and negative expansion rates, for comparison with the results for $d = 4$ and $d = 5$ in Figs. 11.1, 11.2 and 11.3. The locations of the freezing and melting points in $d = 3$ have been determined from free-energy calculations [116] and good approximations to the EOS for both the fluid and crystal phases [54] are known. Our estimates of the freezing and melting points as the densities at the onset of melting of a diluted crystal or of partial crystallization of a densified fluid, respectively, compare favorably to the true values computed from free-energy calculations in $d = 3$. The bump around $\phi_G \simeq 0.59$, analogous to the bump in Fig. 11.3 around $\phi_G \simeq 0.41$, is often cited as the approximate location of the "kinetic" glass transition [70]. Comparing Figs. 11.3 and 1.3 reveals that the melting point and suggested kinetic glass transition are closer for $d = 4$ than for $d = 3$, which is a possible reason why there is a lower tendency to crystallize for $d = 4$ than for $d = 3$. Similar results have been observed for binary hard disks, a model glass former [57], as discussed in Chapter 13.

11.3 Disordered Jammed Packings

Packings representative of the maximally random jammed (MRJ) state are produced by a combination of expansion rates. The expansion rate must be initially high (compared to the average thermal velocity) to suppress crystallization and produce disordered configurations that are trapped in the neighborhood of a

jammed packing. As discussed in more detail in Chapter 9, near the jamming point the expansion rate must be sufficiently slow to allow for particle readjustments necessary for collective jamming. Figure 11.3 shows the final jamming packing fractions of packings created using a variety of expansion rates, as the packing fraction at which the curves intersect the line $\tilde{\phi}_J = \phi$. We see that by increasing the expansion rate, we attain packings with lower jamming packing fractions.

By comparing Fig. 11.3 and to the analogous plot for a $d = 3$ system (Fig. 1.3), where it is widely accepted that $\phi_{MRJ} \simeq 0.64 - 0.65$ [58, 56], we estimate the MRJ density for $d = 4$ to be $\phi_{MRJ} \simeq 0.460 \pm 0.005$. A more accurate calculation of ϕ_{MRJ} demands a better theoretical understanding of order metrics and how the expansion rate in the algorithm affects the ordering in the produced packings; statistical errors are smaller than the effect of the packing-generation protocol. Systematic investigation of different protocol parameters, as done for $d = 4$ in Fig. 11.3, is currently too computationally intensive in higher dimensions. Reasonable estimates of ϕ_{MRJ} for both $d = 5$ and $d = 6$ are obtained using the following less computationally intensive procedure. First, the system of spheres is expanded, starting from zero initial kinetic energy ($T = 0$), until it reached a high pressure (say, $p = 100 - 1000$). Then the system is slowly expanded ($\gamma = 10^{-5} - 10^{-3}$) and periodically cooled to $k_B T = 1$ until a very high pressure (say, $p = 10^{12}$) is attained. The resulting packings are approximately collectively jammed, as demonstrated by very large relaxation times for the pressure during long molecular dynamics runs [75]. Using this method we estimate the MRJ density for $d = 5$ to be $\phi_{MRJ} \simeq 0.310 \pm 0.005$ and for $d = 6$ to be $\phi_{MRJ} \simeq 0.200 \pm 0.01$.

The MRJ densities as well as important equilibrium densities are summarized in Table 11.1. It is useful to compare the MRJ densities for $d = 3 - 6$ to recent estimates of the *saturation* density ϕ_s for the random sequential addition (RSA) packing of hard spheres obtained in Ref. [246]. These authors found that $\phi_s = 0.379 \pm 0.0015, 0.2501 \pm 0.0005, 0.1514 \pm 0.0001$ and 0.0896 ± 0.0006 for $d = 3, 4, 5$ and 6, respectively. The nonequilibrium RSA packing is produced by randomly, irreversibly, and sequentially placing nonoverlapping spheres into a volume. As the process continues, it becomes more difficult to find available regions into which the spheres can be added. Eventually, in the saturation (infinite-time) limit no further additions are possible, and the maximal achievable density is the saturation density ϕ_s [see Ref. [6] and references therein]. As expected, the saturation density in dimension d is substantially smaller than the corresponding MRJ density because, unlike the latter packing, the particles cannot rearrange.

Packing fraction	$d = 3$	$d = 4$	$d = 5$	$d = 6$
ϕ_F	0.494 [6, 116]	$0.32 \pm 0.01^*$	$0.19 \pm 0.01^*$	-
ϕ_M	0.545 [6, 116]	$0.39 \pm 0.01^*$	$0.24 \pm 0.01^*$	-
ϕ_{MRJ}	0.645 ± 0.005 [56]	$0.46 \pm 0.005^*$	$0.31 \pm 0.005^*$	$0.20 \pm 0.01^*$
ϕ_{max}	0.7405 . . . [39]	0.6169 . . . [9]	0.4652 . . . [9]	0.3729 . . . [9]

Table 11.1: Important packing fractions for $d = 3, 4, 5$ and 6. The freezing and melting points for $d = 6$ were not calculated here. *Values computed in this work.

Our estimates for the MRJ density are compared to a theoretical formula proposed by Philipse [239] for the "random jamming density" ϕ_d,

$$\phi_d \simeq \frac{0.046d^2 + 1.22d + 0.73}{2^d}, \tag{11.1}$$

which predicts $\phi_3 \simeq 0.601$, $\phi_4 \simeq 0.397$, $\phi_5 \simeq 0.249$, and $\phi_6 \simeq 0.152$. It is seen that Eq. (11.1) underestimates ϕ_{MRJ} in $d = 3$ and becomes worse with increasing dimension. Following Ref. [246], we obtain a better scaling form by noting that the product $2^d \phi_{MRJ}$ for $3 \leq d \leq 6$ is well approximated by a function linear, rather than quadratic, in d (see Fig. 11.4), i.e., the scaling form for ϕ_{MRJ} is given by

$$\phi_s = \frac{c_1}{2^d} + \frac{c_2 d}{2^d}, \tag{11.2}$$

where $c_1 = -2.72$ and $c_2 = 2.56$. Although the scaling form (11.2) applies only in low dimensions such that $d \geq 3$, theoretical arguments given in Ref. [246] suggest that the general scaling form (11.2) persists in the high-dimensional limit, albeit with different coefficients c_1 and c_2.

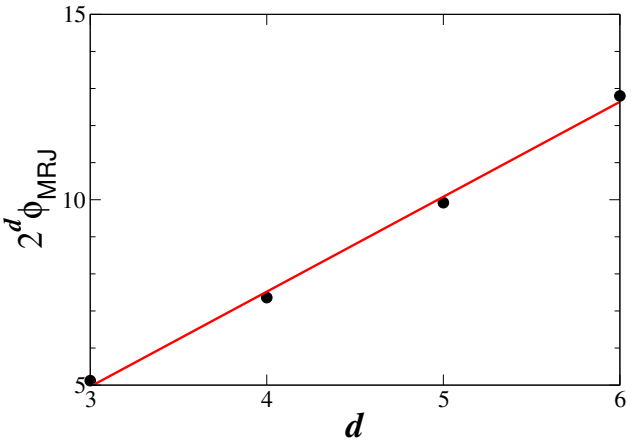

Figure 11.4: Fit of the data for the product $2^d\phi_{MRJ}$ to the linear form (11.2) for $3 \leq d \leq 6$ with $c_1 = -2.72$ and $c_2 = 2.56$.

11.3.1 Pair Correlations

Our main interest is pair correlations in the jamming limit in four, five and six dimensions. We characterize jammed packings statistically using the pair correlation function $g_2(r)$ and structure factor $S(k)$. The pair correlation function measures the probability of finding a sphere center at a given distance from the center of another sphere, normalized by the average number density ρ to go asymptotically to unity at large r,

$$g_2(r) = \frac{\langle P(r) \rangle}{\rho s_1(r)}, \tag{11.3}$$

where $P(r)$ is the probability density for finding a sphere center a distance r from an arbitrary sphere center, $\langle\rangle$ denotes an ensemble average, and $s_1(r)$ is the surface area of a single hypersphere of radius r [6]: $s_1(r) = 2\pi^2 r^3$ in $d = 4$, $s_1(r) = 8\pi^2 r^4/3$ in $d = 5$ and $s_1(r) = \pi^3 r^5$ in $d = 6$. The structure factor

$$S(k) = 1 + \rho\hat{h}(k) \tag{11.4}$$

is related to the Fourier transform of the total correlation function $h(r) = g_2(r) - 1$. It measures spatial correlations at wavenumber k and in particular, large-scale density fluctuations at $k = 0$ [202]. The structure factor can be observed directly via scattering experiments [230].

11.3.1.1 Pair Correlation Function

As discussed in detail in Chapter 9, in the jamming limit, the pair correlation function $g_2(r)$ consists of a δ-function due to sphere contacts and a background part $g_2^b(r)$ due to spheres not in contact:

$$g_2(r) = \frac{\bar{Z}\delta(r - D)}{\rho s_1(D)} + g_2^b(r), \tag{11.5}$$

where $\bar{Z}$ is the average kissing number. Figure 11.5 compares the pair correlation function for jammed packings of $100,000$ spheres in $d = 3$, 4, 5 and 6. Due to periodic boundary conditions, $g_2(r)$ can only be calculated up to half the length of the simulation box, which limits the calculation to $r/D \simeq 3$ for $d = 6$. The well-known split second peak present in $d = 3$ is strongly diminished as the dimension increases, i.e., the amplitude of the split second peak decreases and the sharp cusps become rounded with increasing dimension. The split third peak present in $d = 3$ with considerable structure and two shoulders vanishes almost completely in the higher dimensions. The oscillations are strongly damped with increasing dimension and the period of oscillations might also decrease slightly with increasing dimension; this latter possibility is revealed more vividly in the structure factor through the shift in the location of the maximum, as we will

describe below. The inset to Fig. 11.5 shows the magnitude of the decaying oscillations in $h(r)$ on a semi-log scale. Though at the values of r/D shown, up to about half the length of the simulation box, there is still structure in addition to the oscillations, especially apparent for $d = 3$, it appears that the decay rate of the oscillations in $h(r)$ does not change significantly with dimension, whereas the amplitude of oscillations does. However, further studies with larger r and therefore larger systems are needed to obtain more quantitative results.

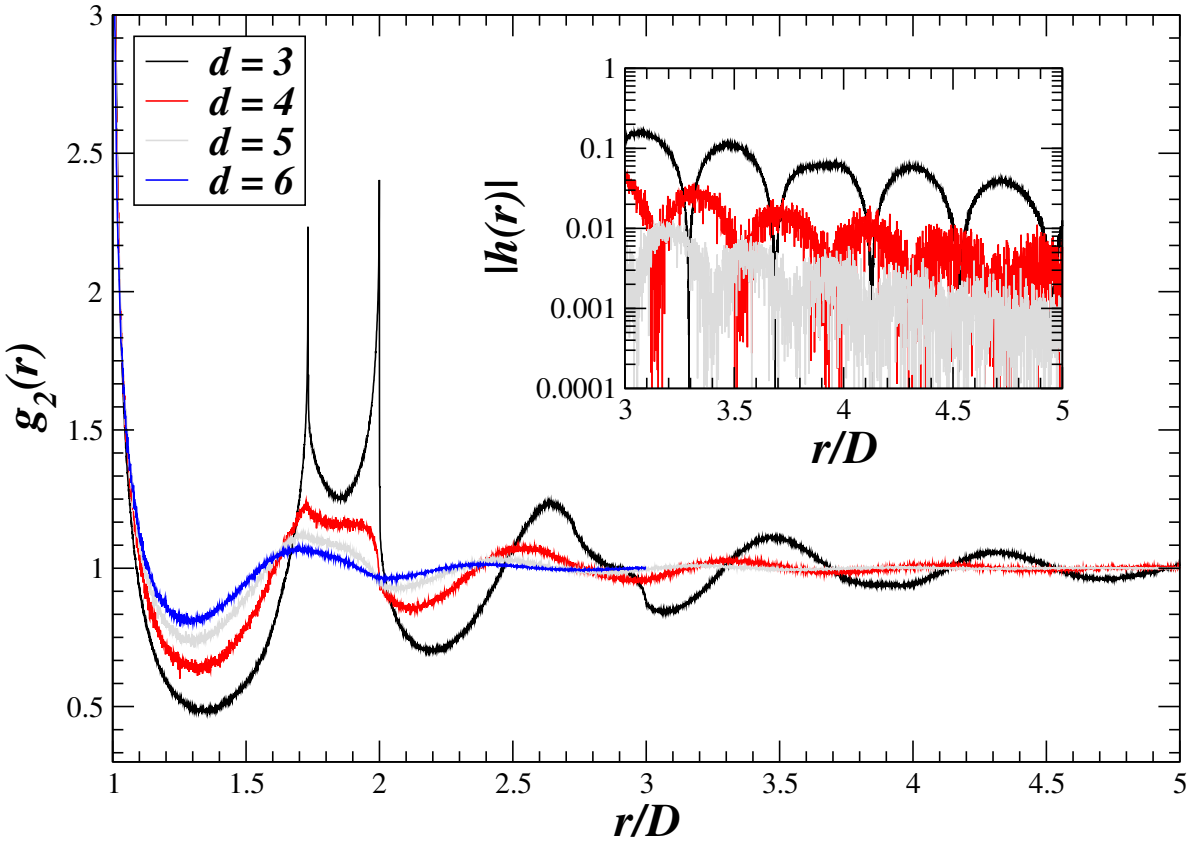

Figure 11.5: The pair correlation function $g_2(r)$ for MRJ packings of $N = 10^5$ hard spheres for $d = 3, 4, 5$ and 6. Pair separation is plotted in units of the sphere diameter D. (For $d = 6$, $g_2(r)$ was only calculated up to $r/D = 3$ due to the system size and periodic boundary conditions). The delta-function contribution [c.f. Eq. 11.5] at contact, of course, is not shown. The inset shows $|h(r)| = |g_2(r) - 1|$ on a logarithmic scale for $d = 3, 4$ and 5.

11.3.1.2 Structure Factor

We can calculate the structure factor $S(k)$, defined in Eq. 11.4, for $d = 4$ and $d = 5$, by

$$S(K) = 1 + 128\phi \int_0^\infty x^3 h(x) \frac{J_1(Kx)}{Kx} dx \tag{11.6}$$

and

$$S(K) = 1 + 480\phi \int_0^\infty \frac{x^4 h(x)}{(Kx)^2} \left[\frac{\sin(Kx)}{Kx} - \cos(Kx) \right] dx, \tag{11.7}$$

respectively, where $\phi = \pi^2 \rho D^4/32$ for $d = 4$ and $\phi = \pi^2 \rho D^5/60$ for $d = 5$, $x = r/D$ and $K = kD$ are the dimensionless radius and wave number, and $J_\nu(x)$ is the Bessel function of order ν. We do not calculate the structure factor for $d = 6$ because at present we do not have $g_2(r)$ over a sufficiently large range of r.

Following the procedure used in chapter 10 for $d = 3$, rather than working directly with $g_2(x)$ as in Eq. (11.4), we consider the average cumulative coordination $Z(x)$ [76],

$$Z(x) = \rho \int_1^x s_1(x')g_2(x')dx'. \tag{11.8}$$

The excess coordination $\Delta Z(x)$,

$$\begin{aligned} \Delta Z(x) &= 1 + 64\phi \int_0^x (x')^3 h(x')dx' \\ \Delta Z(x) &= 1 + 160\phi \int_0^x (x')^4 h(x')dx', \end{aligned}$$

for $d = 4$ and $d = 5$, respectively, is the average excess number of sphere centers inside a spherical window of radius x centered at a sphere, compared to the ideal gas expectations, $16\phi x^4$ for $d = 4$ and $32\phi x^5$ in $d = 5$. We can rewrite Eq. (11.4) in terms of $\Delta Z(x)$ using integration by parts to get

$$S(K) = -2 \int_0^\infty \Delta Z(x) \frac{d}{dx} \frac{J_1(Kx)}{Kx} dx$$

and

$$S(K) = -3 \int_0^\infty \Delta Z(x) \frac{d}{dx} \left[\frac{\sin(Kx)}{(Kx)^3} - \frac{\cos(Kx)}{(Kx)^2} \right] dx,$$

for $d = 4$ and $d = 5$, respectively. Note that accurate evaluations of the integrals of $\Delta Z(x)$ require extrapolations of its large-x tail behavior, for which we have used an exponentially-damped oscillating function

$$\Delta Z(x) = a_1 x e^{-a_2 x} \cos(a_3 x + a_4), \tag{11.9}$$

where a_1, a_2, a_3, and a_4 are fitting parameters, as in Chapter 10.

Figure 11.6 shows $S(k)$ for jammed packings of $N = 10^5$ spheres in three, four and five dimensions. Qualitatively, $S(k)$ is somewhat similar for $d = 3$, 4, and 5. However, with increasing dimension, the height of the first peak of $S(k)$ decreases, the location of the first peak moves to smaller wavelengths, and the oscillations become damped. The width of the first peak also increases with increasing dimension, which could indicate that the correlation length decreases with increasing dimension. The inset to Fig. 11.6 shows $S(k)$ for a jammed packing and a fluid near the freezing point in four dimensions. The relation between the structure factor for the fluid and jammed packing is strikingly similar to what is found for $d = 3$, except that the peaks of both curves for $d = 4$ appear scaled down relative to $d = 3$. Overall, our results for both $g_2(r)$ and $S(k)$ are consistent with a recently proposed "decorrelation" principle [240]. We note that similar pair decorrelations are observed for RSA packings as the dimension increases up to $d = 6$ [246].

It is of interest whether infinite wavelength density fluctuations $S(k = 0)$ vanish, i.e., whether MRJ packings are hyperuniform [202] in dimensions higher than three as well. As we observed in Chapter 10 for $d = 3$, $S(k)$ for $d = 4$ appears to go to zero faster near the origin for the jammed packing than for the fluid. However, we cannot reliably determine whether $S(k)$ vanishes at the origin because our calculation of $S(k)$ for small k involved an extrapolation of the large-x tail of $\Delta Z(x)$. Nevertheless, using larger system sizes of one million spheres, in Chapter 10 we demonstrated that saturated MRJ packings for $d = 3$ are hyperuniform to a high accuracy [76] and the comparison of $d = 4$ and $d = 5$ to $d = 3$, shown in Fig. 11.6, suggests MRJ packings for $d = 4$ and $d = 5$ are also hyperuniform.

11.3.2 Isostaticity

We study the near-contact contribution to $g_2(r)$, i.e., interparticle distances r that are very close to the sphere diameter D, using the cumulative coordination number $Z(x)$, where as before $x = r/D$ is the dimensionless radius and $x - 1$ is the dimensionless interparticle gap. Figure 11.7 shows $Z(x)$ for jammed packings of $10,000$ spheres for $d = 4$ and $d = 5$ with rattlers removed. Due to computational constraints, our packings for $d = 6$ were produced with a relatively high expansion rate ($\gamma = 10^{-3}$) and were not grown to sufficiently high pressures, as necessary to properly distinguish between true contacts and near contacts; therefore, we

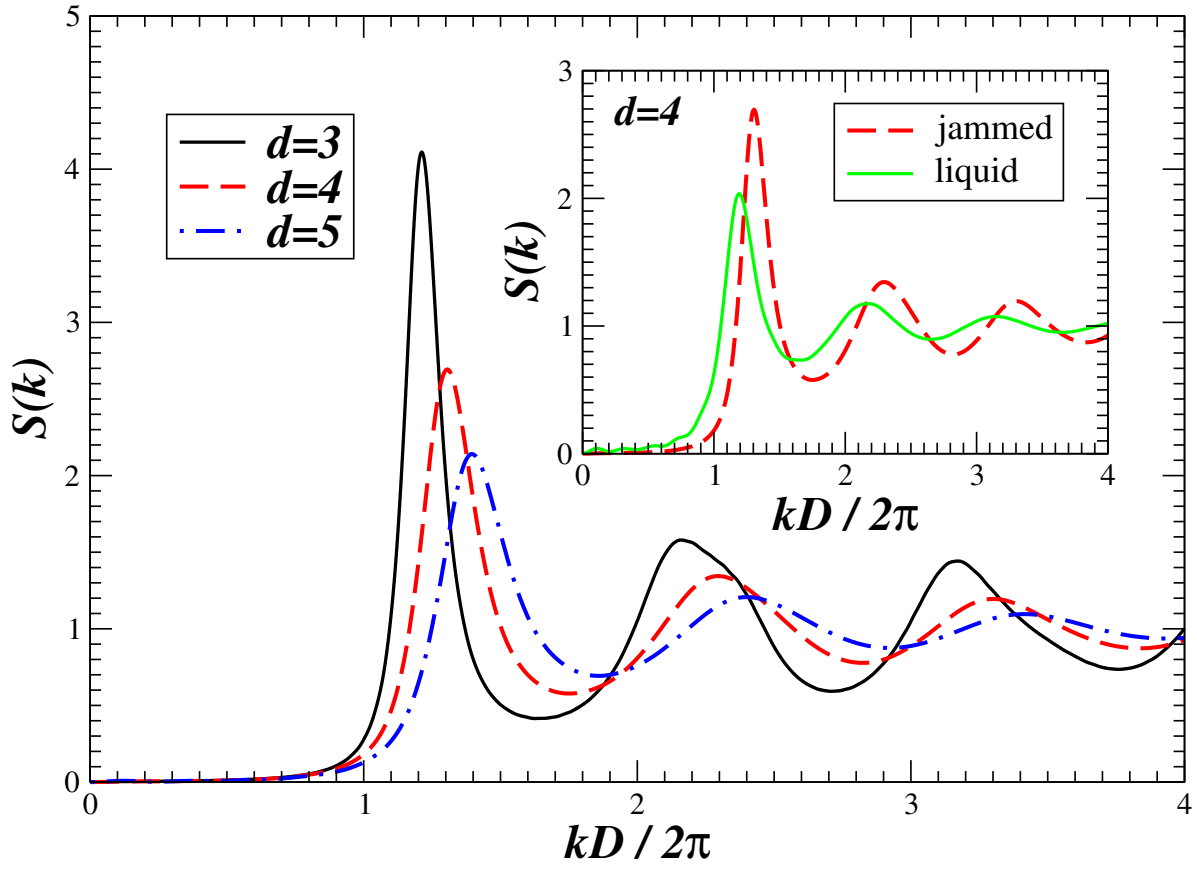

Figure 11.6: The structure factor $S(k)$ for jammed packings of $N = 10^5$ hard spheres for $d = 3$, 4 and 5. Inset: A comparison for $d = 4$ of $S(k)$ for a jammed packing and for a fluid near the freezing point.

do not show results for $d = 6$ in this section. The plateaus at $Z = 8$ in Fig. 11.7 (a) and $Z = 10$ in Fig. 11.7 (b) show that both packings are isostatic (see Section 5.4).

Packings produced by the LS algorithm almost always contain a nonzero fraction of "rattlers", which are spheres trapped in a cage of jammed neighbors, but free to move within the cage. We find approximately $\sim 1\%$ rattlers for $d = 4$ and $\sim 0.6\%$ rattlers for $d = 5$, as compared to $\sim 2 - 3\%$ rattlers for $d = 3$ [75]. Rattlers can be identified as having less than the required $d + 1$ contacts necessary for local jamming and are removed to study the jammed backbone of the packing, which we focus on in this section.

The insets in Fig. 11.7 show $Z(x) - 2d$, along with a power-law fit for intermediate interparticle gap $x - 1$,

$$Z(x) = \bar{Z} + Z_0(x - 1)^{\alpha}, \tag{11.10}$$

where $\bar{Z} = 2d$. Since the packings are generally slightly subisostatic, we apply a small correction ($< 0.1\%$) to the isostatic prediction of $2d$ by using the midpoint of the apparent plateau in $Z(x)$. The best-fit exponent is $\alpha \simeq 0.6$ in both $d = 4$ and $d = 5$, in agreement with that found for $d = 3$ [75]. The coefficients of the power law, $Z_0 \simeq 11$ in $d = 3$, $Z_0 \simeq 24$ for $d = 4$, and $Z_0 \simeq 40$ for $d = 5$ are close to the corresponding kissing numbers of the densest packings, $Z = 12$ for $d = 3$, $Z = 24$ for $d = 4$, $40 \leq Z \leq 46$ for $d = 5$ and $72 \leq Z \leq 80$ for $d = 6$. Motivated by this observation, we measured the value of the gap $x - 1$ at which the cumulative coordination $Z(x)$ equals the kissing number of the densest packing to be: $x - 1 \simeq 0.35$, 0.34, $0.31 - 0.36$ and $0.33 - 0.36$ in $d = 3$, 4, 5 and 6, respectively, which we can define to be the cutoff for the near-neighbor shell. This definition produces results similar to that of the more common definition of the cutoff for the near-neighbor shell as the value of the gap $x - 1$ at the first minimum in g_2, which occurs at $x - 1 \simeq 0.35$, 0.32, 0.30 and 0.28 in $d = 3$, 4, 5 and 6, respectively. It is also interesting to observe

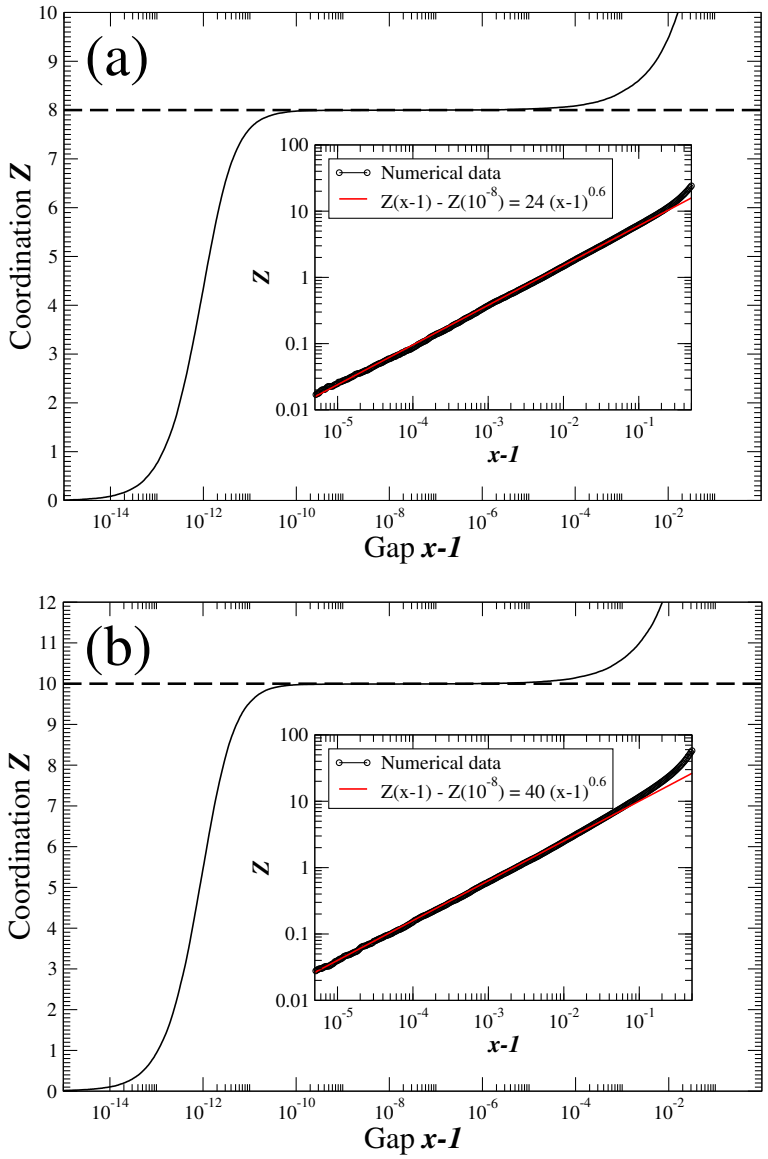

Figure 11.7: The near-contact cumulative coordination $Z(x)$ for 10,000-sphere MRJ packings for $d = 4$ (a) and for $d = 5$ (b), with rattlers removed. The inset shows $Z(x)$ on a log-log scale along with power-law fits for intermediate interparticle gap $x - 1$ beyond contact. Packings of $N = 10^5$spheres in $d = 5$ with final expansion rates of $\gamma = 10^{-4}$ give similar results; such packings with final expansion rates of $\gamma = 10^{-5}$ are (presently) too computationally expensive. Compare these plots to the equivalent results for $d = 3$ in Fig. 9.6.

that the power-law fit to $Z(x)$ is good over a rather wide range of gaps, almost up to the first minimum in g_2. We should, however, emphasize that the minimum of g_2 is not very precisely defined, especially due to decorrelation in high dimensions, and the choice of the gap at the minimum of g_2, or at which $Z(x)$ equals the kissing number of the densest packing, as a special point is somewhat arbitrary and not theoretically justified at present.

11.4 Conclusions

We have presented the first numerical results characterizing random jammed hard-sphere packings in four, five and six dimensions. We find disordered packings, representative of the maximally random jammed state, to be isostatic and have packing fractions $\phi_{MRJ} \simeq 0.46$, $\phi_{MRJ} \simeq 0.31$ and $\phi_{MRJ} \simeq 0.20$ for $d = 4$, 5 and 6, respectively. For equilibrium sphere packings, we estimate the freezing and melting packing fractions for the fluid-solid transition in four dimensions to be $\phi_F \simeq 0.32$ and $\phi_M \simeq 0.39$, respectively, and in five dimensions to be $\phi_F \simeq 0.19$ and $\phi_M \simeq 0.24$, respectively. Additionally, a signature characteristic of the kinetic glass transition is observed around $\phi_G \simeq 0.41$ for $d = 4$. We observe a significantly lower tendency to crystallize for $d = 4$ than in $d = 3$, which is likely due to the closer proximity of the melting and kinetic glass transition densities for $d = 4$ [57].

We find that in high dimensions the split-second peak in the pair correlation function g_2, present for $d = 3$, gets dramatically diminished and oscillations in both g_2 and the structure factor $S(k)$ get significantly dampened. These findings are consistent with a recently proposed "decorrelation principle" [240], stating that unconstrained spatial correlations vanish asymptotically in the high-dimensional limit. Accordingly, in this limit the pair correlation function $g_2(r)$ would be expected to retain the delta-function contribution from nearest-neighbor contacts, but the extra structure representing unconstrained spatial correlations beyond a single sphere diameter would vanish. Figures 11.5 and 11.6 show dramatically the decorrelation principle already taking effect in four, five and six dimensions. We note that decorrelation principle is also apparent in the same dimensions for RSA packings [246].

A particularly interesting property of jammed hard-sphere packings is hyperuniformity, the complete suppression of infinite wavelength density fluctuations, i.e., the vanishing of the structure factor $S(k)$ as $k \to 0$. It has been recently conjectured that all saturated strictly-jammed packings are hyperuniform [202] and calculations of the structure factor near $k = 0$ for $d = 3$ using one million particle systems have strongly suggested that MRJ packings for $d = 3$ are indeed hyperuniform [76]. Though the system sizes used in this paper were too small to probe such large-scale density fluctuations without relying on dubious extrapolations, our numerical results for the structure factor for $d = 4$ and $d = 5$, as shown in Fig. 11.6, are consistent with hyperuniformity.

As in three dimensions, disordered jammed sphere packings show no signs of crystallization, are isostatic, and have a power-law divergence in $g_2(r)$ at contact. Interestingly, all three dimensions (3, 4 and 5) share the same power law exponent $1 - \alpha = 0.4$ when rattlers are removed, and show the first minimum of $g_2(r)$ close to where the cumulative coordination $Z(r)$ equals the kissing number of the densest lattice packing. Such a relation between the kissing numbers of the densest packings and MRJ packings for $d = 3$, 4, 5 and 6, if not coincidental, is very surprising and may be a consequence of the geometrical structure of MRJ packings. It suggests that disordered packings might be deformed crystal packings, in which the true contacts are deformed into near contacts, and only the minimal number of contacts necessary for jamming is preserved. This interpretation is to be contrasted with the usual interpretation of disordered packing in $d = 3$ in terms of tetrahedral or icosahedral packings, without relation to the crystal (FCC) packing. This interpretation is similar to the interpretation of the MRJ state for binary hard-disks as a random partitioning of the monodisperse triangular crystal into "small" and "large" disks, i.e., a deformed monodisperse triangular disk crystal in which a randomly chosen fraction of the particles have grown in size, proposed in Ref. [57].

It is important to point out that hard-sphere packings behave rather differently in two dimensions than in three and higher dimensions. For $d = 2$, jammed hard-sphere systems are polycrystalline and there is a very weak, nearly continuous fluid-solid phase transition. Hence, there is no glassy behavior for $d = 2$ and consequently no amorphous jammed packings. Glassy behavior, due to geometrical frustration arising from the inconsistency of local optimal packing rules and global packing constraints, first appears in three dimensions [6]. It is likely that geometrical frustration generally increases with dimension, consistent with

our observation that nucleation is suppressed with increasing dimension.

Computational costs rise dramatically with increasing dimension and theoretical understanding based on observations in moderate dimensions is necessary. We believe that the numerical results presented in this work provide tests and motivations for such theories.

Chapter 12

Tetratic Order in Systems of Hard Dominos

Previous Monte Carlo investigations by Wojciechowski *et al.* have found two unusual phases in two-dimensional systems of anisotropic hard particles: a tetratic phase of four-fold symmetry for hard squares [247], and a nonperiodic degenerate solid phase for hard-disk dimers [248, 249]. In this Chapter , we study a system of hard rectangles of aspect ratio two, i.e., hard-square dimers (or dominos), and demonstrate that it exhibits a phases with both of these unusual properties [81]. The liquid shows quasi-long-range tetratic order, with no nematic order. The solid phase we observe is a nonperiodic tetratic phase having the structure of a random tiling of the square lattice with dominos with the well-known degeneracy entropy $1.79k_B$ per particle. Our simulations do not conclusively establish the thermodynamic stability of this orientationally disordered solid; however, there are strong indications that this phase is glassy. Our observations are consistent with a KTHNY two-stage phase transition scenario with two continuous phase transitions, the first from isotropic to tetratic liquid, and the second from tetratic liquid to solid. We obtain similar results with both a classical Monte Carlo method using true rectangles and a novel molecular dynamics algorithm employing rectangles with rounded corners.

12.1 Introduction

Hard-particle systems have provided a simple and rich model for investigating phase behavior and transport in atomic and molecular materials. It is long-known that a pure hard-core exclusion potential can lead to a variety of behaviors depending on the degree of anisotropy of the particles, including the occurrence of isotropic and nematic liquids, layered smectic, and ordered solid phases [250]. Through computer investigations of various particle shapes, other phases have been found, such as the biaxial [251] (recently synthesized in the laboratory [252]) and cubatic phases in three dimensions, in which the axes of symmetry of the individual particles align along two or three perpendicular axes (directors). One only need look at simple shapes in two dimensions to discover interesting phases. In recent work, Wojciechowski *et al.* studied hard squares and found the first example of a tetratic liquid phase at intermediate densities [247]. In a tetratic liquid, there is (quasi) long-range orientational ordering along two perpendicular axes, but only short-range translational ordering. The solid phase is the expected square lattice, with quasi-long-range periodic ordering. On the other hand, by studying hard-disk dimers (two disks fused at a point on their boundary), they have identified the first example of a nonperiodic solid phase at high densities [248]. In this phase, the centroids of the particles are ordered on the sites of a triangular lattice. However, the orientations of the dimers are disordered, leading to a high degeneracy entropy of the nonperiodic solid and a lower free energy as compared to periodic solids. An experimental study of orientational ordering in nonequilibrium (quasi)two-dimensional systems was recently performed using vibrated granular monolayers [200]. It was found that the exact particle shape impacts the nature of the orientational ordering substantially, with tetratic ordering appearing for rectangles (monolayers of cylinders) even for high aspect ratios, and nematic ordering appearing for more rounded particles.

We look at systems of rectangles of aspect ratio $\alpha = a/b = 2$, i.e., hard-square dimers (or dominos). Since the aspect ratio is far from unity, it is not clear *a priori* whether nematic or tetratic orientational ordering (or both) will appear. Theoretical investigations of tetratic ordering, similar to Onsager's studies of nematic ordering, were performed by Zwanzig for a restricted-orientation model of a system of hard rods [253]. Zwanzig found a phase transition into a nematic phase at sufficiently high densities for very elongated rods. The theory was however focused on three-dimensional systems of very elongated rods. Furthermore, restricted-orientation models have been shown to be misleading [254]. Recent Density Functional Theory calculations [255], extending previous work based on scaled-particle theory [256], have predicted that for $\alpha = 2$ the tetratic phase is only metastable with respect to the ordered solid phase in which all particles are aligned. However, these calculations are only approximate and the authors point out that tetratic order is still possible in spatially ordered phases. An obvious candidate for forming a stable tetratic phase are dominos: two dominos paired along their long edges form a square, and these squares can then form a square lattice assuming one of two random orientations, thus forming a tetratic phase with degeneracy entropy of $\ln(\sqrt{2})$. In fact, one does not need to pair up the rectangles but rather simply tile a square lattice with dominos which randomly assume one of the two preferred perpendicular directions. The degeneracy entropy of this domino tiling has been calculated exactly to be $(2G/\pi)k_B \approx 0.58313k_B$ [257, 258], where $G = \sum_{n=0}^{\infty}(-1)^n(2n+1)^{-2} \approx 0.91597$ is Catalan's constant.

At high densities, free-volume theory [250] predicts that the configurational entropy (per particle) diverges like

$$s_{\mathrm{FV}} \sim f \ln(1 - \phi/\phi_c) + s_c,$$

where f is the (effective) number of degrees of freedom per particle, ϕ_c is the volume fraction (density) at close packing, and s_c is an additive constant due to collective exclusion-volume effects (see Chapters 6 and 13). Therefore, the densest solid is thermodynamically favored, but if several solids have the same density the additive factor matters, as discussed in see Chapter 6. Therefore, for hard rectangles, for which the maximal density is $\phi_c = 1$ and is achieved by a variety of packings, the degeneracy entropy can dominate s_c and thus the nonperiodic random tiling can be thermodynamically favored. Indeed, our simulations of the hard-domino system produce high-density phases with structures very similar to that of a random covering of the square lattice with dimers. However, additional free-energy calculations and more sophisticated Monte Carlo simulations of the solid phase are necessary for definite answers.

The phase transitions in two-dimensional systems are of interest to the search for continuous KTHNY [47, 48, 49] transitions between the disordered liquid and the ordered solid phase. At present there is no agreement on the nature of the transition even for the hard-disk system. A previous study of the melting of a square-lattice crystal, stabilized by the addition of three-body interactions, found evidence of a (direct) first-order melting [259]. Our observations for the domino system are relatively consistent with a KTHNY-like two-stage transition: a continuous phase transition from an isotropic to a tetratic liquid with long-range tetratic order around $\phi \approx 0.7$, and then another continuous transition from tetratic liquid to tetratic solid with quasi-long-range translational order $\phi \approx 0.8$. However, we cannot rule out the possibility of a weak first-order phase transition between the two phases without more detailed simulations.

This chapter is organized as follows. In Section 12.2, we present the simulation techniques used to generate equilibrated systems at various densities. In Section 12.3, we analyze the properties of the various states, focusing on the orientational and translational ordering in the high-density phases. We conclude with a summary of the results and suggestions for future work in Section 12.4.

12.2 Simulation Techniques

In this section, we provide additional details on the MC and MD algorithms we implemented. It is important to point out that it is essential to implement techniques for speeding up the near-neighbour search, in both MC and MD. For rectangles with a small aspect ratio, we employ the well-known technique of splitting the domain of simulation into cells (bins) larger than a particle diameter $D = \sqrt{a^2 + b^2}$, and consider as neighbors only particles whose centroids belong to neighboring cells. Additional special techniques more suitable for very aspherical particles or systems near jamming are described in Ref. [11].

12.2.1 Monte Carlo

We have implemented a standard MC algorithm in the NVT ensemble, with the additional provision of changing the density by growing or shrinking the particles in small increments. Each rectangle is described by the location of its centroid (x, y) and orientation θ. For increased computational speed the pair $(\sin\theta, \cos\theta)$ may be used to represent the orientation. In a trial MC step, a rectangle is chosen at random and its coordinates are changed slightly, either translationally $(\Delta x, \Delta y)$ or orientationally $(\Delta\theta)$. Every move has an equal chance of being translational or orientational. The rectangle's new position is then compared against nearby rectangles for overlap; if there is no overlap, the trial move is accepted. We call a sequence of N trials a cycle. The simulation evolves through *stages*, defined by a speed $n_{cycles/stage}$. At the end of a cycle, pressure data are collected by the virtual-scaling method of Eppenga and Frenkel [166]. Namely, $p = PV/NkT = 1 + \phi\alpha/2$ where α is the rate at which growing the particles causes overlaps. At the end of a stage, order parameters and other statistics are collected, and then the packing fraction ϕ is changed by a small value $\Delta\phi$; it may be increased, decreased, or not changed at all. If $\Delta\phi > 0$, then ϕ cannot necessarily change by $\Delta\phi$ every stage, because the increase could create overlaps. We scale down the increase by factors of 2 until a $\Delta\phi_{\text{eff}}$ is found that does not cause any overlaps when applied. Typical values for runs are $n_{cycles/stage} = 1000$ and $\Delta\phi = \pm 1 \times 10^{-5}$. Since there is a limit on how fast one can increase the density in such a Monte Carlo simulation, especially at very high densities, we use Molecular Dynamics to compress systems to close packing.

The overlap test is by far the largest computational bottleneck in the MC program. The overlap test for two rectangles is based on the following fact: Two rectangles R_1 and R_2 do not overlap if and only if a separating line ℓ can be drawn such that all four corners of R_1 lie on one side of the line and all four corners of R_2 lie on the other side [133]. The corners of both rectangles are allowed to coincide with ℓ. Without loss of generality, we may assume ℓ is drawn parallel to one of the rectangles' major axes and runs exactly along that rectangle's side. The problem of testing all possible lines ℓ thus reduces to testing the eight lines that coincide with the edges of R_1 and R_2. The test can be optimized somewhat further, as illustrated in Fig. 12.1. An axis $\bar{a}$ of R_1 is chosen. The distance from the center of R_2 to $\bar{a}$ is found. Then the distance d_0 of closest approach of R_2 to a is found by subtracting a sine and a cosine; this distance corresponds to the corner of R_2 that is closest to $\bar{a}$. By comparing d_0 with the length b of the other semiaxis of R_1, two possible lines ℓ, corresponding to two opposite sides of R_1, can be tested at once. If $d_0 < b$, there is an overlap. In this way, four different values of d_0 are calculated; one for each axis of each rectangle. If no comparison finds an overlap, there is no overlap.

12.2.2 Molecular Dynamics

MC simulations are typically the most efficient when one is only interested in stable equilibrium properties. Using the overlap potentials described in Section see Chapter 6, we have implemented the MD algorithm described in Chapter 3 for superellipses and superellipsoids. A superellipse with semiaxes a and b is given by the equation

$$\left[\left|\frac{x}{a}\right|^{2\zeta} + \left|\frac{y}{b}\right|^{2\zeta}\right]^{1/\zeta} \leq 1,$$

where $\zeta \geq 1$ is an exponent. When $\zeta = 1$ we get the simple ellipse, and when $\zeta \to \infty$ we obtain a rectangle with sides $2a$ and $2b$. The higher the exponent the sharper the corners become. The floating-point cost of the algorithm increases as the exponent increases, while the numerical stability decreases. We have used an exponent $\zeta = 7.5$ for the studies presented here (for this exponent the ratio of the areas of the superellipse and the true rectangle is 0.9934). Figure 12.2 gives an illustration of the particle shape.

There are some advantages of the MD simulation over MC. The shapes of the particles can change arbitrarily fast in an easily controlled manner by simply adding a dynamic growth rate $\gamma = da/dt = \alpha db/dt$. If $\gamma > 0$, i.e., the density is increasing, two colliding particles simply get an extra repulsive boost that ensures no overlaps are created. The velocities are periodically rescaled to $T = 1$ to compensate for the induced heating or cooling due to the particle growth [11]. In general, (common) MC methods do not work well near close packing, while MD methods, especially event-driven ones, can successfully be used to study the neighbourhood of jamming points. Additionally, pressure measurement is more natural in the MD method, as the pressure can be directly obtained from time averages of the momentum exchange in binary collisions

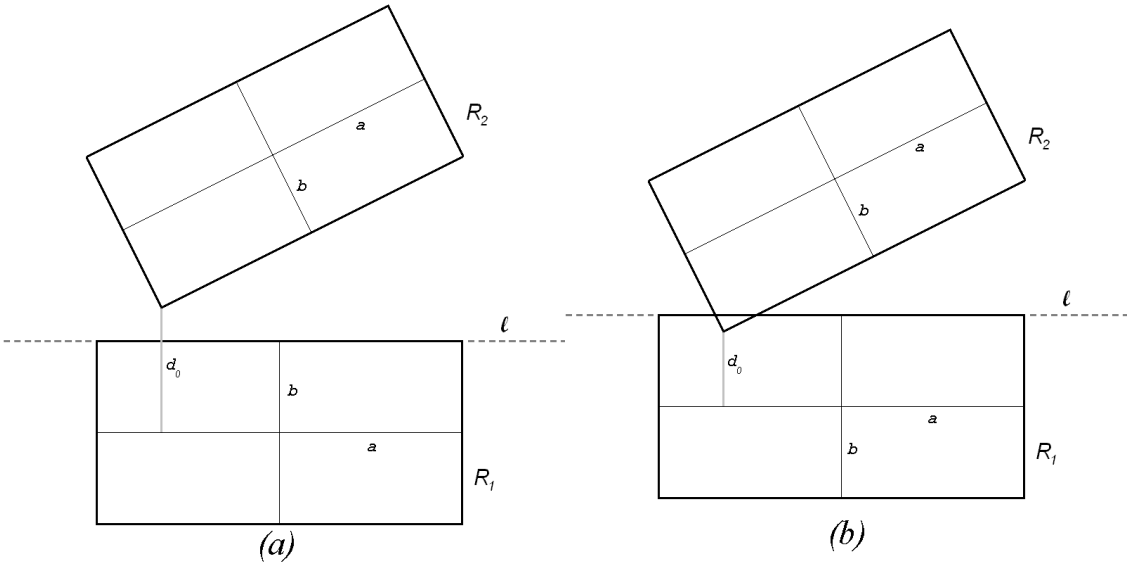

Figure 12.1: Illustration of the optimized overlap test for two rectangles. The axes are $\bar{a}$ and $\bar{b}$, with semiaxes a and b, and the length from $\bar{a}$ to the closest corner of R_2 is d_0. (*Left)* If $d_0 \geq b$, then the rectangles do not overlap. (Right) If $d_0 < b$, then the rectangles overlap.

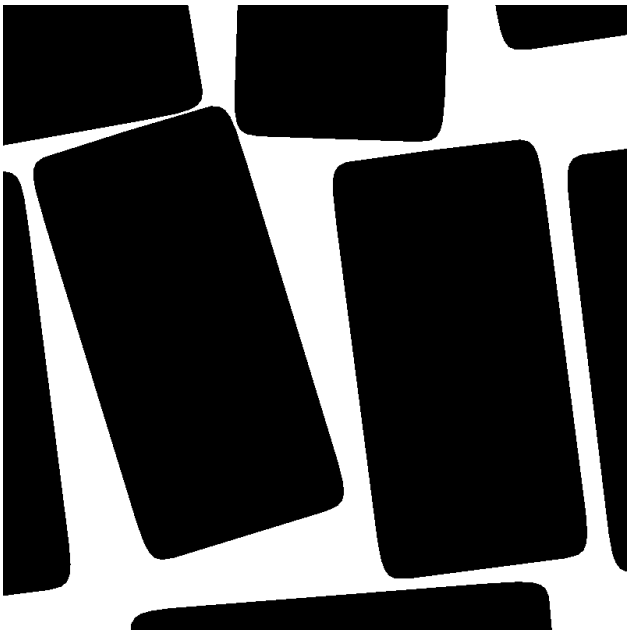

Figure 12.2: An snapshot of a few superellipses (exponent $\zeta = 7.5$) used in the MD simulations. It can be seen that the particle shape is very close to a rectangle.

between particles. We have found this pressure measurement to be much more precise than using virtual particle scaling in MC simulations.

12.3 Results

By using either the MC or the MD algorithm with small particle growth rate ($\Delta\phi$ or γ), we have traced the (quasi)equilibrium phase behavior of systems of dominos over a range of densities. In this section, we present several techniques for measuring orientational and translational order for a given configuration of particles, as well as the results of such measurements for the generated states. We have tested our codes by first applying them to hard squares and comparing the results to those in Ref. [247], and we have observed good quantitative agreement throughout. Our MC pressure measurement systematically slightly underestimates the pressure compared to the NPT ensemble used in Ref. [247] and to our MD simulations. We present some of the results for the MC, and others for the MD simulations, marking any quantitative differences. The two techniques always produced qualitatively identical results.

Describing the statistical properties of the observed states would require specifying all of the n-particle correlation functions. The most important is the pair correlation function $g_2(r, \psi, \Delta\theta)$. Given a particle, $g_2(r, \psi, \Delta\theta)$ is the probability density of finding another particle whose centroid is a distance r away (from the centroid of the particle), at a displacement angle of ψ (relative to the first particle's coordinate axes), and with a orientation of $\Delta\theta$ (relative to the particle's orientation). The normalization of g_2 is such that it is identically unity for an ideal gas. We will use an equivalent representation where we fix a particle at the origin such that the longer rectangle axis is along the x axes, and represent pair correlations with $g_2(\Delta x, \Delta y, \Delta\theta)$, giving the probability density that there is another particle whose centroid is at position $(\Delta x, \Delta y)$ and whose major axis is at a relative angle of $\Delta\theta$. Since a three-dimensional function is rather difficult to calculate accurately and visualize, we can separate the translational and orientational components and average over some of the dimensions to reduce it to a one- or two-dimensional function.

The symmetry and nature of ordering in condensed phases is most easily accessed by using order parameters; specifically these would be scalar *order metrics*, which are typically averaged forms of local order parameters. Several types of order metrics are commonly applied in studies of two-dimensional hard-particle systems: *orientational*, *bond-orientational*, and *translational* order metrics, or combinations of orientational and bond-orientational order metrics [260]. We will present results for each of these types of ordering for the domino system. It would be useful to construct bond-orientational and translational order metrics that do not depend on splitting each domino into two squares in the future.

12.3.1 Equation Of State

The pressure as a function of density can be most accurately measured in the MD simulations. There is no exact theory that can predict the entire equation of state (EOS) for a given many-particle system. However, there are two simple theories that produce remarkably good predictions for a variety of systems studied in the literature. For the isotropic fluid (gas) phase of a system of hard dominos, scaled-particle theory (SPT) [261] predicts

$$p = \frac{1}{1-\phi} + \frac{9}{2\pi}\frac{\phi}{(1-\phi)^2}, \tag{12.1}$$

and modifications to account for possible orientational ordering are discussed in Refs. [256, 255]. For the solid phase, free-volume (FV) theory predicts a divergence of the pressure near close packing of the form [c.f. Eq. (5.24)]

$$p = \frac{3}{1-\phi/\phi_c}, \tag{12.2}$$

and (liquid-state) density functional theory can be used to make quantitative predictions at intermediate densities [255]. For superellipses with exponent $\zeta = 7.5$ the maximal density is somewhat less than 1 and we take it to be equal to the ratio of the areas of the particle and a true rectangle, $\phi_c \approx 0.9934$.

The numerical EOS from the NVT MD simulation are shown in Fig. 12.3 for both a slow compression starting from an isotropic liquid and a decompression starting from a perfect random domino tiling generated with the help of random spanning trees, as explained in Ref. [262]. We note that the random domino tiling

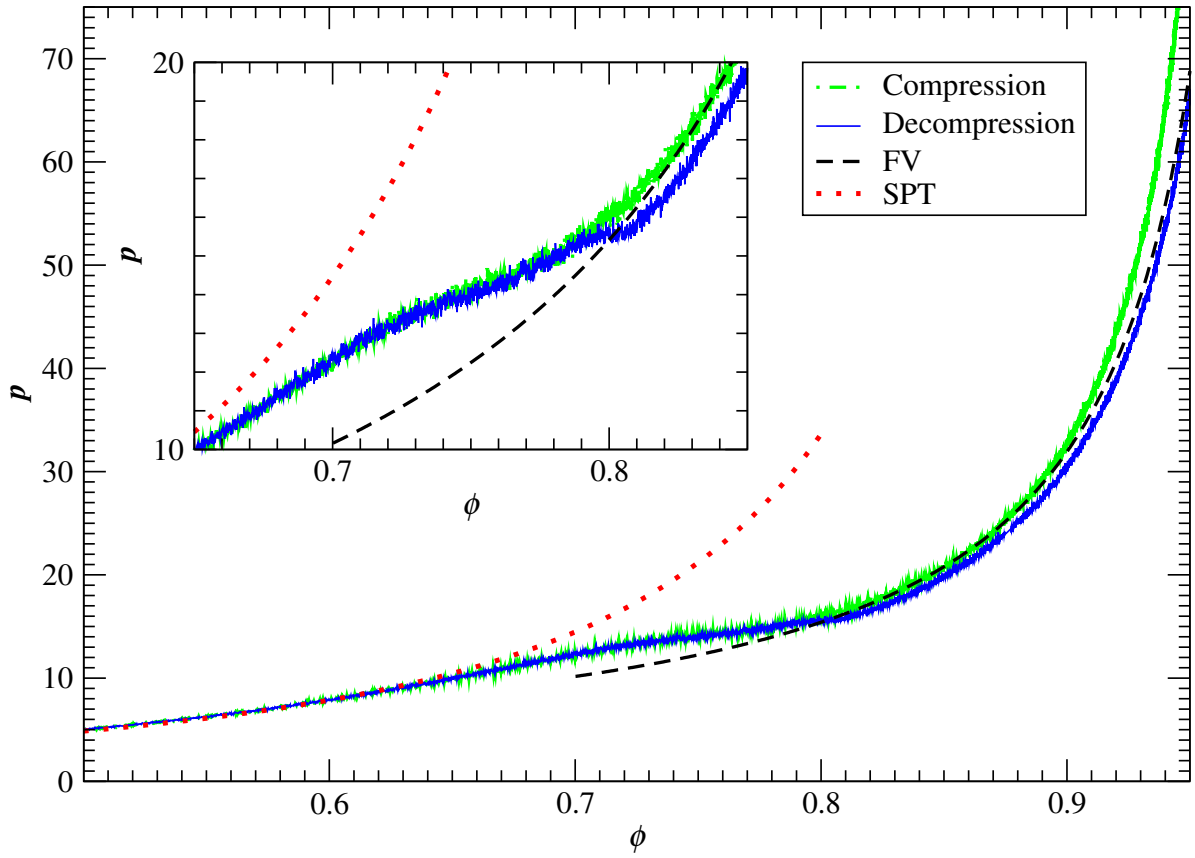

Figure 12.3: Reduced pressure $p = PV/NkT$ in a system of $N = 5000$ superellipses with exponent $\zeta = 7.5$ during MD runs with $\gamma = \pm 2.5 \cdot 10^{-5}$. The predictions of simple versions of SPT and FV theory are also shown for comparison. The agreement with FV predictions is not perfect; a numerical fit produces a coefficient 2.9 instead of 3 in the numerator of Eq. (12.2). Particularly noticeable are the change in slope around $\phi \approx 0.72$ and also the transition onto a solid branch well-described by free-volume approximation around $\phi \approx 0.8$. Starting the decompression from an ordered tiling in which all rectangles are aligned produces identical pressure to within the accuracy available. Systems of $N = 1250$ and $N = 10000$ particles, as well as a wide range of particle growth rates, were investigated to ensure that there were no strong finite-size or hysteresis effects. In faster compressions of an isotropic liquid one gets smaller final densities due to the occurrence of defects such as vacancies or grain boundaries.

used was generated inside a square box (see Fig. 12.11) even though periodic boundary conditions were used in the actual simulation. We expect this to have a very small effect [263]. It is clearly seen from the figure that there is a transition from the liquid to the solid branch in the region $\phi \approx 0.7$ and $\phi \approx 0.8$, although no clear discontinuities or a hysteresis loop are seen (which would be indicative of a first-order phase transition). Compressing an isotropic liquid invariably freezes some defects and thus the jamming density is smaller (and the pressure is thus higher) than in the perfect crystal.

12.3.2 Orientational Order

Orientational order can be measured via the orientational correlation function of order m

$$G_m(r) = \langle \cos(m\Delta\theta) \rangle_r, \tag{12.3}$$

where m is an integer and the average is taken over all pairs of particles that are at a distance between r and $r + dr$ apart from each other. The one-dimensional function $G_m(r)$ can be thought of as giving normalized

Fourier components of the distribution of relative orientations versus interparticle distance. When $m = 2$, it measures the degree of nematic ordering (parallel alignment of the particles' major axes), and when $m = 4$ it measures the degree of tetratic ordering (parallel alignment of the particles' axes). The infinite-distance value $\lim_{r\to\infty} G_m(r) = S_m$ gives a scalar measure of the tendency of the particles to align with a global coordinate system; S_2 is the usual nematic order parameter, and S_4 is the tetratic order parameter. They can be very easily calculated from an alternative definition

$$S_m = \max_{\theta_0} \langle \cos[m(\theta - \theta_0)] \rangle \,, \tag{12.4}$$

which can be converted into an eigenvalue problem (in any dimension) for the case $m = 2$ [264]. When $m = 4$, we can rewrite it in the same form as $m = 2$ by replacing θ with 2θ. The vector $\mathbf{n}_m = (\cos\theta_0, \sin\theta_0)$ determines a natural coordinate system for orientationally ordered phases. It is commonly called the director for nematic phases ($m = 2$), and we will refer to it as a bidirector for tetratic phases ($m = 4$).

In two-dimensional liquid-crystalline phases, it is expected that there can be no long-range orientational ordering, but rather only quasi-long-range orientational ordering [265]. Based on elasticity theory with a single renormalized Frank's constant $\tilde{K} = \pi K/(8k_BT)$, it is predicted [266] that there will be a power-law decay of the correlations at large distances, $G_m(r) \sim r^{-\eta}$, where

$$\eta = m^2/16K. \tag{12.5}$$

This would imply that S_m vanishes with increasing system size,

$$S_m \sim N^{-\eta/4}. \tag{12.6}$$

We note that this prediction is based on literature for the nematic phase. We are not aware of any theoretical work explicitly for a tetratic phase.

The KTHNY theories predict that the isotropic liquid first undergoes a defect-mediated second-order transition into an orientationally quasi-ordered but translationally disordered state when $\tilde{K} = 1$ by disclination pair binding. At higher densities there is another second-order phase transition into a solid that has long-range orientational order and quasi-long range translational order, mediated by dislocation pair binding. The validity of this theory is still contested even for hard disks [267], and its applicability to systems where there is strong coupling between orientational and translational molecular degrees of freedom is questionable. Additionally, the basic theory needs to be modified to include three independent elastic moduli as opposed to only two in the case of six-fold rotational symmetry.

The observed change in S_4 as an isotropic liquid is slowly compressed is shown in Fig. 12.4 for both MD and MC runs. It is clearly seen that tetratic order appears in the system around $\phi \approx 0.7$ and increases sharply as the density is increased, approaching perfect order ($S_4 = 1$) at close packing. Throughout this run S_2 remains close to zero and thus no spontaneous nematic ordering is observed. It is important to note that superellipsoids are not perfect rectangles and have rounded sides. It is therefore not unexpected that they show less of a tendency toward tetratic (right-angle) ordering, and have the isotropic-tetratic (IT) transition at slightly higher densities. Additionally, the MD runs show more (correlated) variability due to the strong correlations between successive states (snapshots), and MD compressions lead to states with more pronounced defects. Therefore, we prefer to consider the MC results, other than at very high densities when we have to resort to MD studies. We have also performed runs decreasing the density of a random domino tiling, which has no nematic but has perfect tetratic order, and the resulting S_4 is also shown in the figure. Only a mild hysteresis is seen, especially for the MC runs, which would be indicative of a continuous IT transition, or at least a weakly discontinuous one. We note that we have never observed a phase boundary between a crystallized region and a disordered liquid, which would be indicative of a first-order phase transition.

Figure 12.5 shows $G_4(r)$ for a collection of states in the vicinity of the IT transition, thoroughly equilibrated using MC, on both a log-linear (lower densities) and a log-log (higher densities) scale. It is seen that there is a clear change in the long-range behavior of $G_4(r)$ as the density crosses above $\phi_c \approx 0.70$, from an exponential decay typical of an isotropic liquid, to a slower-than exponential decay at higher densities. The decay tails at higher densities are rather consistent with a power-law decay, and the fitted exponents η are shown in Fig. 12.6. It can be seen that η crosses the value $\eta_c = 1$ predicted by KTHNY theory when

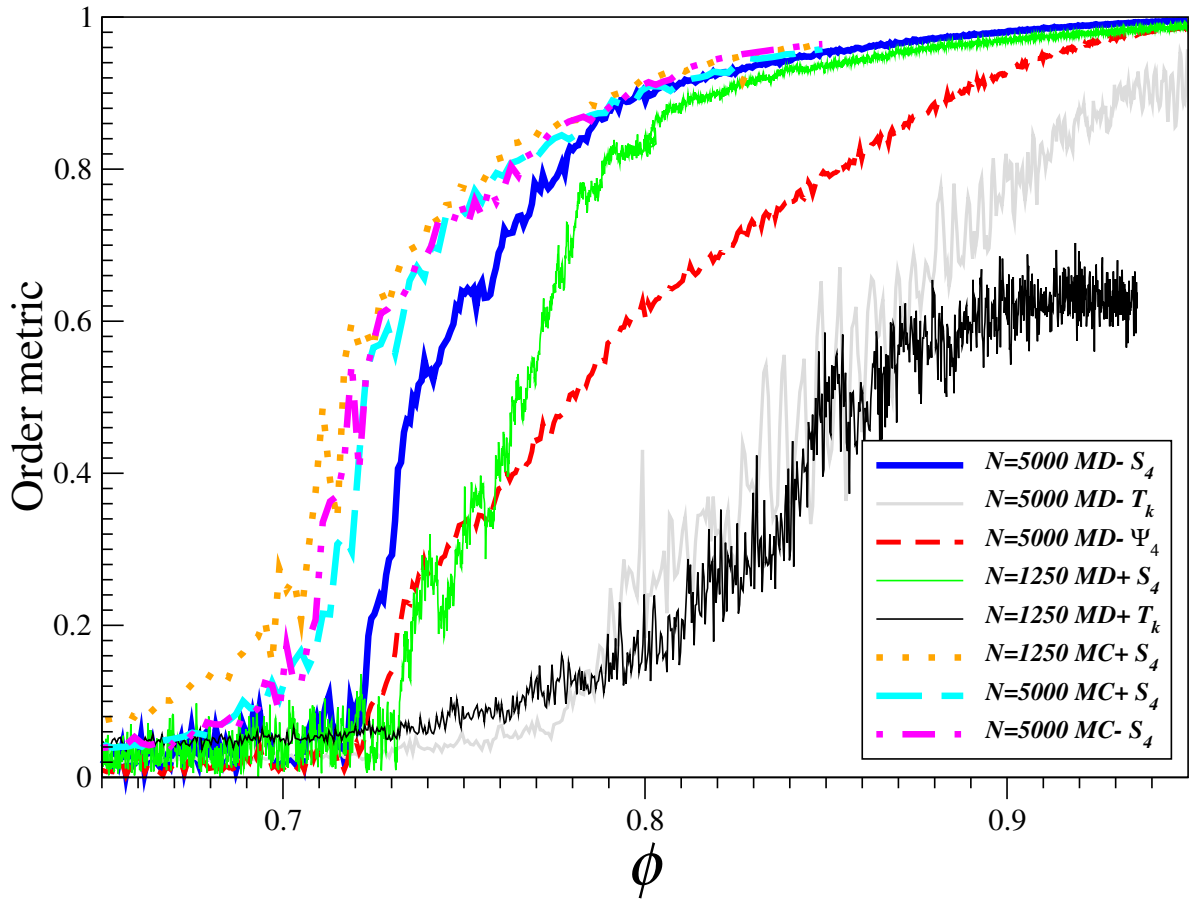

Figure 12.4: Values of the order metrics S_4 [c.f. Eq. (12.4)], T_k [c.f. Eq. (12.8)], and Ψ_4 [c.f. Eq. (12.7)] for snapshot configurations along compression (marked with a plus sign in the legend) and decompression (marked with a minus sign) MC and MD runs. A transition in S_4 is visible around $\phi \approx 0.7$, and a transition in T_k is indicated around $\phi \approx 0.8$. The hysteresis between compression and decompression runs is stronger the larger the system size and the larger the expansion rate γ, especially in MD runs.

$\phi \approx 0.71$, which is very consistent with the estimates of the location of the IT transition through the other methods above. It is not clear to us why the authors of Ref. [247] used the value of the exponent predicted by KTHNY theory for the bond-bond orientational order in the hard-disk system, $\eta_c = 1/4$, instead of $\eta_c = 1$. The somewhat higher values for S_4 for the system with $N = 1250$ relative to the system with $N = 5000$ particles are quantitatively well-explained by Eq. (12.6) using the values of η from Fig. 12.6. However, in order to unambiguously determine whether the decay in $G_4(r)$ is of a power-law form, larger systems should be studied and the scaling with system size determined carefully. Such studies are numerically very challenging and conflicting observations have been made even for the simplest case of hard disks. Unfortunately, for rectangles and superellipses, computational limitations presently prevent us from thoroughly equilibrating samples of more than 10000 particles.

12.3.3 Bond-Orientational Order

It is geometrically intuitive that orientational ordering is related to translational ordering, especially for elongated particles which need to coordinate orientational and translational degrees of freedom in order to avoid overlap at higher densities. From the observations above, we are motivated to look for translational order of the kind present in a random domino tiling. Looking at the centroids of the dominos themselves

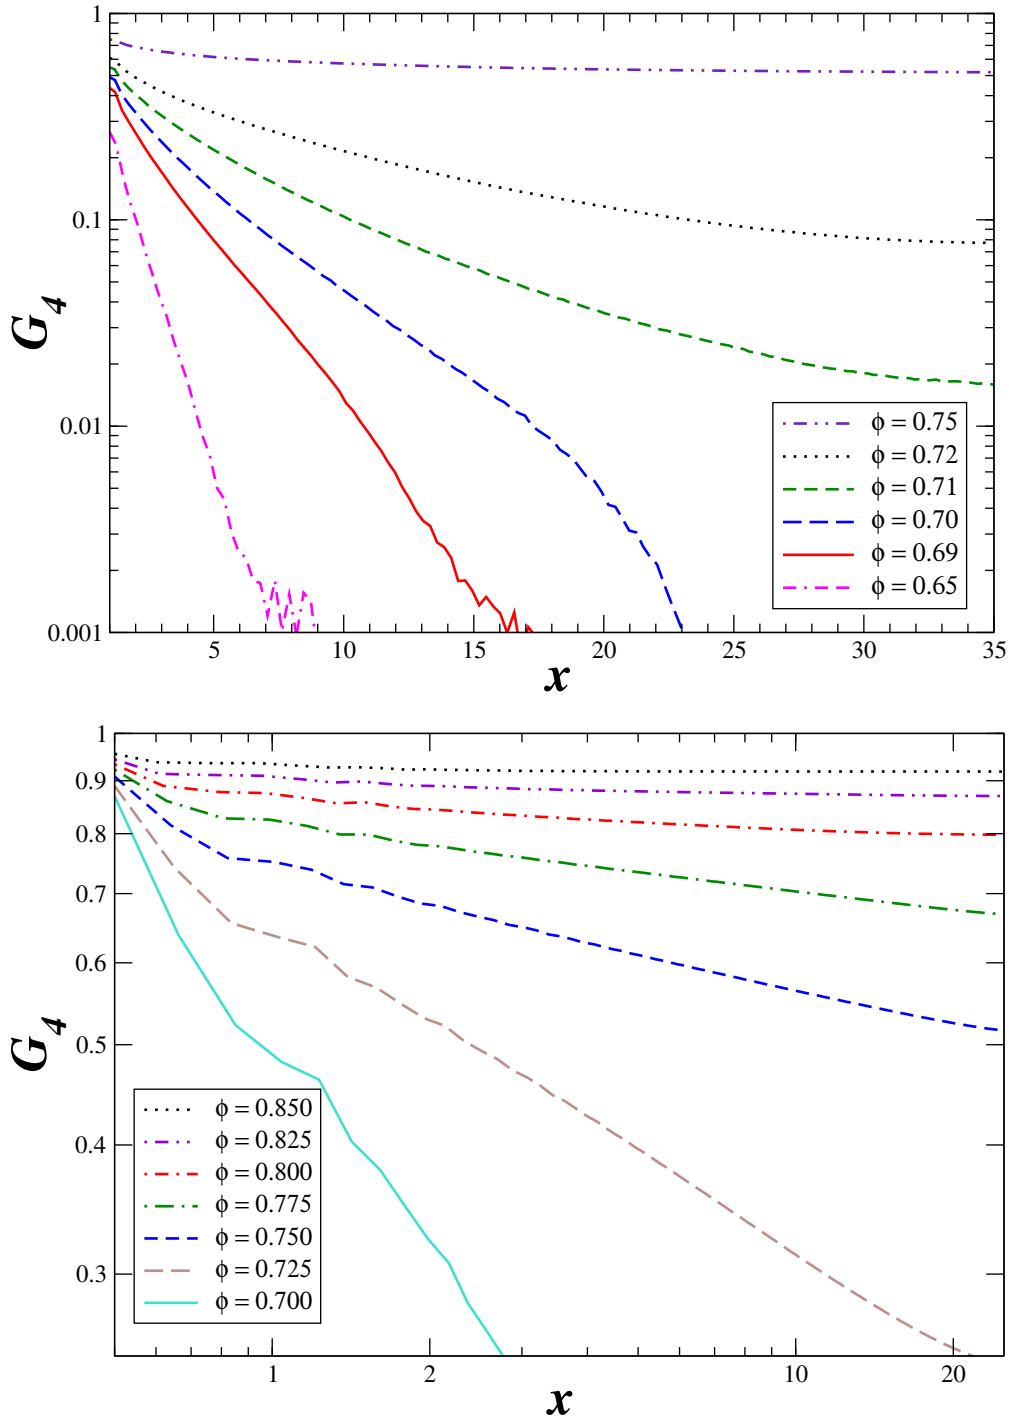

Figure 12.5: (*Left*): Log-linear plot of $G_4(x)$ for thoroughly equilibrated samples of $N = 10000$ particles, showing the decay of orientational ordering with distance $x = r/D$. The isotropic-tetratic transition occurs between $\phi = 0.69$ and 0.70, when the tail behavior of $G_4(r)$ changes from exponential (short-ranged) to slower-than-exponential. (*Right*): Log-log plot of $G_4(x)$ for equilibrated systems of $N = 5000$ particles, showing power-law decay indicative of quasi-long-range tetratic order. The fitted values of the power-law exponent are shown in Fig. 12.6.

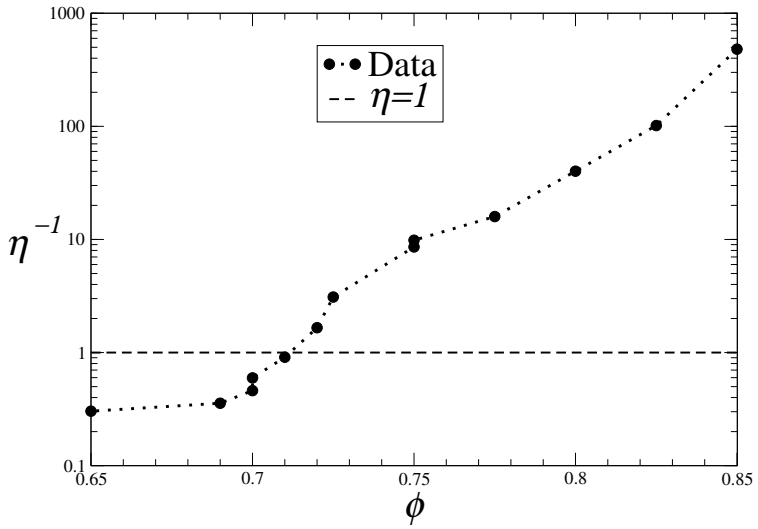

Figure 12.6: Log-linear plot of $1/\eta$, where η is the exponent of decay of $G_4(r)$ found by fitting the $G_4(x)$ data in Fig. 12.5 to a power-law curve, $G_4(x) = Cx^{-\eta}$.

does not reveal a simple pattern. However, if we split each rectangle into two squares and look at the centroids of the $2N$ squares, translational order will be manifested through the appearance of an underlying square-lattice.

From studies of other two-dimensional hard-particle systems, it has become clear that correlations between the orientations of the bonds connecting nearby particles can be (quasi) long-ranged. We have measured the four-fold bond-orientational order in the system of $2N$ half-domino squares via the scalar order metric Ψ_4 defined by

$$\Psi_4 \exp(i\omega) = \frac{1}{2N} \sum_{i=1}^{2N} \frac{1}{N_{\text{neigh}}} \sum_{j=1}^{N_{\text{neigh}}} \exp(4i\theta_{ij}), \tag{12.7}$$

where θ_{ij} is the orientation of the line connecting the centroids of nearby squares, and the sum for each particle is over the 4 particles whose centroids are closest, i.e., $N_{\text{neigh}} = 4$. Here ω is an angle giving the orientation of the global axes of four-fold alignment of the bonds. Other definitions of neighbors are possible. For example, Delaunay neighbours in the Voronoi tessellation of the point pattern formed by the centroids [260]; however, ambiguity is always present and the results are not very sensitive on the exact definition of neighbors so long as the average number of neighbors is close to 4. We find that for dominos, Ψ_4 qualitatively follows the behavior of S_4 very closely, as illustrated in Fig. 12.4, i.e., orientational ordering and bond-orientational ordering appear simultaneously, just as for the hard-square system [247].

12.3.4 Translational Order

Measuring (quasi)long-range translational order is more difficult than measuring orientational or bond-orientational ordering. Translational order is typically manifested via the appearance of a long-ranged periodic pattern for the centroids of the particles, in our case, the centroids of the half dominos. Such periodicity is most easily quantified by the Fourier transform of the square (half-domino) centroids, i.e., the structure factor

$$S(\mathbf{k}) = \frac{1}{N} \left| \sum_{j=1}^{N} \exp\left(i\mathbf{k}^T \mathbf{r}_j\right) \right|^2 .$$

In a translationally disordered state, $S(\mathbf{k})$ is of order one and decays to unity for large k. For long-ranged periodic systems, $S(\mathbf{k})$ shows sharp Bragg peaks at the reciprocal lattice vectors, while for quasi-long-range

order the peaks have power-law wings. It is however difficult to exactly determine when true peaks replace the finite humps that exist due to short-range translational ordering in the liquid state.

It would be convenient to have a scalar metric of translational order similar to S_4 for tetratic order. We use the averaged value of $S(\mathbf{k})$ over the first four Bragg peaks

$$T_k = \frac{1}{2N}\left[S(\frac{2\pi}{\tilde{a}}\mathbf{n}_\parallel) + S(\frac{2\pi}{\tilde{a}}\mathbf{n}_\perp)\right], \tag{12.8}$$

where $\tilde{a} = a/\sqrt{\phi}$ is the expected spacing of the underlying square lattice, and $\mathbf{n}_\parallel$ and $\mathbf{n}_\perp$ are two perpendicular unit vectors determining the orientation of the square lattice. For a liquid $T_k \approx 0$, and in the tiling limit $T_k = 1$ (however, defects can disrupt periodicity and significantly reduce the value of T_k below unity). When decompressing a prepared tiling, we already know $\mathbf{n}_\parallel = (1, 0)$ and it is best to use this known value. However, when compressing a liquid, we have no way of knowing the final orientation of the lattice. We have tried using the bidirector $\mathbf{n}_\parallel = \mathbf{n}_4$, as determined during the measurement of S_4. This method does not appear to work well because even small fluctuations in the director cause large fluctuations in T_k. Better results were obtained by using $\mathbf{n}_\parallel = (\cos\omega, \sin\omega)$ as determined from Eq. 12.7, as was done in Ref. [260]; however, the large fluctuations remained. We therefore chose to do a brute-force search for the unit vector $\mathbf{n}_\parallel$ that maximizes T_k, i.e., we rotate the presumed square lattice in small increments and find the optimal orientation.

A fundamental problem with T_k measured in this way is that it is smaller than unity away from the tiling limit even in a perfectly periodic domino system, since the two squares forming the domino are always closer than they would be if they were not glued together. In Fig. 12.4, we show the values of T_k along with S_4. It is seen that for the decompression run, T_k starts at unity and decays continuously until it apparently goes to zero around $\phi \approx 0.8$. We are therefore led to believe that there is a second transition from tetratic liquid to tetratic solid at $\phi \approx 0.8$. However, the transition is not sharp and the value of T_k is already too small to confidently distinguish it from zero. It is therefore possible that translational ordering appears as soon as tetratic ordering does, around $\phi \approx 0.7$, as would be the case if a (mild) first-order phase transition existed around this density.

In addition to reciprocal space $S(\mathbf{k})$, one can also look at the center-to-center-distance distribution function $g_2(r)$ for the squares (half dominos). However, quantitative analysis of $g_2(r)$ is made difficult because of oscillations due to exclusion effects and also due to the coupling to orientation. Instead of presenting such a one-dimensional pair correlation function, we present $g_2(\Delta x, \Delta y)$, which is simply the orientationally-averaged $g_2(\Delta x, \Delta y, \Delta\theta)$. In Figs. 12.7, 12.8 and 12.9 we show a snapshot of a system of $N = 5000$ dominos, along with the corresponding $g_2(\Delta x, \Delta y)$ and $S(\vec{k})$, for three densities, corresponding to an isotropic liquid, a tetratic liquid [i.e., a state with (quasi-) long range tetratic but only short-range translational order], and a tetratic solid [i.e., a state with (quasi-) long range tetratic and translational order]. For the $g_2(\Delta x, \Delta y)$ plots, we have drawn the expected underlying square lattice at that density. Note that $g_2(\Delta x, \Delta y)$ always has two sharp peaks corresponding to the square glued to the one under consideration in the dimer (domino).

12.3.5 Solid Phase

The investigations of ordering presented so far suggest that the domino system is very similar to the hard-square system [247] when dominos are split into two squares. Essentially, around $\phi \approx 0.7$ the centroids of the squares form a disordered square lattice and their orientations align with the orientation of the square lattice, and the lattice itself orders over large distances at higher densities, leading to a tetratic solid. However, a fundamental question remains concerning the way in which the squares are joined to form the dominos in the thermodynamically-stable solid phase. There are two likely possibilities: The tiling shows (translational) ordering itself, or the tiling is "random". In the context of a discrete system like domino tilings, the concept of a random tiling is mathematically well-defined in terms of maximizing entropy [263, 268]. This random tiling has a positive degeneracy entropy $0.58313k_B$, unlike ordered tilings such as the nematic tiling (in which all dominos are aligned).

Our compressions of isotropic liquids have invariably led to apparently disordered domino tilings upon spontaneous "freezing", albeit with some frozen defects. This suggests that the disordered tiling has lower free energy than ordered tilings. However, it is also possible that the disorder is simply dynamically trapped when the tetratic liquid freezes. In fact, starting a decompression run from an aligned nematic tiling shows

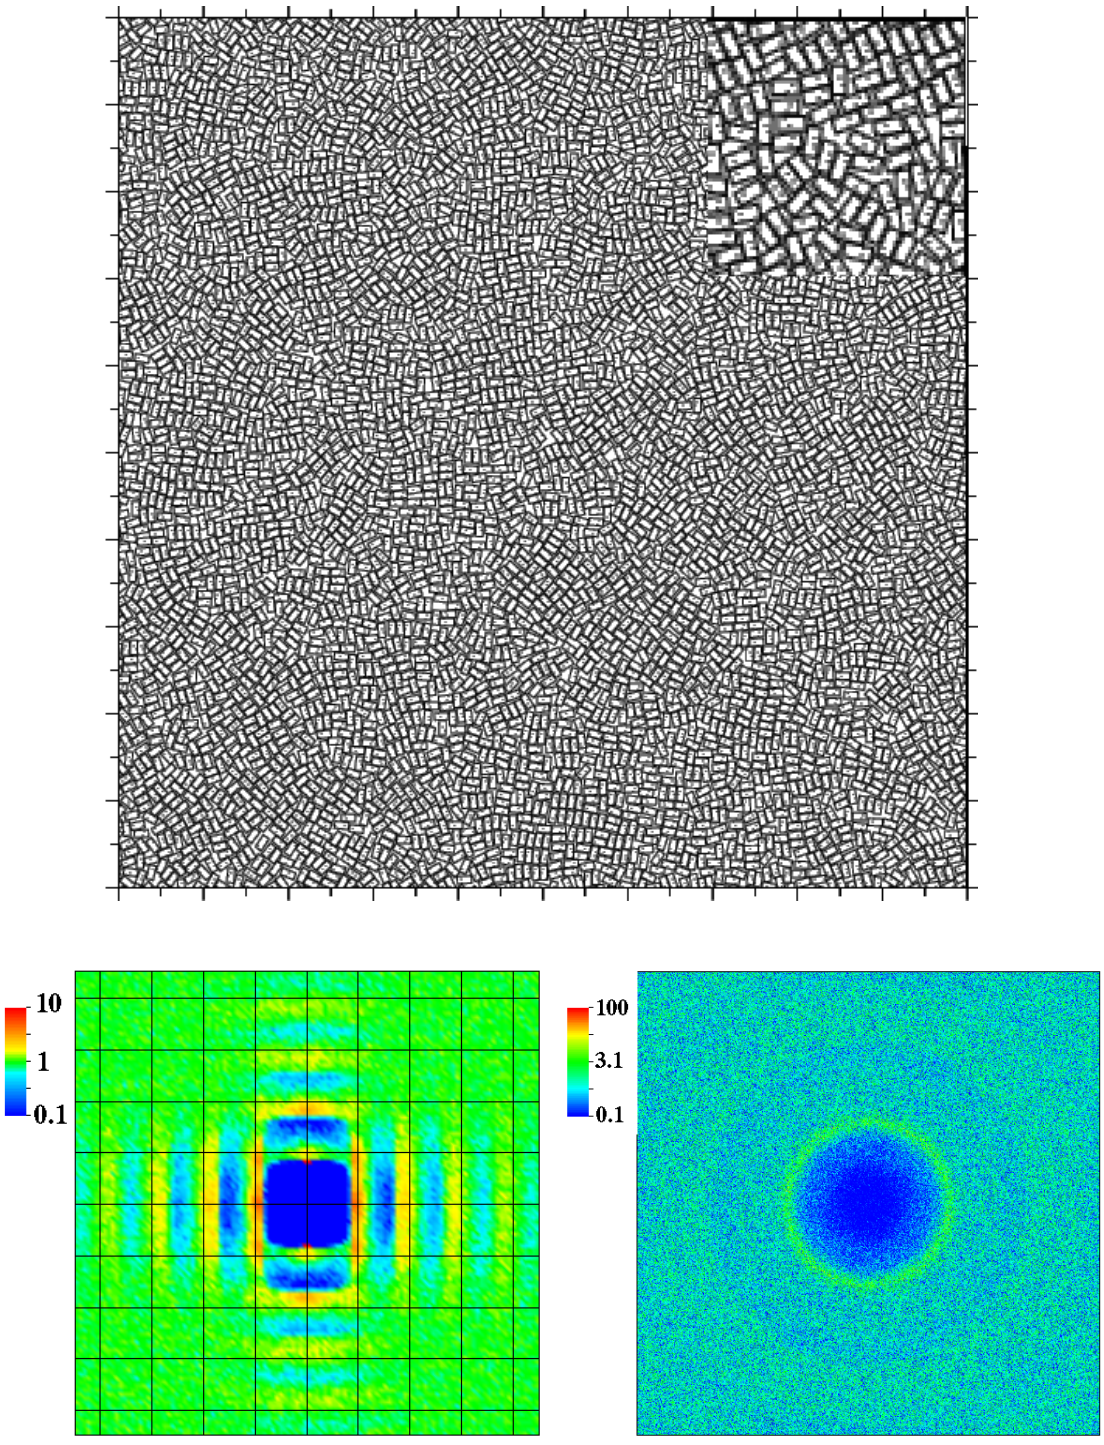

Figure 12.7: A snapshot configuration of a system of $N = 5000$ dominos at $\phi = 0.7$ (top) with inset with threefold magnification showing local packing structure, along with $g_2(\Delta x, \Delta y)$ overlayed over the underlying square lattice (bottom left) and $S(\vec{k})$ (bottom right), obtained after splitting each domino into two squares. It is clear that the system is isotropic from the rotational symmetry of $S(\vec{k})$. Only short-range order is visible in $g_2(\Delta x, \Delta y)$, confirming that this is an isotropic liquid.

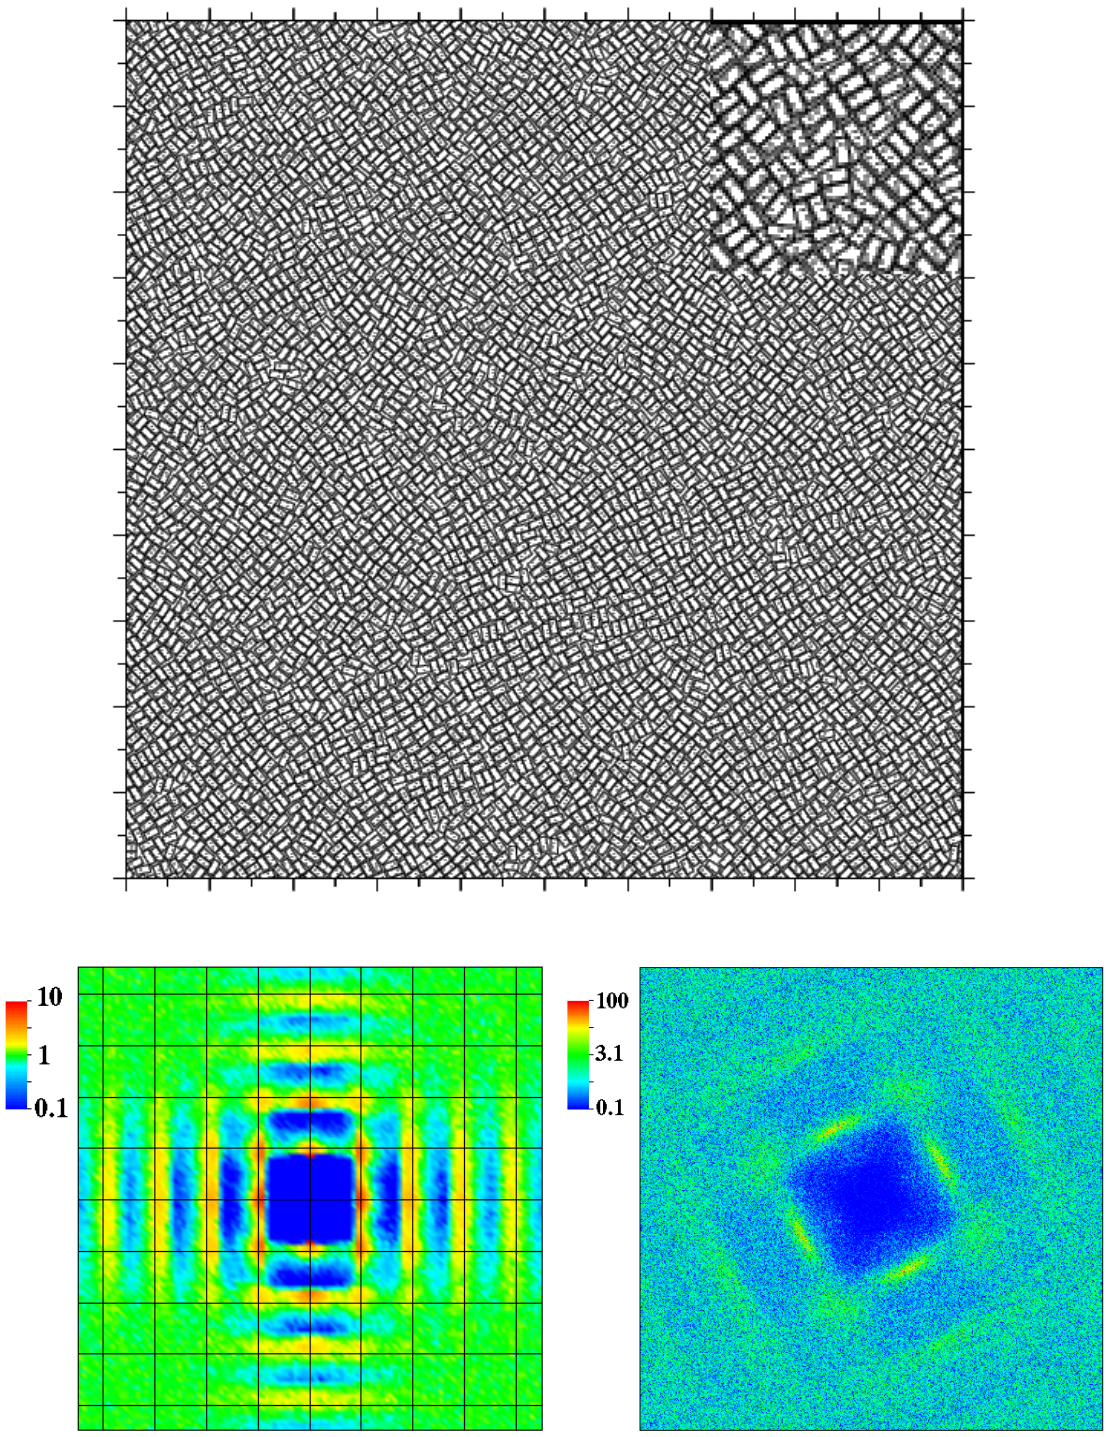

Figure 12.8: A system of $N = 5000$ dominos as in Fig. 12.7 but at $\phi = 0.750$, which shows a tetratic liquid phase. Four-fold broken symmetry is seen in $S(\vec{k})$, but without pronounced sharp peaks. The range of ordering in $g_2(r)$ has increased, but still appears of much shorter range than the size of the system, as seen clearly in the plot of the actual domino configuration. It is interesting that $g_2(\Delta x, \Delta y)$ is very anisotropic, being much stronger to the side of a square relative to its diagonals. No phase boundary characteristic of first-order transitions is visible.

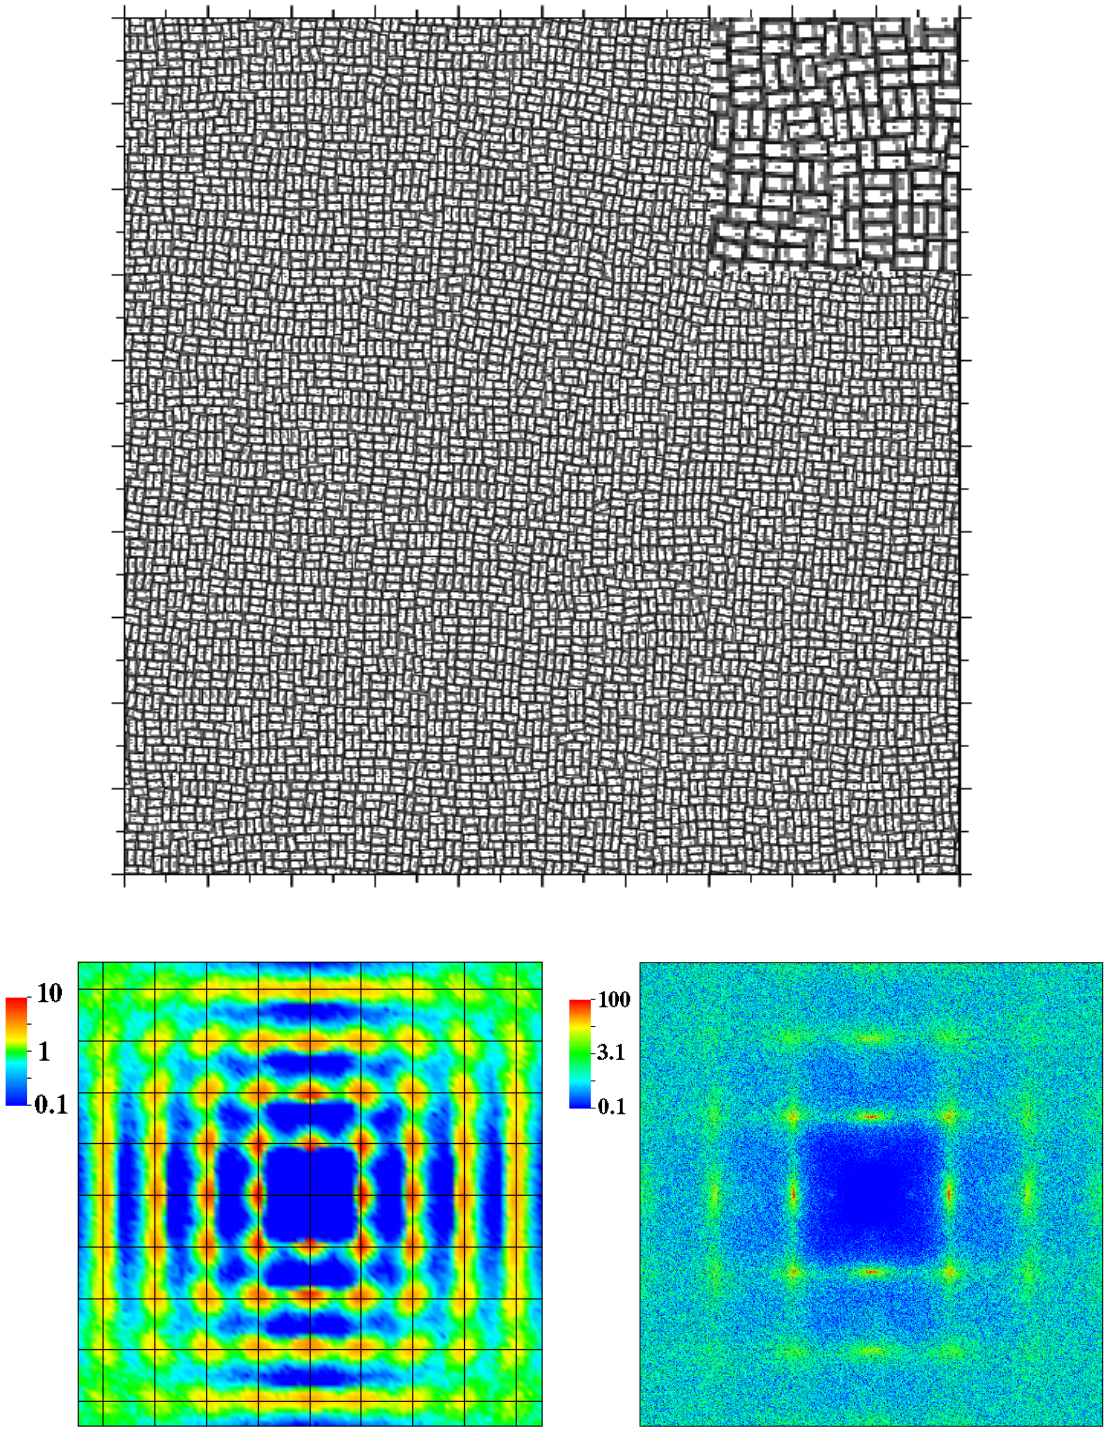

Figure 12.9: A system of $N = 5000$ dominos as in Figs. 12.7 and 12.8 but at $\phi = 0.825$. The structure factor shows sharp peaks (maximum value is above 10) on the sites of a (reciprocal) square lattice, and $g_2(r)$ shows longer-ranged translational ordering, indicating a solid phase. Visual inspection of the configuration confirms that the translational ordering spans the system size and shows some vacancies consisting of only a single square (half a particle).

that the tiling configuration is preserved until melting into a tetratic liquid occurs somewhat below $\phi \approx 0.8$. This is demonstrated in Fig. 12.10, where both S_4 and S_2 as well as T_k are shown along a decompression run starting with both a disordered and an ordered tiling. It is seen that S_2 drops sharply around $\phi \approx 0.8$ while S_4 remains positive until $\phi \approx 0.7$, clearly demonstrating the thermodynamic stability of the tetratic liquid phase in the intermediate density range. Subsequent compression of this liquid would lead to a disordered tiling without any trace of the initial nematic ordering.

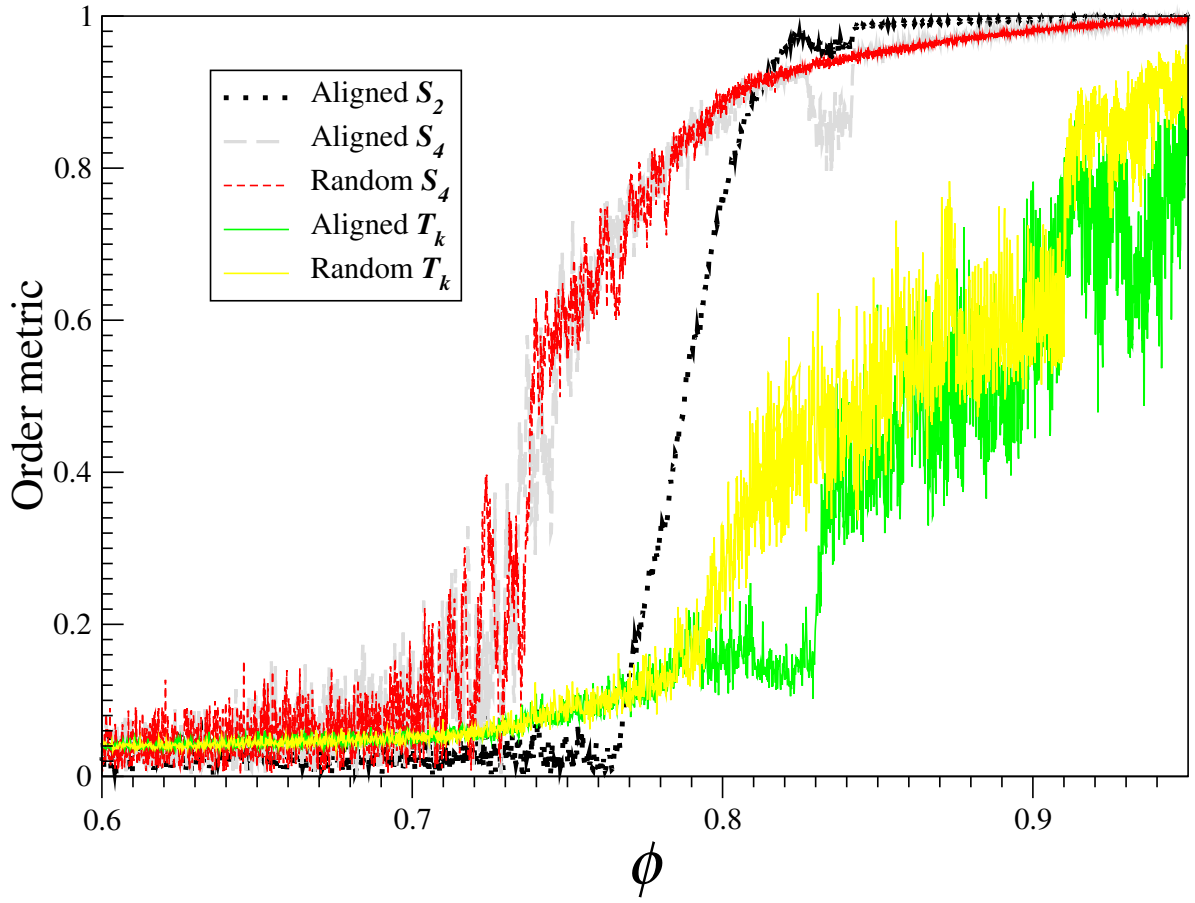

Figure 12.10: Nematic, tetratic, and translational order metrics as a domino tiling in which all rectangles are aligned is slowly decompressed from close-packing. The nematic crystal spontaneously realigned to a different orientation of the director from the starting one at around $\phi \approx 0.84$, causing some fluctuations and a drop in T_k which are likely just a finite-size (boundary) effect.

It is intuitive to expect that the free-volume contribution to the free energy is minimized for ordered tilings at high densities. However, we also expect that solid phase is ergodic in the sense that transitions between alternative tiling configurations will occur in long runs of very large systems, so that in the thermodynamic limit the space of all tilings will be explored. This amounts to a positive contribution to the entropy of the disordered tiling due to its degeneracy, and it is this entropy that can thermodynamically stabilize the disordered tiling even in the close-packed limit. A closer analysis similar to that carried for hard-disk dimers in Refs. [248, 249] is necessary. In particular, including collective Monte Carlo trial moves that transition between different tiling configurations, as well as relaxation of the dimensions of the unit cell (important for smaller solid systems), is important. Furthermore, only free-energy calculations can determine the free-volume entropies of various tiling patterns. The algorithm presented in Chapter 6 could be used to measure those free energies, however, this requires developing robust overlap potentials for one particle contained within another (its bounding neighborhood), which has not yet been achieved even for ellipses, as discussed

in Chapter 2.

We conjecture that, just like the hard-disk dimer system, the hard-square dimer system has a thermodynamically stable nonperiodic solid phase. Even if this conjecture is false and the nematic phase is thermodynamically favored at sufficiently high densities, our simulations indicate that the dynamics is glassy; that is to say, the system gets trapped in disordered tilings even for very long runs, unlike the hard-disk system, where crystallization occurs easily. As in the hard dumbbell (fused hard-disk dimers) system, we expect that for aspect ratios close to, but not exactly, two, the nonperiodic solid will be replaced by a nematic (and possibly periodic) phase at the highest densities [249, 269]. This is because reaching the maximal density $\phi = 1$ seems to require aligning the rectangles. It is interesting, however, that at least for rational, and certainly for integer aspect ratios such as $\alpha = 3$, there is the possibility of disordered solid phases being stable even in the close-packed limit. On physical grounds we expect the phase diagram to vary smoothly with aspect ratio, rather than depending sensitively on the exact value of α, and the determination of the phase diagram around $\alpha = 1$ and $\alpha = 2$, even if only qualitatively, is an important challenge for future research.

Accepting for a moment the existence of a nonperiodic solid phase, it remains to verify that the compressed systems we obtain in our simulations are indeed like (maximal entropy) random tilings of the plane with dimers. This is hard to do rigorously, as it requires comparing all correlation functions between a random tiling and our compressed systems. Figure 12.11 shows a visual comparison of a random tiling of a large square, generated using random spanning trees by a program provided to us by the authors of Ref. [262], and a system of superellipses compressed to $\phi = 0.95$ (close to the achievable maximum for our MD program for such high superellipse exponents). While the translational ordering in the compressed solid is clearly not perfect as it is for the true tiling, visual inspection suggests close similarity between the the local tiling patterns of the two systems. Note that the primary type of defect that we observe are single-square vacancies, i.e., half-dominos missing from the true tiling of the plane. The number of vacancies observed during slow compressions is small and we do not expect it to affect thermodynamic properties significantly. In Fig. 12.12, we show $g_2(\Delta x, \Delta y)$ for the true tiling, along with the difference in g_2 between the true tiling and the compressed solid. Here we do not split the rectangles into two squares, i.e., the figure shows the probability density of observing a centroid of another rectangle at $(\Delta x, \Delta y)$ given a rectangle at the origin oriented with the long side along the x axis. It can be seen that there is a close match between the random tiling and the compressed solid, at least at the two-body correlation level.

12.4 Conclusions

The results we presented highlight the unusual properties of the simple hard-rectangle system when the aspect ratio is $\alpha = 2$, hopefully stimulating further research into the hard-rectangle system. For square dimers (dominos), in addition to the expected low-density isotropic liquid phase, a stable tetratic liquid phase is clearly observed, in which there is four-fold orientational ordering but no translational ordering. A tetratic solid phase closely connected to random domino tilings is observed and we conjecture that it is thermodynamically stabilized by its positive degeneracy entropy. The transitions between the phases are consistent with a KTHNY-like sequence of two continuous transitions. If this is indeed the case, then the hard dimer system provides an excellent model for the study of continuous transitions, with a rather large gap in density between the two presumed transitions $\Delta\phi \approx 0.1$, unlike the hard-disk system. Random jammed packings of rectangles seem to be translationally ordered, similar to the behavior for disks [72] but unlike spheres which can jam in disordered configurations [58]. However, unlike disks, the systems of rectangles show orientational disorder, once again illustrating the geometric richness of even the simplest hard-particle models.

Further investigations are needed for the domino system to conclusively determine its phase behavior. Improved MC with collective moves that explore multiple tilings, as well as allow for relaxation of the boundary conditions, should be implemented. Additionally, the free energies of the different phases should be computed so that the exact locations of the phase transitions could be identified. The final goal is to completely characterize the phase diagram of the hard rectangle system in the $\alpha - \phi$ plane, as has been done, for example, for diskorectangles [266] and ellipses [271]. In addition to nematic and smectic phases, novel liquid crystal phases with tetratic order may be discovered.

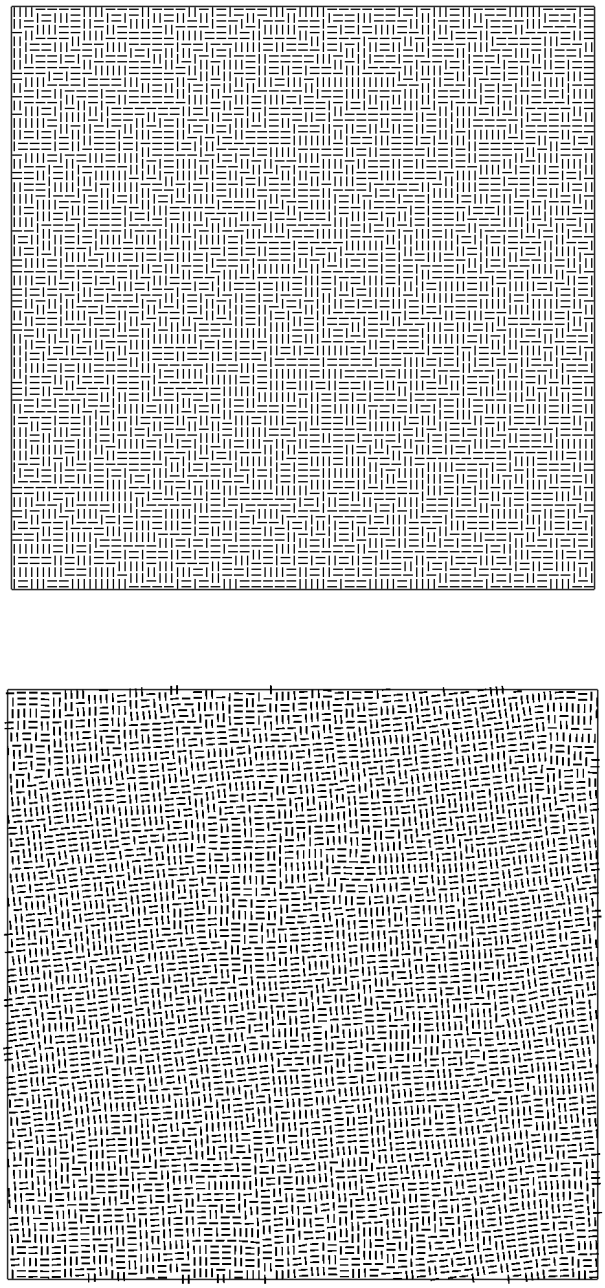

Figure 12.11: A comparison between a true random tiling of a square with dominos [262] (top), and the unit cell of a system of $N = 5000$ superellipses with exponent $\zeta = 7.5$ slowly compressed from isotropic liquid to $\phi = 0.95$ (bottom). The compressed system is not a perfect tiling due its lower density and frozen defects, as well as the rounding of the superellipses relative to true rectangles. Therefore at large scales the two systems look different. However a closer local examination reveals similar tiling patterns in the two systems, typical of "random" tilings.

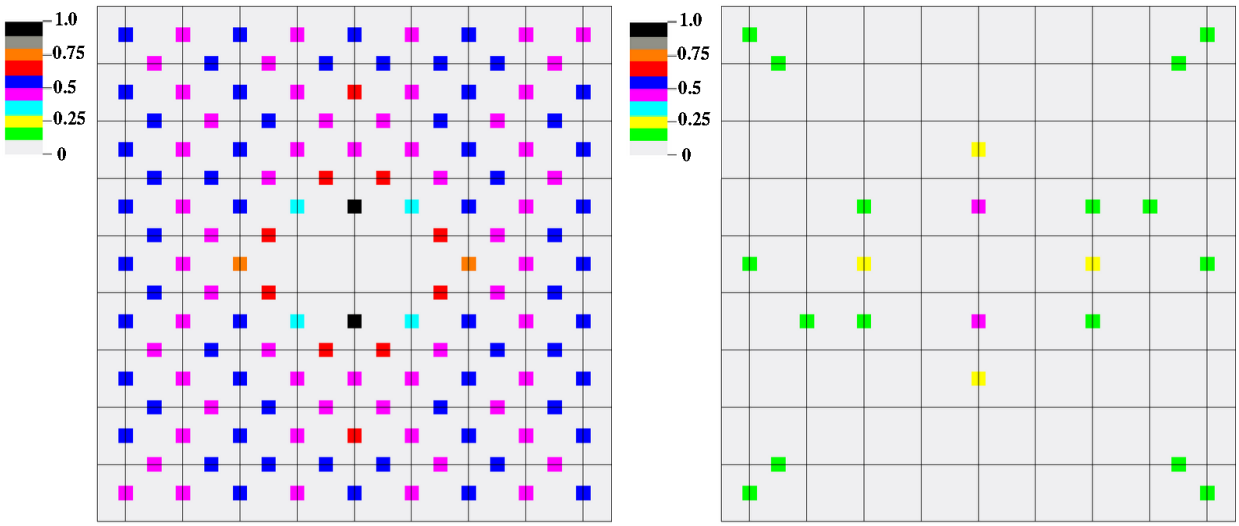

Figure 12.12: (*Left*): Center-center pair correlation function $g_2(\Delta x, \Delta y)$ for the perfect random tiling in Fig. 12.11. This g_2 is a collection of δ-functions whose heights can also be calculated exactly [270] (the calculation is nontrivial and we have not performed it). We have normalized g_2 so that the highest peaks have a value of one. (*Right*): The absolute value of the difference between $g_2(\Delta x, \Delta y)$ for the two systems shown in Fig. 12.11, shown on a coarse-enough scale so that the broadening of the peaks due to thermal motion is not visible. The color table used in this figure is discrete in order to highlight the symmetry and hide small fluctuations due to finite system size. The difference in g_2 is almost entirely within the smallest interval of the color table (less than 0.1, gray), with only some peaks showing differences up to 0.25.

Chapter 13

Configurational Entropy of Binary Hard-Disk Glasses

In this chapter we study the thermodynamics of a binary hard disk mixture where the ratio of disk diameters is $\kappa = 1.4$ [57, 83]. We use the algorithm from Chapter 6 to calculate the free volume entropy of glassy configurations and from these entropies we obtain the configurational entropy (degeneracy) as a function of density. We find that the configurational entropy of the glasses near the kinetic glass transition is very close to the mixing entropy, suggesting that the degeneracy is zero only for the phase-separated crystal. We demonstrate that there is no ideal glass transition in a binary hard-disk mixture by explicitly constructing an exponential number of jammed packings with densities spanning the spectrum from the accepted "amorphous" glassy state to the phase-separated crystal. Thus the configurational entropy cannot be zero for an ideal amorphous glass, presumed distinct from the crystal in numerous theoretical and numerical estimates in the literature. This objection parallels our previous critique of the idea that there is a most-dense random (close) packing for hard spheres [58]. We also perform free energy calculations to determine the equilibrium phase behavior of the system, and predict a first-order freezing transition at a density below the kinetic glass transition, which appears strongly kinetically suppressed and is not observed directly. New simulation techniques are needed in order to gain more complete understanding of the thermodynamic and kinetic behavior of the binary-disk mixture, and in particular of the demixing process during crystallization.

13.1 Introduction

Understanding the glass transition in dense or supercooled liquids remains one of the challenges of condensed matter physics. In particular, considerable effort has been directed at identifying the cause of the dramatic slowdown of the dynamics in the vicinity of the kinetic glass transition, as evidenced in a decrease of the diffusion coefficient and an increase in relaxation times. One possibility is that a thermodynamic transition different from the usual liquid-solid transition underlies the kinetic one. One scenario originally suggested by Adam and Gibbs [272] relates the slow diffusion to a vanishing of the number of alternative configurations available to the liquid, leading to an *ideal thermodynamic glass transition* when the liquid has no choice but to remain trapped in one of few glassy configurations. An important basic assumption in these considerations is that crystalline configurations, which are thermodynamically favored, are kinetically inaccessible and therefore the liquid is restricted to exploring "amorphous" configurations. In particular, the term amorphous has become implicitly attached to the term glass, and crystalline configurations have been assumed to be qualitatively different from glassy ones. In this chapter, we study a specific model glass former, namely, a binary hard disk mixture, and show that, for this model, the presumed "ideal glass" is in fact a phase-separated crystal, and that that there is no special amorphous (random) state, but rather a continuum of states from the most disordered one to the most ordered one [58].

13.1.1 Inherent-Structure Formalism

An inherent-structure formalism was proposed by Stillinger and Weber and has since been used extensively in the analysis of the thermodynamics of supercooled liquids [29, 273]. The "inherent-structures" of hard-particle systems are in fact (collectively) *jammed packings* [75], which are mechanically stable packings where the particles are trapped in a static configuration despite thermal or external agitation. For soft-particle systems, an essential quantity in this thermodynamic analysis is the number of distinct energy minima (basins) with a given energy per particle. For hard-particle systems this becomes the number of distinct jammed packings $N_g(\phi_J) = \exp\left[N s_c(\phi_J)\right]$ with jamming packing fraction (density) ϕ_J, where $s_c(\phi_J)$ is the *configurational entropy*, or *degeneracy*, per particle. It is assumed that the liquid remains in the vicinity of these jamming basins for long periods of time, jumping from one basin to another as it explores the available configuration space. Denser packings are favored in terms of their free-volume; the most favored one being the crystal of density $\phi_{\max}$. However, it is reasonable to assume that the degeneracy $s_c(\phi_J)$ decreases with increasing ϕ_J. The liquid achieves minimum free energy by trading off degeneracy for free volume, so that at a given density ϕ it predominantly samples glasses with jamming density $\hat{\phi}_J(\phi)$. The theory of cooperatively rearranging regions developed by Adam and Gibbs [272] proposes that the structural relaxation time in the metastable liquid is on the order of

$$\tau(\phi) \sim \exp\left\{\frac{C}{T s_c\left[\hat{\phi}_J(\phi)\right]}\right\},$$

and therefore diverges at the ideal glass transition.

We can approximate the *free-volume* contribution to the free-energy (per particle) of a glass close to the jamming point with Eq. (6.2),

$$f_g\left(\phi, \phi_J\right) = -d \ln\left(1 - \frac{\phi}{\phi_J}\right) - f_J\left(\phi_J\right).$$

The (total) volume of configuration space corresponding to jamming density ϕ_J is a sum over all of the $s_c(\phi_J)$ basins, and therefore the contribution to the free energy from the glasses with jamming density ϕ_J is approximately

$$f\left(\phi, \phi_J\right) = f_g\left(\phi, \phi_J\right) + s_c(\phi_J).$$

At a given density ϕ the jamming density that maximizes f is the one that dominates the thermodynamic integrals, and it is found from the solution $\hat{\phi}_J(\phi)$ of the equation

$$\left.\frac{\partial f}{\partial \phi_J}\right|_{\phi_J = \hat{\phi}_J(\phi)} = \left[-\frac{d}{\phi\left(1 - \frac{\phi}{\phi_J}\right)} - \frac{\partial f_J\left(\phi_J\right)}{\partial \phi_J} + \frac{\partial s_c\left(\phi_J\right)}{\partial \phi_J}\right]_{\phi_J = \hat{\phi}_J(\phi)} = 0. \tag{13.1}$$

As expected, the pressure of the metastable glass is equal to the pressure of just one of the jamming basins and is not affected by the fact that the packing explores multiple (statistically identical) basins

$$p = \phi \frac{df}{d\phi} = \phi \frac{\partial f}{\partial \phi_J}\frac{\partial \phi_J}{\partial \phi} + \phi \frac{\partial f}{\partial \phi} = \phi \frac{\partial f}{\partial \phi} = \frac{d}{\left(1 - \frac{\phi}{\phi_J}\right)}.$$

The configurational entropy $s_c(\phi_J)$ must vanish above some density $\phi_J^{\max}$, if nothing else than because $\phi_J^{\max} \leq \phi_{\max} = \pi/\sqrt{18}$. The conjectured *ideal glass state* corresponds to the point where the number of available basins becomes subexponential, that is, $s_c(\phi_J^{\mathrm{IG}}) = 0$. The usual assumption in the literature is that $s_c(\phi_J)$ is an inverted parabola and that f_J is constant, and this assumption gives a monotonically increasing $\hat{\phi}_J(\phi)$ [274, 50, 273]. At densities above an *ideal glass transition* density ϕ_{IG}, defined via $\hat{\phi}_J(\phi_{\mathrm{IG}}) = \phi_J^{\mathrm{IG}}$, the liquid becomes permanently trapped in the ideal glass state. A crucial unquestioned assumption has been that $\phi_J^{\mathrm{IG}} < \phi_{\max}$, i.e., that there is a gap in the density of jammed states between the amorphous and crystal ones. We will explicitly show that this assumption is flawed for the binary hard-disk mixture we study, and suggest that this is the case in other similar models, contrary to numerous estimates for ϕ_J^{IG} in the literature [50, 181, 180, 275, 183, 276].

13.2 Equilibrium Phase Diagram

In this section we use molecular dynamics to observe the equation of state for monodisperse and bidisperse hard disk systems and also to calculate the excess free energy per particle relative to the ideal gas at different densities. We use the MD data to estimate the location of the freezing transition in a binary hard disk mixture. We predict that at a freezing density $\phi_F \approx 0.775$, a crystal of density 0.8415 composed of predominantly large particles should start precipitating from the liquid mixture. Our study here is similar to that carried out in significant detail for soft disks (interacting via an inverse 12-th power potential) in Ref. [51].

13.2.1 Monodisperse Hard-Disk Systems

Whether the liquid-solid transition for the monodisperse hard disk system is a continuous (second-order) transition or a discontinuous first-order transition is still disputed [267]. We will not try to resolve this question here, however, we must briefly examine the thermodynamics of monodisperse disks as this will be necessary in order to study mixtures. In Fig. 13.1 we show the equation of state (EOS) for monodisperse hard disks, obtained through molecular dynamics at different particle growth rates γ. A more detailed description of the procedure is in Section 1.2 and an analogous plot for three dimensions is given in Fig. 1.3.

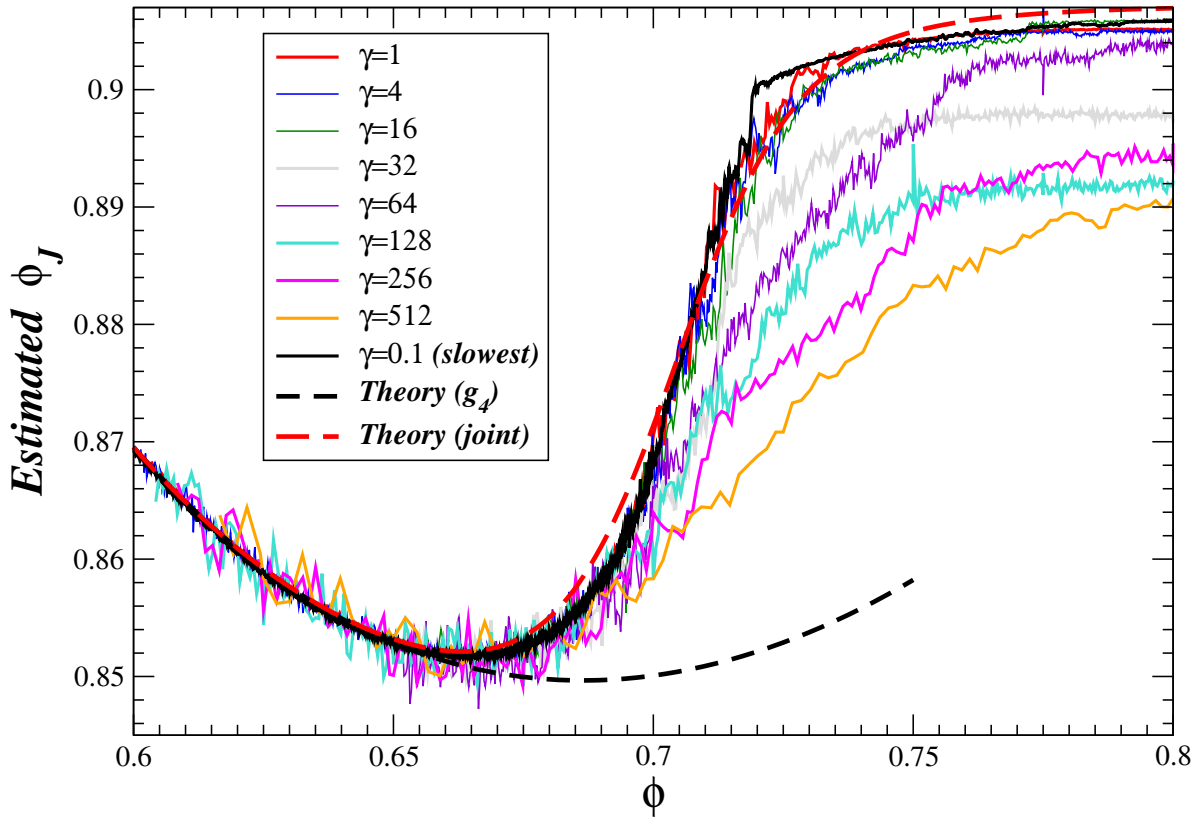

Figure 13.1: The equation of state for the monodisperse system of $N = 4096$ hard disks, as is shown in Fig. 1.3 for hard spheres. Molecular dynamics runs are started with from isotropic liquid and the density increased slowly at different particle growth rates γ, as shown in units of 10^{-6} in the legend. The g_4 equation of state for the liquid [c.f. Eq. (2) in Ref. [277]] and the joint liquid-crystal EOS from Ref. [277] are also shown for comparison. It is seen that for $\gamma < 10^{-5}$, the system has sufficient time to equilibrate for all densities shown and the pressure follows the true equilibrium EOS. We do not show the equivalent curves for runs starting with a perfect triangular crystal and decreasing the density (i.e., $\gamma < 0$), since there is virtually no hysteresis observed from the EOS obtained by increasing the density.

Unlike in three dimensions, there is no clear discontinuity between the EOS of the isotropic liquid and the triangular crystal. Slowing the rate of density increase allows us to find the true equilibrium EOS, as demonstrated by the fact that the observed EOS in Fig. 1.3 barely changes even though γ is decreased by more than an order of magnitude. The transition between the liquid and solid phases occurs in the density range $\phi \approx 0.70 - 0.72$, which would also be the best estimate for the coexistence region assuming that the transition is first order. Even if the transition is first-order, however, the change in entropy between the liquid and solid is very small and therefore for the purposes of free-energy calculations we can assume that there is a continuous transition, i.e., that the free energy per particle $f_{\rm mono}(\phi)$ is a unique and smooth function of density. We will therefore not explicitly distinguish between the liquid and solid (and possible hexatic phase) of the monodisperse hard disk system, but rather consider them as a single phase.

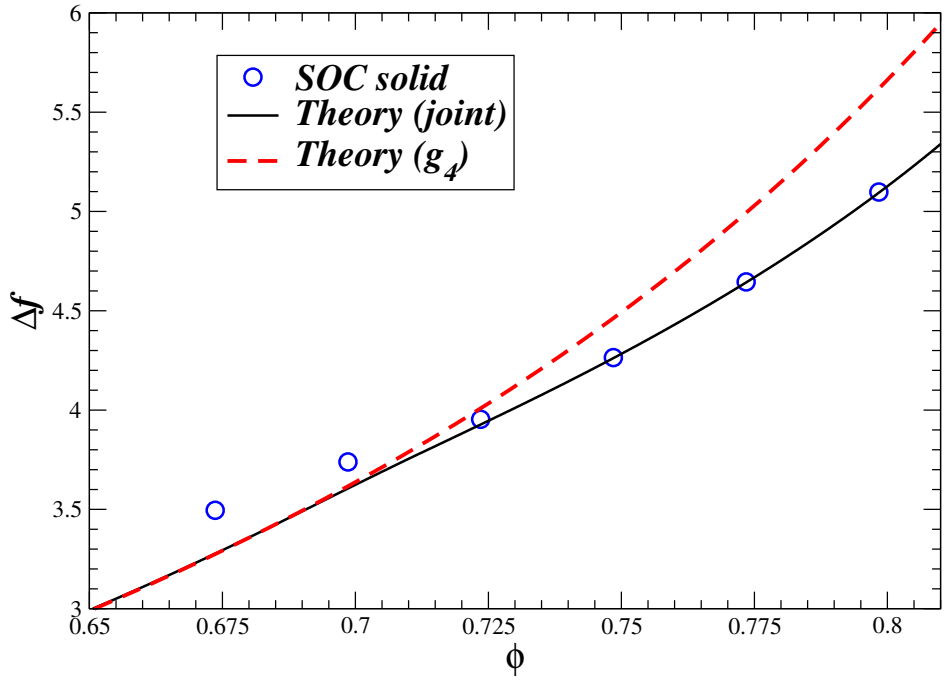

Figure 13.2: The excess free energy per particle Δf for a monodisperse system of hard disks, obtained through integration of the joint (global) liquid/solid EOS proposed in Ref. [277], through the BCMD algorithm for SOC systems, and from the most accurate EOS from liquid-state theory [277]. It is seen that in the solid phase, $\phi > 0.72$, the free energy of the SOC model closely matches that obtained by assuming a continuous EOS, demonstrating that the entropy jump between the liquid and solid phases at the transition, is too small to be measured, if it exists at all (see also Fig. 13.12).

The numerical EOS $p_{\rm mono}(\phi)$ is well-fitted by the semi-empirical joint liquid-solid EOS proposed in Ref. [277]. In Fig. 13.2 we show the free energy obtained by using the BCMD algorithm from Chapter 6 with bounding cells of diameter twice larger than the diameter of the disks, i.e., $\Delta\mu = 1$. As discussed in more detail in Section 6.3.5, the BCMD algorithm calculates the free energy of a single-occupancy cell (SOC) solid in which each particle is restricted to remain within its bounding cell. For sufficiently high densities this gives a very good approximation to the free-volume contribution to the free-energy. For the crystal the configurational entropy is identically zero and therefore the free-volume term is the only contribution to the thermodynamic free energy. The figure shows that $f_{\rm mono}(\phi)$ is indeed (nearly) continuous when going from the liquid to the solid state. We note that one can avoid analytical approximations completely and obtain $f_{\rm mono}(\phi)$ numerically with high accuracy by simply integrating the pressure from the low-densities (where a low-order virial expansion is accurate), using the pressure from an MD simulation with sufficiently small γ.

13.2.2 Binary Mixtures of Hard Disks

We study a binary mixture of disks with a third (*composition* $x_B = 1/3$) of the disks having a diameter (size *dispersity*) $\kappa = 1.4$ times larger than the remaining two thirds ($x_A = 2/3$). Bidisperse disk packings with this aspect ratio and $x_A = x_B = 1/2$ have been studied[1] as model glass formers [51]. For this κ, it is believed that the high-density phase is a phase-separated crystal [61]. It can be proved that the highest density achievable with two disks of size ratio $\kappa < 1.348$ is the same as for monodisperse disks, $\phi_{\max} = \pi/\sqrt{12}$ [278], although we are not aware of any proof that this highest density is only achievable in phase-separated configurations. For large size dispersity, denser packings exist where the small and large disks are mixed [61, 67]. For small size dispersity, a substitutional triangular crystal, in which the large and small disks are randomly mixed, will be thermodynamically favored over the phase-separated crystal at intermediate densities because of its higher degeneracy entropy [279, 280]. We will assume here that for $\kappa = 1.4$ the crystal phase is a phase-separated mixture of monodisperse triangular crystals and neglect any solubility of one type of disks into the crystal of the other type, since such solubility is expected to be negligible due to the large difference in size between the small and large disks [280].

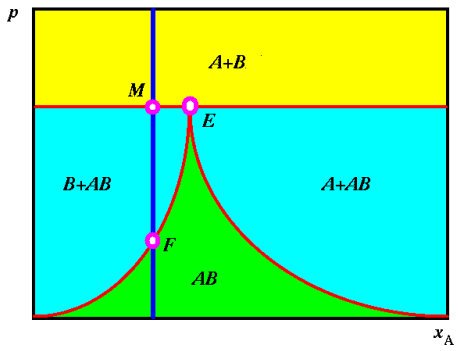

Figure 13.3: An schematic illustration of the Eutectic phase diagram assumed to apply to binary mixtures of hard disks with dispersity $\kappa = 1.4$. The monodisperse phases are denoted by A and B, and the liquid mixture by AB. The horizontal axes is the composition x_A, and the vertical axes is the pressure P. The location of the particular composition we study here is marked by a blue vertical line, along with the freezing point F, the melting point M, and the eutectic point E.

We also expect that the full (over all compositions) equilibrium phase diagram for this size dispersity will be of the eutectic type [280], as schematically illustrated in Fig. 13.3. At low pressures, the equilibrium phase is a mixed isotropic liquid AB. Upon increasing the pressure a freezing point is reached, and phase separation begins by precipitation of a monodispersed phase B composed of large particles in mechanical equilibrium with the surrounding liquid AB (depleted in large particles). When the pressure exceeds the melting point complete phase separation between the large and small disks occurs and the mixed liquid phase ceases to exist. It is important to note that the equilibrium phase diagram is typically presented at constant pressure and composition, whereas our simulations are carried out at constant volume and composition. For sufficiently large systems the two ensembles must agree, however, for finite systems (we typically use $N = 4096 = 64^2$ disks) coexistence is difficult to observe directly due to the surface-tension between the coexisting phases.

The chemical potential in a monodisperse system is $\mu = f + p$, where f is the free energy per particle and p is the reduced pressure [116], while the pressure is $P = p\phi/V_p$, where V_p is the volume of a particle and we have assumed that $kT = 1$. The free-energy f can be calculated using the BCMD algorithm from Chapter 6. The freezing point is the equilibrium point for phase B and phase AB at a given composition $x_A = 1 - x_B$. It can be determined by equating the (relative) chemical potential of particle type B inside the mixture AB, $\mu_{AB}^{(B)}$, with the chemical potential of the pure B phase, $\mu_B^{(B)}$, at equal pressures of the the

[1] Our choice of composition is closer to the estimated eutectic point for disk mixtures with $\kappa = 1.4$ [51] than the commonly-used $x_A = x_B = 1/2$, and also leads to equal area fractions of the large and small disks.

two phases,

$$\begin{aligned} \mu_{AB}^{(B)} &= \mu_B^{(B)} \\ P_{AB} &= P_B. \end{aligned} \tag{13.2}$$

Carrying out several detailed calculations gives the following expressions for the required chemical potentials:

$$\begin{aligned} \mu_B^{(B)}(\phi_B) &= \Delta f_B + p_B + \ln \phi_B - d \ln \kappa \\ \mu_{AB}^{(B)}(\phi_{AB}) &= \Delta f_{AB} + p_{AB} \frac{\kappa^d}{x_A + x_B \kappa^d} - x_A \tilde{\mu}_B + \ln \phi_{AB} - x_A \frac{\kappa^d - 1}{x_A + x_B \kappa^d} + \ln \frac{x_B}{x_A + x_B \kappa^d}. \end{aligned}$$

Calculating the relative chemical potential $\mu_{AB}^{(B)}$ requires calculating the sensitivity of the excess free energy of the mixture with respect to the composition (at constant density),

$$\tilde{\mu}_B = \left(\frac{\partial \Delta f_{AB}}{\partial x_A} \right)_\phi,$$

in addition to the excess free energies and pressure at a fixed composition. This can be done numerically by calculating the excess free energy for mixtures with slightly differing compositions via thermodynamic integration starting at low densities (where a liquid theory approximate EOS works well enough). The condition of equal pressures gives

$$\pi = \frac{\phi_{AB} p_{AB}(\phi_{AB})}{x_A + x_B \kappa^d} = \frac{\phi_B p_B(\phi_B)}{\kappa^d},$$

which can be used to express both ϕ_{AB} and ϕ_B as functions of π and then solve the equation $\mu_{AB}^{(B)}(\pi) = \mu_B^{(B)}(\pi)$ [c.f. Eq. (13.2)].

We do not give the full details of this calculation here, and merely state the result. Our free-energy calculations predict that at the freezing density $\phi_F \approx 0.775$, a crystal of density 0.8415 composed of predominantly large particles should start precipitating from the liquid mixture. As we will see shortly, nucleation is kinetically strongly suppressed due to the need for large-scale diffusion of large disks toward the nucleus [281], and in fact, we have not observed spontaneous crystallization even in simulations lasting tens of millions of collisions per particle well above the estimated freezing density.

We note in passing that had we used the best theoretical liquid-state predictions for the EOS of the liquid mixture, as discussed in Ref. [282], instead of the numerical EOS, we would not predict a freezing transition at $\phi_F \approx 0.775$. Instead, at all pressures the mixture would be predicted to be more stable, $\mu_B(AB) < \mu_B(B)$. As we will show in Section 13.3.1, the liquid-state theoretical prediction for the EOS is not sufficiently accurate at the densities above $\phi \approx 0.75$. We are not aware of any better analytical form of the EOS for mixtures, and therefore prefer to use an explicit numerical EOS for the liquid state. Also note that predicting the melting point and eutectic points requires knowing the EOS for the liquid mixture at all compositions, and we have not tried to calculate them in this work as our focus is on the freezing transition and in particular the kinetic glass transition, at a fixed composition. In Ref. [51] an approximate EOS based on an effective single-component system was constructed and the eutectic point estimated to occur at a composition of $x_A = 0.75$. It is believed that mixtures closer to the eutectic composition are better glass formers, and this was one of our reasons for choosing a composition $x_A = 2/3$ instead of the commonly used $x_A = 1/2$.

13.2.2.1 Phase-Separated Crystal Phase

We briefly examine the phase-separated crystal $A + B$, which we call the *crystal* phase even though one of the monodisperse phases could in fact be liquid (i.e., not possess long-range translational order). If phase separation is complete, at a given overall density ϕ the density of each of the phases $\phi_A(\phi)$ and $\phi_B(\phi)$ can be determined from the condition of mechanical equilibrium between the phases

$$P = \frac{p_{\text{mono}}(\phi_A)\phi_A}{V_p^A} = \frac{p_{\text{mono}}(\phi_B)\phi_B}{V_p^B} = \frac{p_{\text{mono}}(\phi_B)\phi_B}{\kappa^d V_p^A},$$

along with the condition that the overall density be ϕ,

$$\frac{1}{\phi} = \frac{1}{x_A + x_B\kappa^d}\left(\frac{x_A}{\phi_A} + \frac{x_B\kappa^d}{\phi_B}\right).$$

The solution to these equations is shown in Fig. 13.4. It is seen that at a given density the phase B, composed of large particles, is at a higher density and thus higher reduced pressure. At a density of $\phi \approx 0.75$ the small-particle phase A melts to a liquid, i.e., loses its translational order.

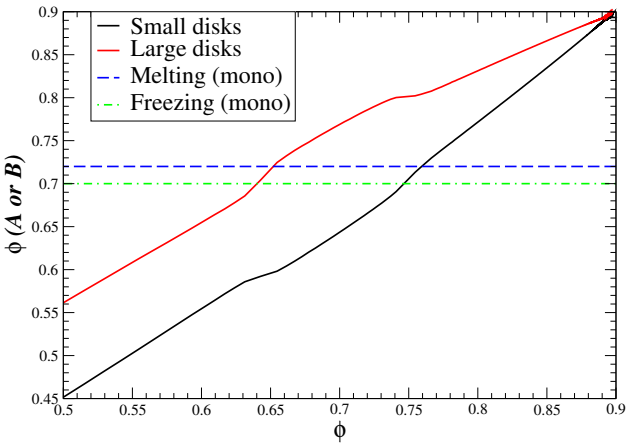

Figure 13.4: The densities $\phi_A(\phi)$ and $\phi_B(\phi)$ of the large- and small-particle phases in a fully phase-separated mixture at overall density ϕ. It is seen that the large-particle solid melts into a liquid at around $\phi \approx 0.75$, and at $\phi \approx 0.65$ the small-particle solid also melts. Note that this calculation assumes that there is no mixing between the small and large particles, which is only true at very high densities.

Below a certain density the large particles should start diffusing into the phase of small particles, forming a mixture AB. We used molecular dynamics to observe this melting of the phase-separated crystal, starting from a high-density phase-separated mixture and reducing the density slowly. However, as Fig. 13.5 illustrates, even the slowest MD runs did not achieve true equilibrium, as seen by the strong dependence of the observed EOS on γ. It is seen that below a density of $\phi \approx 0.8$, the phase-separated crystal is no longer stable and large particles start diffusing in the small-particle phase. This diffusion is very slow and even tens of millions of collisions per particle cannot equilibrate the phase-separated systems properly. A similar observation was made for a binary mixture of soft disks in Ref. [51] and it was concluded that "heterogeneous simulations can no longer [below the glass transition temperature] help us to identify the thermodynamically stable phase." In our simulations, long MD runs at a fixed density observed complete melting at a density $\phi = 0.765$ and therefore it is clear that at this density the stable phase is the mixed liquid. However, at $\phi = 0.775$ only partial melting occured and a crystallite of large particles remained stable for very long periods of time.

13.3 Is There an Ideal Binary Disk Glass?

Previous simulations have cast doubt on the existence of ideal glass transitions in hard-particle systems [217, 283]. In particular, it has already been suggested that the slope of $s_c(\phi_J)$ at $\phi_J^{\rm IG}$ dramatically affects the location of the presumed transition; an infinite slope shifts the transition to zero temperature [284]. Questions have also been raised about the validity of extrapolations into temperature/density regions that are inaccessible to accurate computer simulations [182], as well as the impact of finite-size effects [285]. Here, we present clear evidence that the concept of an ideal glass transition is flawed for distinctly different reasons. Specifically, for our model, we *explicitly* construct an exponential number of jammed packings with jamming densities ϕ_J that span from the accepted "amorphous" state with $\phi_J^g \approx 0.84$ to that of the crystal, with $\phi_{\rm max} = \pi/\sqrt{12} \approx 0.91$, thus clearly showing that the configurational entropy cannot be zero for the

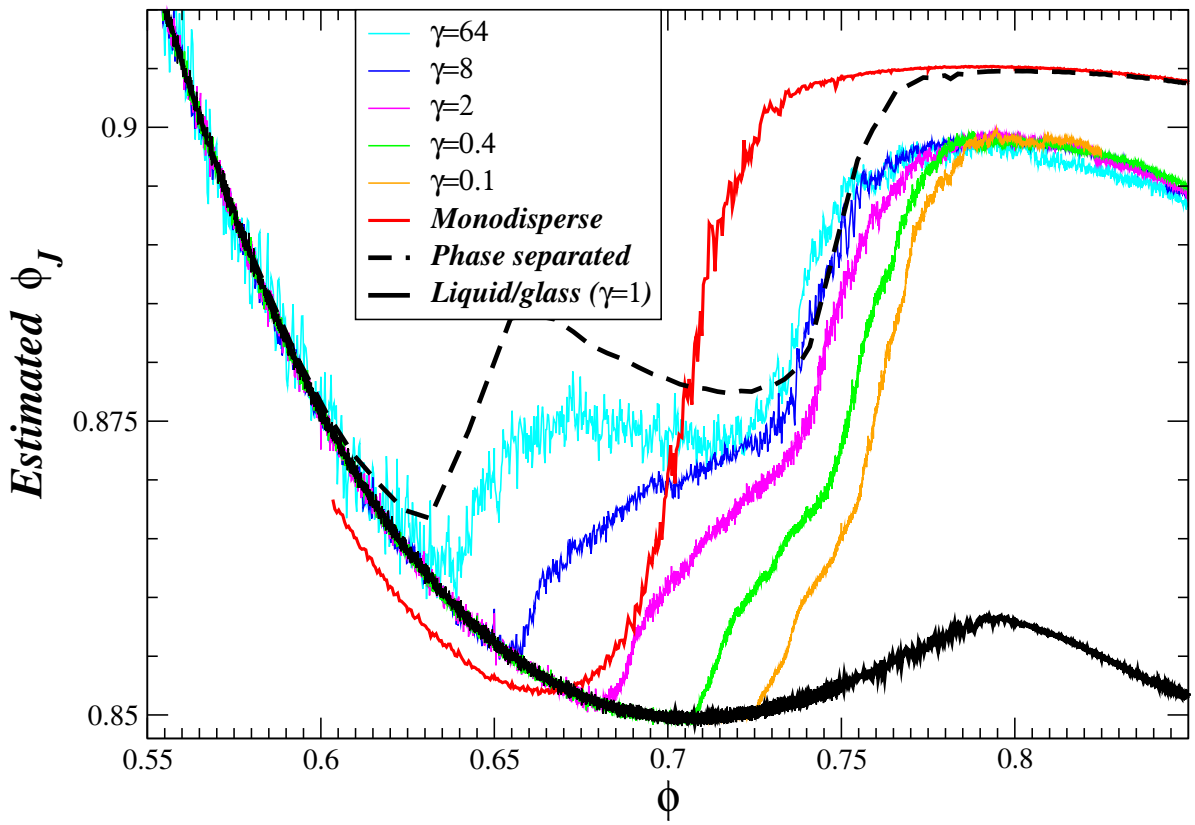

Figure 13.5: Equation of state as observed at different negative rates γ (shown in units of 10^{-6} in the legend), starting from a phase separated crystal. We also show the EOS obtained by compressing a dilute liquid with $\gamma = 10^{-6}$, forming a glass at high densities. It is seen that the phase-separated samples fall out of equilibrium at densities below $\phi \approx 0.8$, when sluggish diffusion of large particles into the small-particle solid begins. At sufficiently low densities complete melting into a mixed isotropic liquid occurs and the EOS matches the one measured by compressing a liquid. For comparison, we also show the EOS for the monodisperse hard disk system from Fig. 13.1. Note that the EOS of the phase-separated crystal does not perfectly match that of the monodisperse crystal because of the finite-size effects coming from the interface between the large- and small-particle solids.

hypothetical most-dense amorphous (ideal) glass distinct from the crystal. This objection is in the same spirit as the critique of the concept of random close packing (RCP) raised in Ref. [58], namely, that there is a continuous tradeoff between disorder (closely linked to degeneracy) and density, making the definition of a most-dense random packing ill-defined. Instead, Ref. [58] replaces RCP with the maximally random jammed (MRJ) state, i.e., the most disordered of all jammed states.

13.3.1 Kinetic Glass Transition

The calculation of the true equilibrium liquid equation of state (EOS) is not possible inside the glassy region with conventional simulation methods, especially for large system sizes [286, 182, 287, 285]. We produce glasses by starting with a low-density liquid and growing the particle diameters at a growth rate $\gamma = dD/dt \ll 1$ [75], for a very wide range of compression rates γ, as shown in Fig. 13.6. As seen in the figure, at densities below $\phi \approx 0.775$ the runs at different expansion rates are all in quasi-equilibrium and follow approximately the same EOS, namely, the EOS of the isotropic mixed liquid. After this density, fast compressions fall out of equilibrium and follow a glassy EOS, leading to a disordered jammed packing. Up to a density of about $\phi_g \approx 0.8$, the slowest runs follow the same EOS, which suggests that this is the EOS

of the super-compressed liquid, i.e., the metastable extension of the liquid branch.

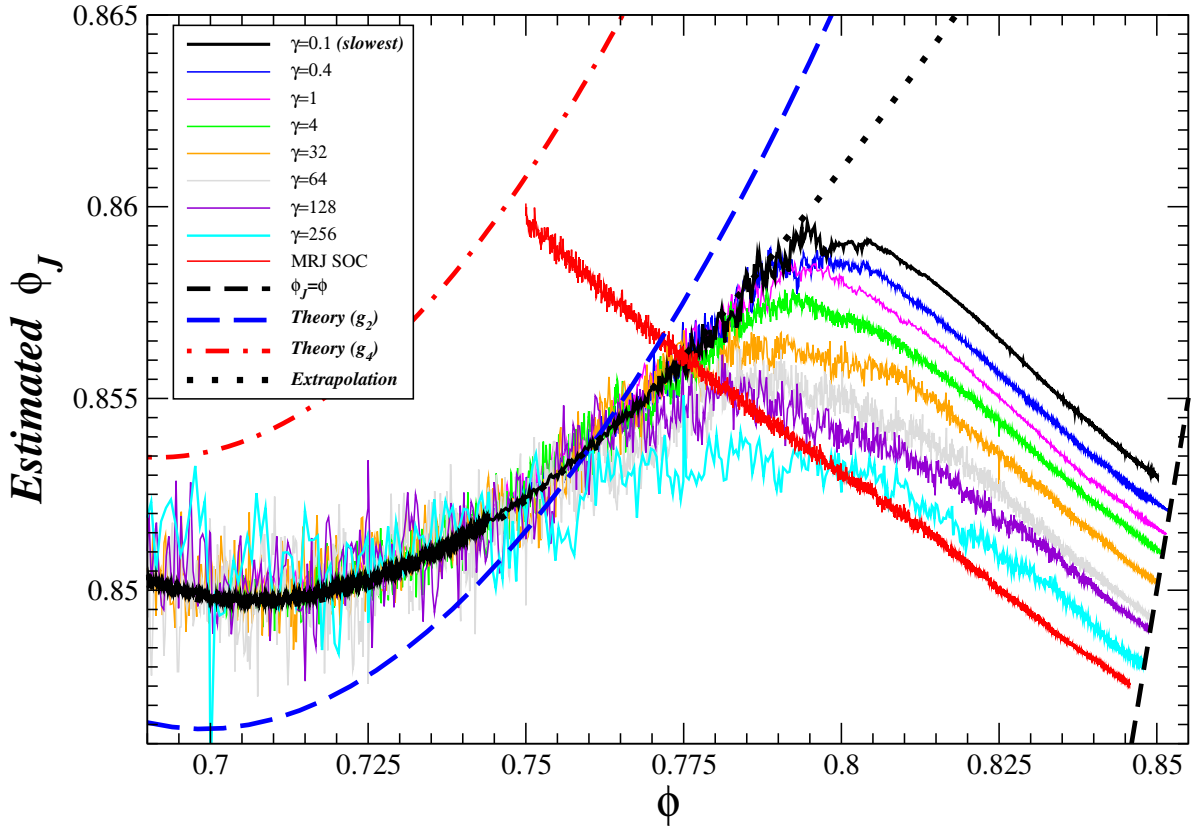

Figure 13.6: The equation of state $\tilde{\phi}_J(\phi)$ for $N = 4096$ disks as observed by compressing a liquid with different expansion rates γ (shown in units of 10^{-6} in the legend). At densities below $\phi \approx 0.775$ the runs at different expansion rates follow the EOS of the mixed liquid, however, after a kinetic glass transition density $\phi_g \approx 0.8$ the systems become trapped in glassy configurations even for the slowest runs. Note that we have run many more expansion rates over different density ranges and here we only show a representative sample. For comparison, we also show the theoretical liquid mixture EOS from Ref. [282], using either the Henderson (g_2) or improved (g_4) EOS for the monodisperse liquid, as given by Eq. (2) in Ref. [277]. It is seen that the theoretical liquid theory prediction is not sufficiently accurate at these densities. The EOS for the SOC-constrained (estimated) MRJ glass is also shown.

This kind of liquid-branch extension cannot be obtained using MD for monodisperse spheres in three dimensions since slow compressions crystallize after the melting density is surpassed, and therefore theoretical predictions about the existence or analytical form of such a hypothetical branch [276, 274] cannot be verified computationally. For the binary disk system, where crystallization does not occur, the results in Fig. 13.6 suggest that one can numerically study the liquid branch with high accuracy at least up to a density of $\phi_g \approx 0.8$. In order to analytically extend the liquid EOS beyond this density we have fitted a cubic function to $\tilde{\phi}_J(\phi)$ for the slowest runs up to the density where slowing down the compression by an order of magnitude does not change the observed pressure (within statistical variability). The fit

$$\tilde{\phi}_J(\phi) = 3.136 - 8.4826\phi + 10.277\phi^2 - 4.0356\phi^3 \tag{13.3}$$

is shown in Fig. 13.6, however, it should be emphasized that it is just a fit and there is no reason to believe it is quantitatively accurate much beyond $\phi \approx 0.8$. It is important to point out that in order to prepare a system in an (metastable) equilibrium liquid configuration at such high densities, one must compress the (stable) liquid from lower density (at least $\phi \approx 0.75$) very slowly. Quenching the liquid fast to a high density

produces states that are clearly not in any kind of thermodynamic equilibrium, even though they will appear stable due to very large relaxation times.

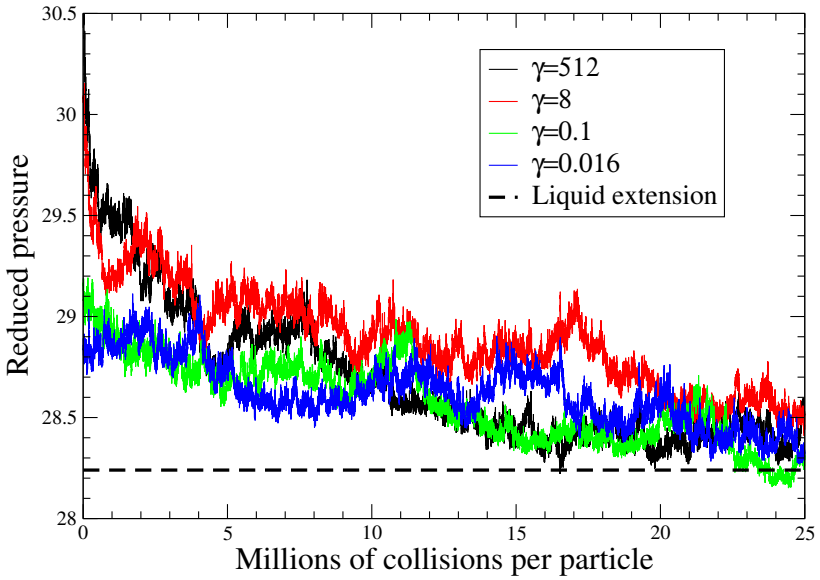

Figure 13.7: The relaxation of the pressure during long molecular runs at a fixed density $\phi = 0.8$, for several of the glasses produced during the compressions shown in Fig. 13.6 (γ is shown in units of 10^{-6} in the legend). For comparison, we show the pressure predicted by the liquid branch extension in Eq. (13.3).

As seen in Fig. 13.6, above the *kinetic glass transition density* $\phi_g \approx 0.8$, the systems become trapped in glassy configurations even for the slowest runs and jam in disordered jammed packings with jamming densities $\phi_J \approx 0.85$. Note that different definitions can be used for what the glass transition density is, here, we take it to be the maximal density at which our simulations can equilibrate (not necessarily in the true equilibrium state, but possibly also in a metastable liquid state) the binary mixtures. We also see in Fig. 13.6 that the nonequilibrium glassy EOS is very well described by an empirical linear relation

$$\tilde{\phi}_J = (1 + \alpha)\phi_J - \alpha\phi, \tag{13.4}$$

where $\alpha \approx 0.133$, over a wide range of $\phi > \phi_g$. We do not yet have a theoretical understanding of this relation. It is clear from the figure that even the slowest compressions fall out of equilibrium at a density around ϕ_g, so that equilibrating the liquid in reasonable time is not possible beyond this kinetic glass-transition density. Very long MD runs, with as many as 50 million collisions per particle, have failed to equilibrate our samples at a fixed $\phi = 0.8$, and in fact very different microstructures all remained stable for very long periods of time. This is shown in Fig. 13.7, where we show the evolution of the pressure during long molecular dynamics runs at $\phi = 0.8$ for several of the states in Fig. 13.6, at fixed density. For the glasses produced with faster expansion rates the initial pressure is higher and then decays more rapidly, however, a very slow residual decay of the pressure is seen in all of the samples, indicating the occurrence of very slow structural relaxation.

The final jamming densities of the glasses compressed at different rates are shown in Fig. 13.8. Note that slower compressions consistently yield denser packings with no hints of the existence of a *densest* glass. Fast compressions produce packings that are not truly jammed [75] and subsequent relaxation of these systems increases the density to around $\phi_J \approx 0.847$. This behavior of our hard-disk systems is closely related to the observation that supercooled liquids sample saddle points with the saddle index diminishing only below the temperature where even the slowest cooling schedules fall out of equilibrium [288, 289], i.e., the kinetic glass transition temperature. Observations similar to those in in Fig. 13.8 have already been made for systems of soft particles, e.g., the energy of the lowest inherent-structure sampled has been shown to continuously decrease for slower cooling [290].

In Fig. 13.9 we show the EOS for the SOC-constrained solids obtained by taking a snapshot of a configuration liquid at densities ranging from 0.7 (well within the equilibrated liquid density range) to

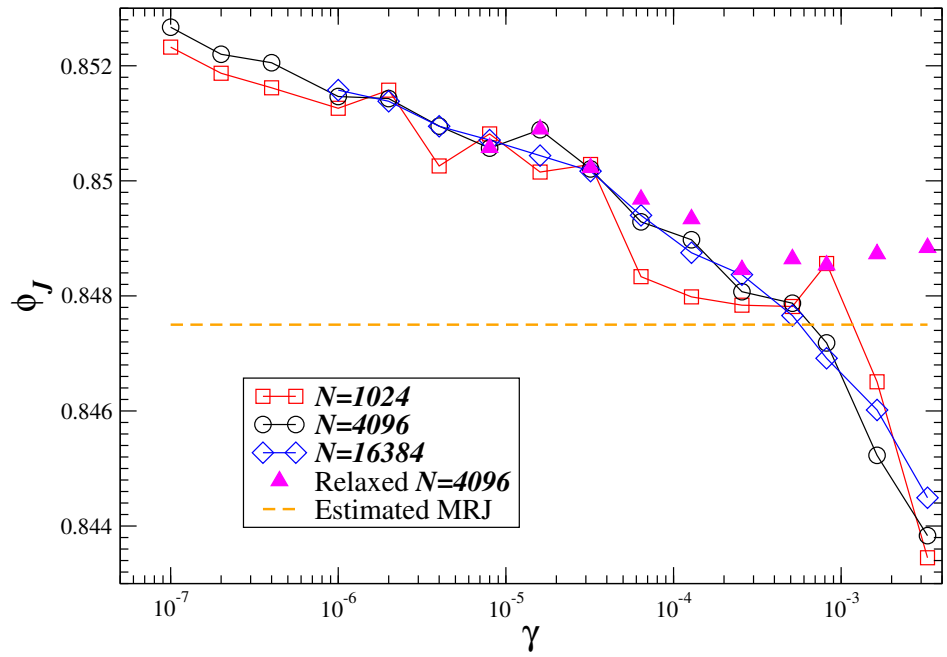

Figure 13.8: Final jamming density ϕ_J for different numbers of particles N, with and without additional relaxation (and subsequent slow compression) to ensure a truly jammed packing has been reached.

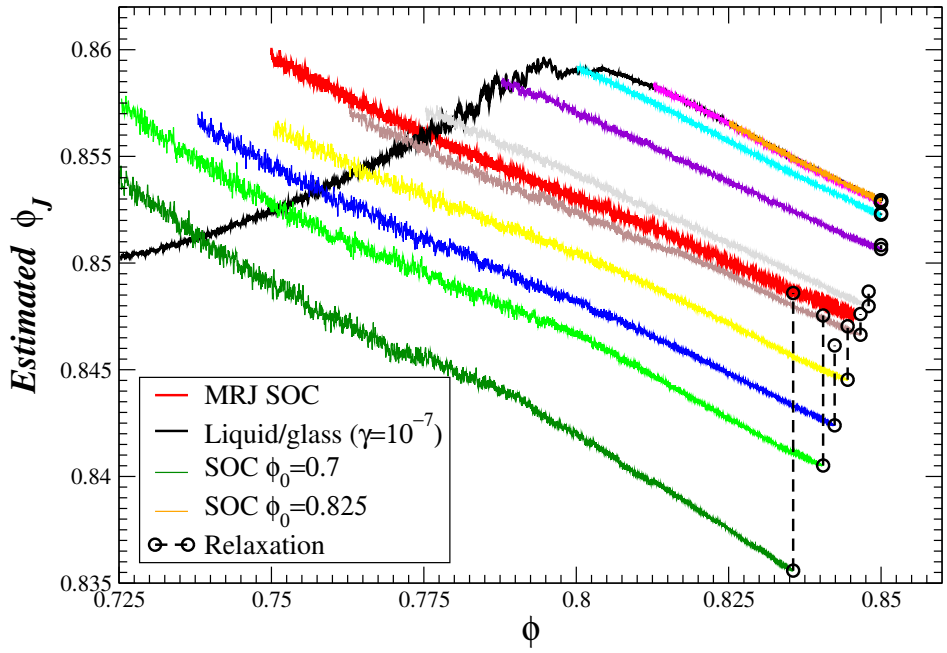

Figure 13.9: The EOS for SOC solids obtained by bounding each disk within a cell of twice its size, starting with liquid or glass configurations generated by saving snapshot configurations during one of the slowest compression shown in Fig. 13.6 ($\gamma = 10^{-7}$, replicated from Fig. 13.6 for comparison) and then compressing the SOC solid at $\gamma = 10^{-6}$ to near jamming. The initial densities ϕ_0 go from 0.7 to 0.825 and in the legend we only mark the first and last curves since ϕ_0 can be read from the starting point of the each curve. Also shown is the relaxation of the pressure (i.e., increase in ϕ_J) after the cell constraints are removed.

0.825 (well within the out of equilibrium glassy density range), and enclosing it with a bounding cell with $\Delta\mu = 1$. The liquid/glass configurations were generated by saving snapshot configurations during the slowest compression shown in Fig. 13.6 ($\gamma = 10^{-7}$). In the SOC models we measure the reduced pressure p through the momentum exchange during interparticle collisions, and do not include the pressure on the walls of the cells p_c. The SOC solid has a lower pressure than the unconstrained liquid because of the presence of cells, and the difference between the two diminishes as the density increases, becoming virtually negligible beyond the kinetic glass transition density. It is interesting to observe that the EOS of the SOC solids also follows Eq. (13.4) closely. The compressions of the SOC solids generated from liquid snapshots seem to produce jammed packings at densities as low as $\phi_J \approx 0.83$. However, these packings are unstable once the bounding cells are removed and molecular dynamics is run at a constant density, as shown in Fig. 13.6. In particular, this subsequent relaxation leads to glasses with $\phi_J \approx 0.85$, which is our best estimate for the MRJ density, as discussed in more detail shortly.

13.3.2 Configurational Entropy of Glasses

In this section we focus on calculating the free-volume and configurational contributions to the entropy of the dense liquid and glassy states obtained by compressing a liquid at different rates γ.

The quantity $N_g(\phi_J) = \exp\left[N s_c(\phi_J)\right]$ has recently been estimated via explicit enumeration for binary mixtures of relatively small numbers of hard disks [66]. These studies have observed an approximately Gaussian $N_g(\phi_J)$ that is peaked at a density $\phi_{\rm MRJ} \approx 0.842$, interpreted to correspond to the MRJ state for this system. Such a Gaussian $N_g(\phi_J)$ corresponds to an inverted parabola for $s_c(\phi_J)$, as has been assumed in previous studies of the thermodynamics of binary disk glasses [50]. For large systems, such enumeration is not yet possible and thermodynamics has been used to obtain estimates of $s_c(\phi_J)$, namely, it is estimated as the difference between the entropy (per particle) of the liquid $s_L(\phi)$ and the entropy of the "glass" $s_g(\phi)$,

$$s_c\left[\hat{\phi}_J(\phi)\right] = s_L(\phi) - s_g(\phi). \tag{13.5}$$

Here s_L is obtained via thermodynamic integration of the *equilibrium* liquid equation of state (EOS) from the ideal gas limit, while s_g is defined as the entropy of the system constrained to vibrate around a single basin with jamming density ϕ_J, without the possibility of particle rearrangements. There is significant ambiguity in defining these constraints; however, at least in the truly glassy region, the system is typically spontaneously constrained (jammed) by virtue of a very slow rearrangement dynamics, so that s_g can be defined reasonably precisely.

Formally, one can always partition configuration space into disjoint basins, each basin centered around a jammed configuration. For soft spheres such a partitioning can be defined by associating with each energy minimum (inherent structure) the basin of states for which gradient descent leads to the energy minimum under consideration. Such a partition is only useful, however, if the configurational volume (free energy) of a given basin $s_g(\phi, \phi_J)$ can be estimated easily, so that the number of basins can be calculated from Eq. (13.5). Note that the number of basins can only be estimated using thermodynamics up to exponential factors in N. We define the glass free-volume entropy $f_g = -s_g$ as the free energy of the SOC-constrained glass, where the cell is sufficiently large so that the pressure on the cell walls p_c is negligible, and small enough to prevent particle rearrangements[2]. The measured f_g obviously depends on the chosen bounding cell scaling (relative to the particles) $\mu = 1 + \Delta\mu$, unless the pressure on the cell walls $p_c(\Delta\mu)$ decays sufficiently rapidly so that truncating its integral at a given $\Delta\mu_{\rm max}$ does not substantially increase the free energy.

In Fig. 13.10 we show $p_c(\Delta\mu)$ for SOC-constrained glasses at several different densities (as in Fig. 13.9, but this time shrinking the bounding cells rather than growing the particles). We see that at densities above the kinetic glass transition, within numerical accuracy, $p_c(\Delta\mu)$ goes to zero as $\Delta\mu$ increases. However, for densities below the kinetic glass transition $p_c(\Delta\mu)$ clearly remains positive and therefore the f_g measured for the SOC glass will show significant dependence on the choice of cell cutoff. Closely related methods have previously been used to calculate s_c [181, 183, 291], with similar, though less accurate results. For soft-particle glasses an alternative method is to use the harmonic approximation to the vibrational entropy at an energy minimum as an estimate of s_g [180, 275]. All methods are rigorous only in the jamming or

[2]We have found $\Delta\mu = 1$, i.e., a cell diameter twice the diameter of the disk it bounds, sufficiently small to prevent particle rearrangements, see Fig. 13.10.

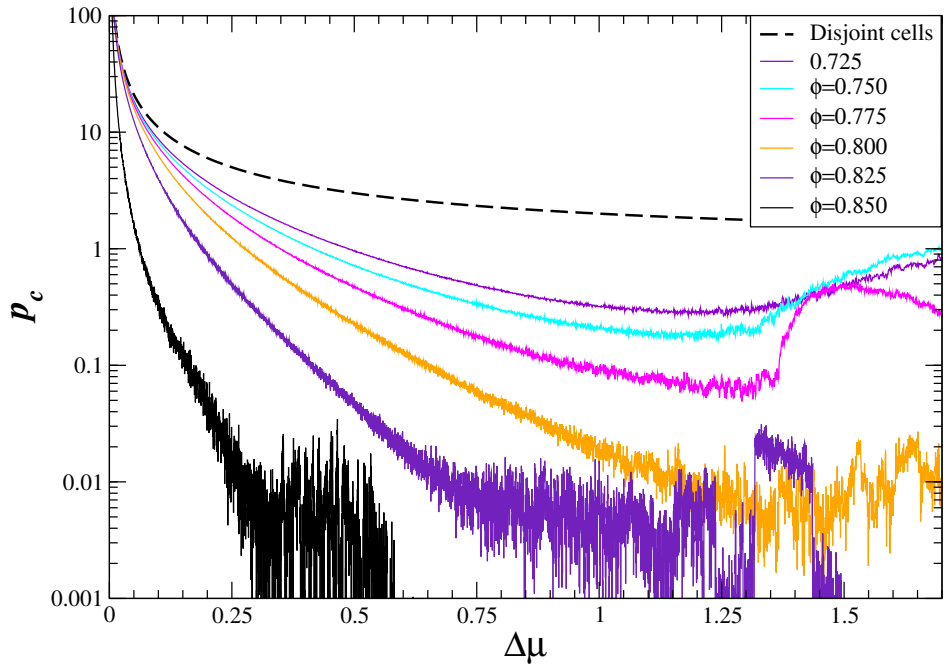

Figure 13.10: The averaged pressure on the bounding cell walls $p_c(\Delta\mu)$ as a function of the size of the cells for SOC glasses at several different densities. Similar results are shown in Ref. [184]. For comparison, the cell pressure that would be measured if the bounding cells are disjoint is shown. Note that when p_c becomes very small the simulation is not able to measure it accurately within the time interval over which particle-cell collisional momentum transfer is averaged.

$T \to 0$ limit, and are approximate for truly equilibrated liquids, so the quantitative results at low ϕ should be interpreted with caution.

The excess free energy of the SOC glassy mixtures is shown in Fig. 13.11, along with the free energy of the most-equilibrated liquid/glass $f_L(\phi)$ obtained by integrating the numerical EOS from the ideal-gas limit. We see that the free energy of the SOC glasses is substantially higher than that of the unconstrained liquid at all densities, however, the difference becomes approximately constant at high densities and seems to approach the entropy of mixing $s_{\rm mix} = x_A \ln x_A + x_B \ln x_B \approx 0.6365$. With this observation in mind, we show the measured $s_c(\phi) = s_L(\phi) - s_g(\phi)$ for the different glass compressions in Fig. 13.12. For comparison, the results for a slow compression of a monodisperse system are also shown, and the entropy of mixing $s_{\rm mix} = x_A \ln x_A + x_B \ln x_B$ has been subtracted from s_c. It is seen that for the monodisperse case $s_c - s_{\rm mix}$ ($s_{\rm mix} = 0$ in this case) becomes very nearly zero after the liquid freezes (around $\phi \approx 0.7$), indicating a continuous or a very mildly discontinuous liquid-solid phase transition (see also Fig. 13.2).

More interesting is the fact that $s_c - s_{\rm mix}$ also becomes nearly zero for the binary glasses around the kinetic phase transition (around $\phi \approx 0.8$). This important observation has not been made before. It means that the estimated number of packings that the liquid samples near the glass transition is very close to $s_{\rm mix}$, which is also the entropy of the uncorrelated ensemble of discrete states in which a fraction x_A of the particles is chosen to be large and the remaining particles are chosen to be small. It is interesting to observe that the parabolic fit to $s_c(\phi_J)$ from the work in Ref. [66], if constrained to equal the mixing entropy at the maximum, passes through zero at $\phi \approx 0.9$, much higher than the extrapolation in [50] and close to the crystal jamming density. We note that all measurements of s_c in the literature that we are aware of are above or close to s_c near the kinetic glass transition, and all estimates of the zero crossing of s_c are based on extrapolations beyond this point without numerical support [50, 181, 180, 275, 183, 276].

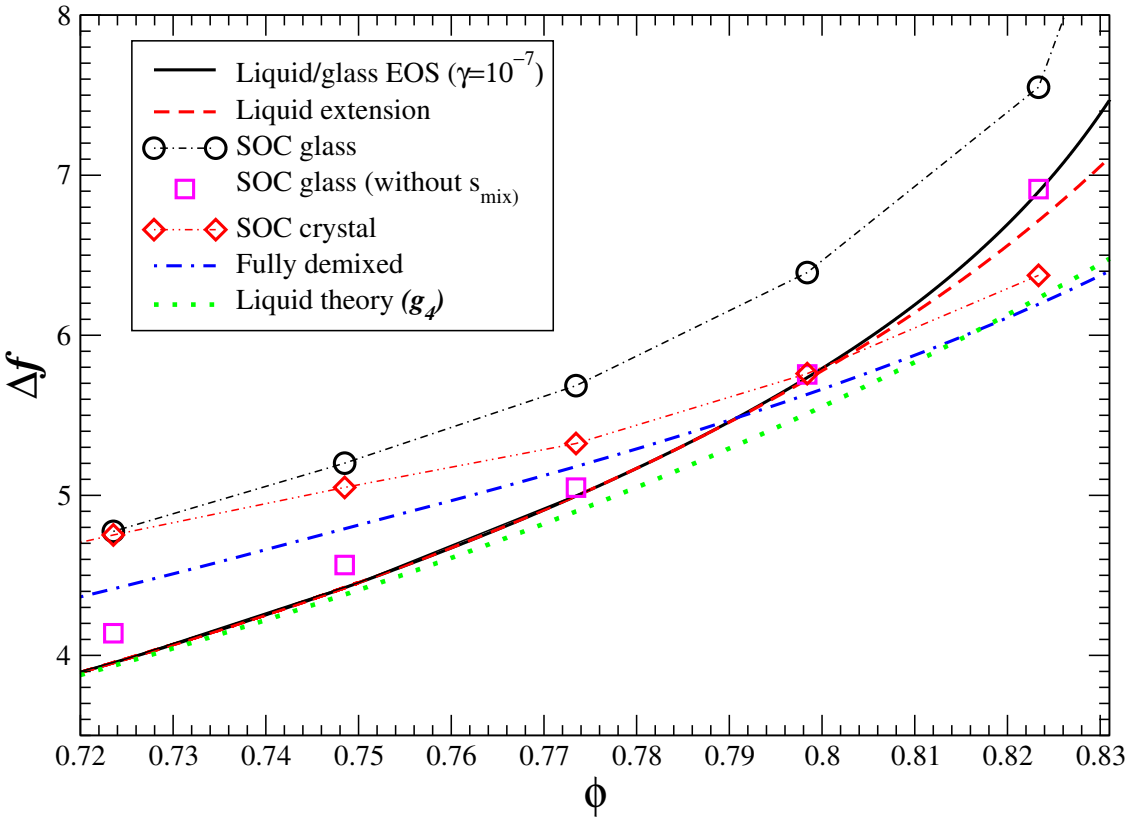

Figure 13.11: The excess free energy of the compressed liquid/glass phase (with and without $s_{\rm mix}$), and of phase-separated SOC-constrained crystals (note that the cell constraints prevent mixing at any density). Also shown are the predictions for an ideal phase-separated mixture, which does not perfectly match with that of the SOC solids even at high densities because of the entropic cost of the interface between the large- and small-particle solids.

13.3.3 Micro-Segregated Glasses

The observations made in Fig. 13.12 strongly suggest that extrapolations of s_c above the kinetic transition, predicting an ideal glass transition at density below the maximal possible density, are flawed. The only way to get zero configurational entropy is to get rid of the entropy of mixing, i.e., to fully demix the two types of disks. In fact, an exponential number of amorphous jammed packings exist over the whole density range from that accepted as the MRJ density $\phi_{\rm MRJ} \approx 0.84$ to that of the phase-separated crystal $\phi_{\rm max} \approx 0.91$. Lower-density jammed packings also exist [66]; however, they do not have thermodynamic significance and thus our simulations do not generate them. In our simulations we observe that higher ϕ_J implies micro-segregation in the form of increased clustering of the large particles. This has been most vividly demonstrated in Ref. [51]. This observation suggests that one can generate denser packings by artificially encouraging clustering, i.e., increasing the amount of (spatial) ordering in the packings.

To achieve clustering, we start from a monodisperse ($\kappa = 1$) triangular crystal at pressure $p = 100$ in which a third of the particles has been selected as being "large". The large particles then slowly grow in diameter while the system is kept in (quasi)equilibrium at a constant (isotropic) pressure using a Parinello-Rahman-like variation of the MD algorithm [11], as described in Section 3.2.4.2. When $\kappa = 1.4$ we stop the process and then slowly compress the system to jamming. By spatially biasing the initial partitioning into large and small disks, we can achieve a desired level of clustering and higher jamming densities for the final packings. Figure 13.13 illustrates two different jammed packings, one with an uncorrelated random choice of large disks, and another with correlations encouraging micro-segregation. The packing produced

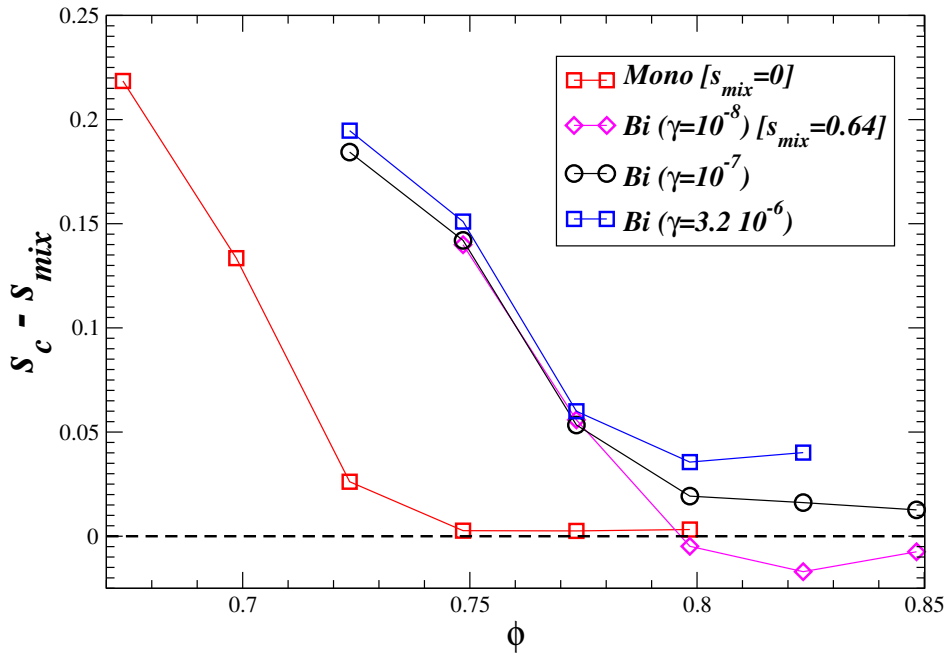

Figure 13.12: Estimated $s_c(\phi) - s_{\rm mix}$ for monodisperse and bidisperse systems of $N = 4096$ disks, as obtained from (sufficiently slow) compressions with a range of γ's.

by an uncorrelated random assignment of small versus large particles is the most disordered packing, i.e., it is representative of the MRJ state for this binary hard disk mixtures.

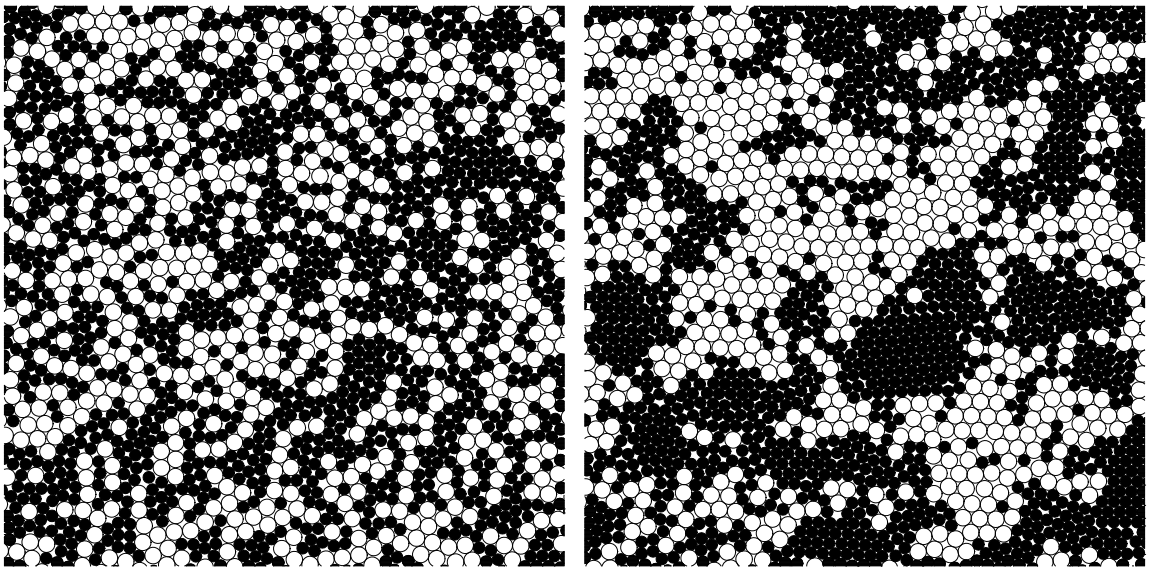

Figure 13.13: The microstructure of a packing without (left, $\phi_J \approx \phi_{MRJ} \approx 0.846$) and with moderate clustering (right,$\phi_J \approx 0.850$).

For the purpose of creating clustered initial assignments of "small" or "large" (i.e., A or B) labels, we use a level-cut of a Gaussian random field (GRF) [292]. Specifically, we construct a discretized GRF on a square lattice of $4096 = 64^2$ [293], and assign label A to all sites where the filed has value larger than a

certain cutoff (chosen so that two thirds of the disks are labeled A), and label B otherwise. By using suitably chosen parameters for a flexible family of pair correlation functions originally proposed by Matern [294] we were able to generate different levels of clustering, as illustrated in the inset in Fig. 13.15. Specifically, the two parameters for the Matern correlation are the correlation length R and the interface smoothness parameter $0 < \nu < 1$. Increasing R or ν increases the clustering, and we used five values of R going from 1 to 5, along with five values of ν going from 0.1 to 0.5, for a total of 25 different types of micro-segregated initial configurations. Higher values of R and ν produce denser packings as the B disks are grown in size at constant pressure, as illustrated in Fig. 13.14.

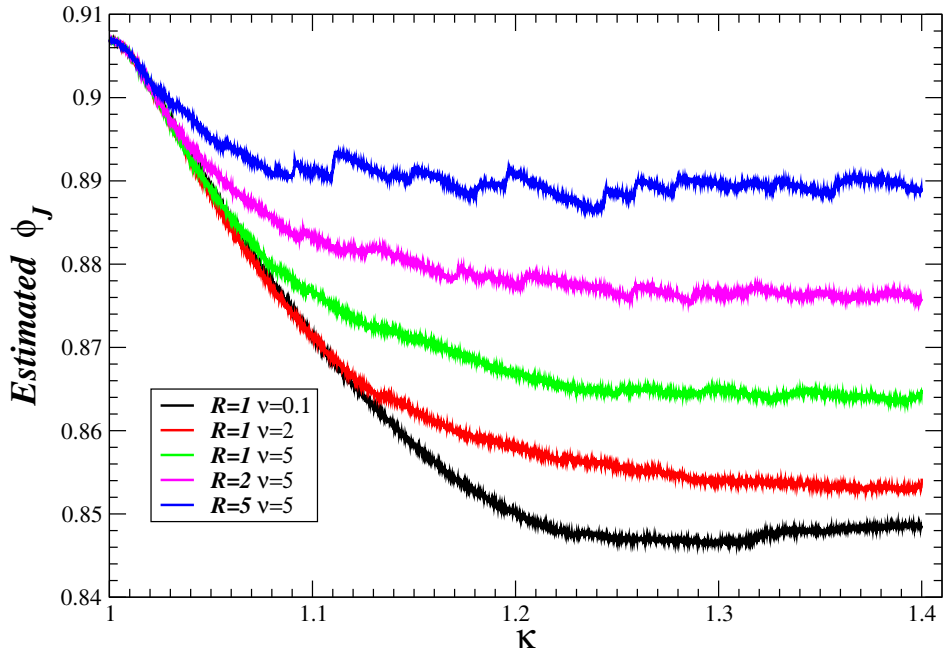

Figure 13.14: Converting an initially monodisperse disk packing of $N = 4096$ disks into a jammed binary disk packing by slowly growing a chosen third of the disks at a growth rate $\gamma = 10^{-5}$ while keeping the pressure at $p \approx 100$. Here we used a leveled GRF with the Matern correlation in order to generate clustered initial configuration. It is seen that the estimated jamming density $\tilde{\phi}_J(\kappa)$ decreases from $\tilde{\phi}_J \approx 0.91$ as the size dispersity κ grows to the final value of $\kappa = 1.4$. The final jamming density is larger the more clustered the initial configuration is.

To determine the configurational entropy (degeneracy) for a given choice of the GRF parameters, we use a recently-developed algorithm for obtaining numerical approximations of the entropy (per site) of lattice systems [295]. In principle the true degeneracy can be calculated by measuring the probability $p(\mathcal{C})$ of observing a particular configuration $\mathcal{C}$ of a rectangular window of $n \times m$ sites, and then calculating the entropy per site

$$s_{n,m} = \lim_{n,m\to\infty} \frac{S_{n,m}}{nm} = \lim_{n,m\to\infty} \left\{ -\frac{1}{nm} \sum_{C} [p(\mathcal{C}) \ln p(\mathcal{C})] \right\},$$

where the sum is over all of the possible 2^{nm} configurations. In Ref. [295] the above limit is approximated accurately and efficiently with small windows by exploiting a Markov approximation, to obtain

$$s_{n,m} \approx S_{n,m} - S_{n-1,m} - S_{n,m-1} + S_{n-1,m-1}.$$

This approximation is seen to converge relatively fast, as demonstrated in Fig. 13.15 for the 25 different choices of parameters R and ν for the Matern correlation function. We have used windows of 4×4 sites, since calculating $S_{n,m}$ requires generating many GRFs for the same correlation function and counting the

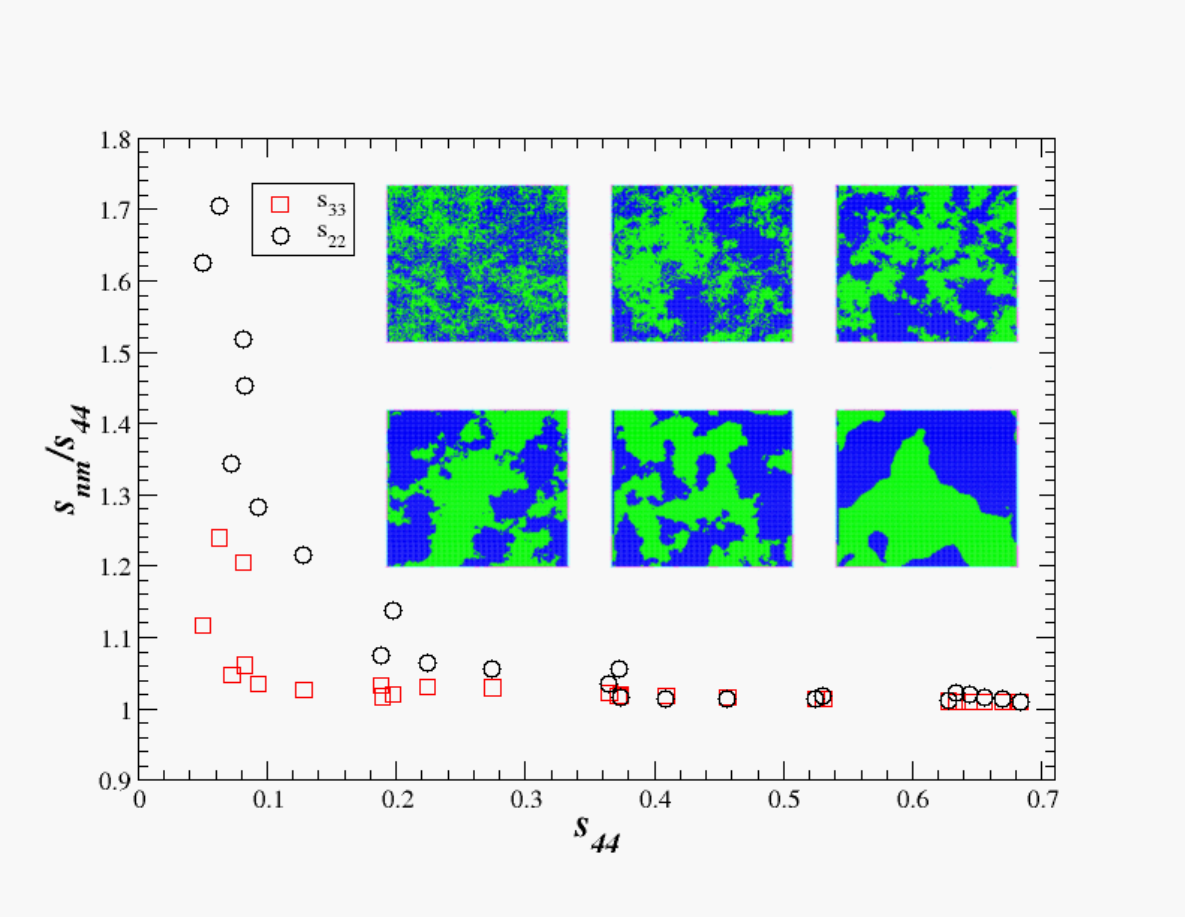

Figure 13.15: The Markov approximation $s_{n,n}$ to the entropy per site for $n = 2-4$ [295], for different choices of the parameters R and ν for the Matern correlation function in the GRF. It is seen that $s_{3,3}$ is close to $s_{4,4}$, suggesting that $s_{4,4}$ is a good approximation to the true entropy per site, especially for less clustered configurations (i.e., higher $s_{4,4}$). (*Inset*) Sample realizations of the partitionings into large and small sites on a 64^2 grid, for several values of R and ν.

probabilities of observing different configurations of a window of size $n \times m$ sites. This process becomes prohibitively expensive for $n = m = 5$.

We will assume that a different assignment of A and B labels will produce a distinct jammed packing, i.e., that the configurational entropy s_c for the jammed packings generated with a particular choice of GRF is well-approximated by $s_{4,4}$ from Fig. 13.15. Figure 13.16 shows our results for s_c versus the jamming density ϕ_J, for the 25 different choices of the GRF parameters. The results clearly show that in order to increase ϕ_J one must sacrifice degeneracy (s_c). The figure also shows the first *measured*, rather than extrapolated, estimate of $s_c(\phi_J)$. This observed $s_c(\phi_J)$ only goes to zero for the phase-separated crystal state, rather than the hypothetical amorphous ideal glass state postulated by extrapolations.

It is not *a priori* obvious that $s_c(R, \nu)$ is strongly correlated with $\phi_J(R, \nu)$, since they both depend on both R and ν. This is demonstrated to be the case in Fig. 13.17, where we show color plots of $s_c(R, \nu)$ and $\phi_J(R, \nu)$ over the grid of 25 values for (R, ν), and a strong correlation between the two is clearly seen. This may be however due to the particular choice of the correlation function in the GRF. Ideally, what we are interested in thermodynamically is the highest s_c at a given ϕ_J, i.e., the type of micro-clustering that decreases the degeneracy the least in order to increase the jamming density by a given amount from ϕ_{MRJ}. We do not know how to calculate the true $s_c(\phi_J)$, or how to construct samples representative of the most disordered samples at densities other than ϕ_{MRJ}. The results obtained for the particular way we generated

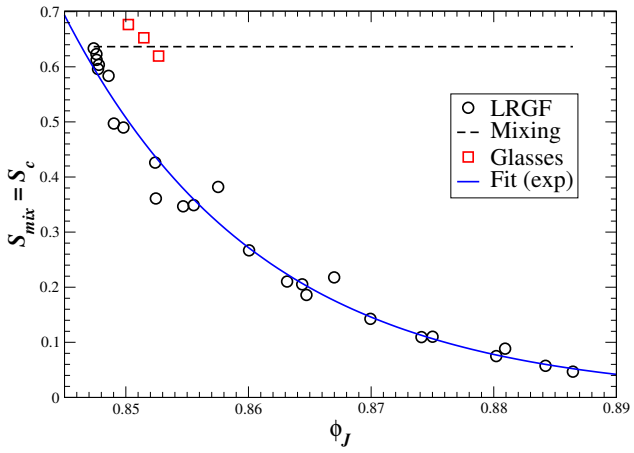

Figure 13.16: The measured degeneracy of packings of $N = 4096$ disks obtained by using different parameters of a random Gaussian field with Matern correlations [292], as a function of the jamming density. For comparison, we have shown $s_c(\phi = 0.825)$ for the three compressions shown in Fig. 13.12, and an exponential fit to the data.

micro-segregated samples, shown in Fig. 13.16, show a rapid drop in s_c away from the MRJ point. That is, one must cluster significantly before seeing an appreciable increase in the jamming density.

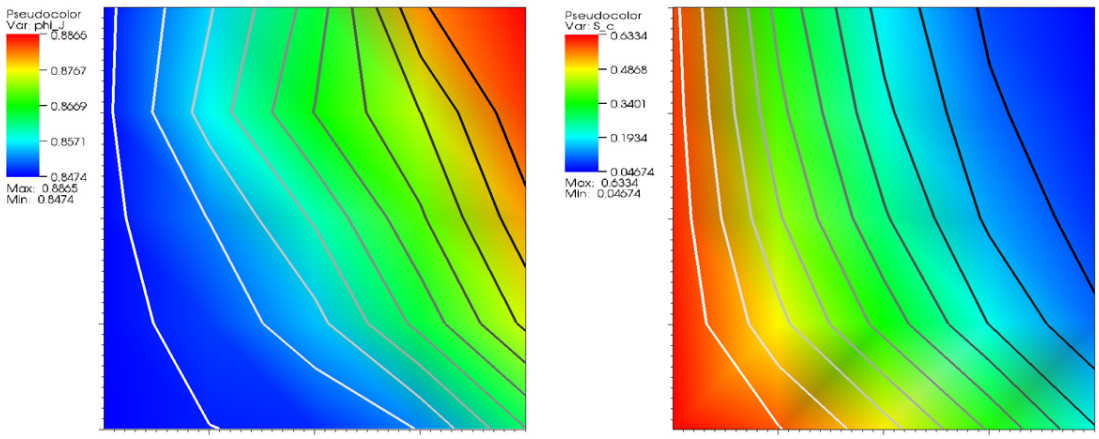

Figure 13.17: A demonstration that $s_c(R, \nu)$ (right) is strongly correlated with $\phi_J(R, \nu)$ (left). The x-axes of the color plots is ν, and the y-axes is R. If a similarly strong correlation exists for different choices of clustering correlations, it is important to determine whether the $s_c(\phi_J)$ we calculate here continues to apply.

A calculation of f_J for the different micro-segregated glasses, using the BCMD algorithm, shows that f_J is essentially constant independent of (R, ν), at least to within statistical fluctuations. Substituting this in Eq. (13.1) together with the exponential fit $s_c(\phi_J)$ from 13.16 predicts that for densities lower than $\phi \approx 0.8$ the equilibrated liquid samples the MRJ basin, $\hat{\phi}_J(\phi) = \phi_{MRJ}$, and for higher densities the liquid samples the phase-separated crystal basin, $\hat{\phi}_J(\phi) = \phi_{\max}$. On the other hand, the smoothly increasing $\phi_J(\gamma)$ in Fig. 13.8, the spontaneous clustering seen for supercooled soft disks in Ref. [51], and the widely-observed lowering

of the energies of the sampled inherent structures for soft-sphere glasses upon supercooling [290], all suggest that $\hat{\phi}_J(\phi)$ should be continuously increasing for very dense liquids. Such behavior of $\hat{\phi}_J(\phi)$, at least within the inherent-structure formalism we consider here, requires that $s_c(\phi_J)$ decay very slowly around ϕ_{MRJ}, so that the liquid prefers to loose degeneracy by ordering (clustering) in order to gain free volume. In fact, the expected behavior in $s_c(\phi_J)$ is that it would be quadratic around ϕ_{MRJ}, i.e., close to an inverted parabola with a maximum at ϕ_{MRJ}. We expect that there exist an exponential number of ordered jammed states with densities lower than ϕ_{MRJ}, and therefore that $s_c(\phi_J)$ is a smooth function, rather than having a cusp-like maximum or sharp discontinuity at ϕ_{MRJ}.

13.4 Conclusions

Continuing on work in Ref. [58], we explicitly demonstrated that the concept of random close packing as the most-dense jammed amorphous packing is flawed. By trading off degeneracy for density in a continuous manner, we constructed an exponential number of amorphous jammed packings with densities spanning the range from the most disordered to most ordered jammed states. We explicitly calculated, as opposed to extrapolated, the degeneracy entropy for densities well above that of the postulated ideal glass transition, and found that the degeneracy is positive for all "amorphous" states and vanishes only for the phase-separated crystal. For the maximally random jammed state, the found a degeneracy entropy very close to the mixing entropy. Furthermore, our free-energy calculations predict a thermodynamic crystallization well-below the kinetic glass transition, casting additional doubt on the search for a thermodynamic origin of the glass transition. Although the present study focused on the hard-disk binary mixture, the fundamental principles are general enough to be applicable to a host of related systems, notably, both mono- and bi-disperse with hard-core and soft interactions.

It is important to point out that our argument has nothing to do with mixing macroscopic liquid and crystal domains (with sharp and identifiable interfaces that make for a negligible reduction in density) in order to get mixed states of intermediate densities. Instead, we construct an exponential number of amorphous configurations that show no signs of crystal nuclei. Artificially mixing large crystal domains with large liquid domains severely underestimates the number of available jamming configurations, since near the glass transition the configurational entropy is close to the entropy of the completely mixed system. For the binary hard disk system, there is no sharp boundary between crystal and liquid states. The micro-separated samples we constructed in our work are not mixtures of a liquid and a crystal phase. They are disordered (amorphous) states which have no qualitative difference from the liquid state. In particular, they do not have (quasi)long-range order, and do not have macroscopic domains that could be considered crystal. Perhaps more significantly, these states are not artificial constructions in which we just mixed some crystal and liquid in a trivial manner. Rather, our choice was motivated by careful observations of the actual thermodynamic and kinetic behavior of hard disk mixtures. Namely, as we decreased the compression rate (cooling rate for soft disks), we saw spontaneous microclustering happening (this has been observed in other systems). If we had many more decades of computational power, we believe we would see micro-segregated glassy states appear spontaneously.

Unfortunately, our results do not resolve the mystery of the glass transition. In fact, the complete thermodynamic behavior of hard disk mixtures remains unclear. Free-energy calculations predicted a freezing transition, but it could not be observed directly with classical MD due to the dramatic kinetic slowdown near the glass transition. Such free-energy calculations proceed in reverse order: One assumes what the equilibrium structures are, and then selects the one with the lowest free energy. At high densities, however, it is not clear what the properties of the liquid phase are, end whether it exists at all. it is hoped that the inherent structure formalism, i.e., the partitioning of the available configuration space into jamming basins, can describe the thermodynamic properties well. For hard sphere systems, this requires the identification and counting of distinct configurations in a statistical ensemble of jammed packings. This has been done by direct enumeration for small systems, however, enumeration is not possible for large systems. In this work, we identified distinct packings for hard disk mixtures with distinct partitionings of the disks in a monodisperse triangular crystal into small and large disks. Such an identification converts the difficult geometrical problem of packing disks into a much simpler combinatorial problem of generating partitionings of the triangular lattice. This identification was suggested by the fact that the calculated configurational

entropy near the glass transition is very close to the mixing entropy.

For the purposes of calculating configurational entropy, it is not necessary that the identification be one-to-one, rather, only that the number of jammed packings corresponding to a given partitioning is sub-exponential in N, and vice-versa. Starting with a given partitioning, we described a procedure for generating a corresponding packing using MD. While it is not trivial to prove mathematically that this generates a unique and distinct packing, we expect that this would be the case if the MD algorithm would be run for an infinitely long time at infinite pressure. In the other direction, starting with a given jammed isostatic bidisperse packing, one can shrink the diameters of the large particles and maintain the existing contacts, as discussed in Section 4.5, while also maintaining jamming. Along the way, new contacts will be formed and old contacts broken, and the path of the algorithm is not unique, but it is expected that the number of different choices that can be made is subexponential in N. In the end this procedure will generate a monodisperse jammed packing where the disks are labeled small and large. Our analysis in Chapter 4 suggested that the majority of strictly jammed monodisperse disk packings are the triangular packings with vacancies. These arguments suggest that for hard disk mixtures there may indeed be a strong correspondence between jammed packings and partitionings of the triangular lattice. It would be a useful future exercises to consider adding vacancies to the initial triangular configurations before applying the Gaussian random fields to them. This might increase $s_c(\phi)$ and produce the expected inverted parabolic shape, and in particular, generate jammed packings at densities below ϕ_{MRJ}.

An important avenue of research is the development of algorithms to equilibrate liquids at densities higher than the kinetic glass transition. It is clear that such algorithms must be very different from classical MD, however, despite the fact that several algorithms have helped significantly reduce the simulation times necessary to equilibrate supercooled or super-compressed liquids [286, 182, 287, 285], true thermodynamic (meta) equilibrium for samples of reasonable size has not yet been achieved at sufficiently high densities to properly elucidate the thermodynamics of disordered solids. Finally, the continued failure to identify any thermodynamic origin to the glass transition suggests that the kinetics of supercooled and super-compressed liquids needs to be understood better. We believe that the hard sphere system is an ideal model for such studies, and described at least one avenue of research in Section 4.5.

Chapter 14

Conclusions and Future Directions

In this dissertation we studied the statistical and thermodynamic properties of ordered and disordered jammed packings of hard frictionless particles computationally, theoretically, and experimentally. We reached several novel and important conclusions, which we briefly highlight here.

In the first part of this thesis we focused on computational and theoretical work. In Chapter 2 we showed that overlap potentials originally designed for ellipsoids by Perram and Wertheim for hard ellipsoids can be generalized to arbitrary convex shapes and given a very useful geometric interpretation. We used these overlap potentials to develop a very efficient event-driven molecular dynamics (MD) algorithm for hard particles in Chapter 3 and used this algorithm to generalize the Lubachevsky-Stillinger packing generation algorithm to nonspherical particle shapes. In Chapter 4 we demonstrated that the concept of jamming can be formalized for hard spheres and designed a mathematical programming algorithm to rigorously test for jamming, rather then relying on dubious heuristics or unsupported assumptions as is typically found in the literature. In Chapter 5 we generated for the first time jammed disordered packings of hard ellipsoids and found that they have a surprisingly high density in addition to having a high contact number. We also showed that the generated packings are collectively jammed despite being hypostatic, thus demonstrating by example that the isostatic conjecture, commonly taken for granted in the granular media literature, does not apply to nonspherical particles. Finally, in Chapter 6 we extended our MD algorithm to rigorously calculate the free energy of nearly jammed hard-particle packings.

In the second part of this thesis we applied the computational and theoretical tools from the first part to the study of the statistical and thermodynamic properties of disordered and ordered packings of hard particles. In Chapter 7 we demonstrated that the computationally-generated packings of frictionless particles closely resemble those produced in experiments with hard spheres and ellipsoids. In Chapter 8 we discovered the densest known crystal packing of ellipsoids by using molecular dynamics. In Chapter 9 we observed that disordered jammed hard sphere packings are strictly isostatic and found an unusual power-law divergence in the number of near contacts. In Chapter 10 we showed that disordered jammed hard sphere packings are hyperuniform and additionally have an unusual non-analytic linear behavior of the structure factor near the origin not found in liquids. In Chapter 11 we found that disordered sphere packings in dimensions $d = 4, 5$ and 6 have very similar properties to those in three dimensions, but with a pronounced pair decorrelation in higher dimensions for distances larger than the hard core. In Chapter 12 we discovered that hard dominos form a disordered solid or glass phase with tetratic order, but no crystal periodicity. Finally, in Chapter 13 we demonstrated that there is no ideal glass transition for binary hard disks of size dispersity 1.4.

14.1 Future Directions

This dissertation identified many open problems and suggested numerous avenues for future research. In this section we point to certain promising directions that our work has identified, and also point to questions that have not yet been addressed properly.

14.1.1 Molecular Dynamics

In Chapter 2 we developed PW-type overlap potentials for the case of one convex particle contained within another, however, we found that the inner PW potential is difficult and expensive to calculate numerically in a stable way. The inner PW potential is essential in the EDMD algorithm of Chapter 3 when using near-neighbor lists and when calculating the free energy using the BCMD algorithm from Chapter 6. Therefore, development of alternative overlap potentials or alternative evaluation of the inner PW overlap potential is an important task. Additionally, developing overlap potentials for nonsmooth particle shapes, and in particular for shapes with flat edges and sharp corners, is needed in order to extend the work we presented here to particle shapes other than smooth convex ones.

In Chapter 3 we demonstrated that event-driven molecular dynamics is an essential tool in the study of hard-particle packings, both because it is the only rigorous way to study the dynamics in the jamming limit, and because it can be implemented very efficiently. However, even today's sophisticated single-processor computers are not sufficiently powerful to carry out certain important studies of very large systems, nonspherical shapes, or systems in higher dimensions. Parallelization of the EDMD algorithm is highly non-trivial; however, it is a worthy task to undertake. Predicting the time of collision of two moving nonspherical particles can also be improved, either by using techniques other than bounding spheres in order to narrow the search intervals during which a collision may happen, or by improving the algorithm to search those intervals for a collision. While we have made substantial progress in tailoring the EDMD algorithm to systems of very elongated or very flat particles, the improvements are still not sufficient to efficiently deal with very large aspect ratios. It is an open challenge to develop an algorithm to improve the efficiency of building the near-neighbor lists for very oblate particles. Finally, improving the search for neighbors in systems with large size dispersity between the particles is a barely-explored research area.

14.1.2 Generating Jammed Packings

The generation of jammed packings of hard particles is a non-trivial task, and in this work we primarily focused on variations on the Lubachevsky-Stillinger molecular dynamics algorithm. It is important to explore alternative algorithms, and in particular, algorithms based on energy minimization in soft-particle systems, as well as algorithms based on contact networks. Algorithms that use deformable particles instead of perfectly hard ones are important because realistic particles are not perfectly hard. Furthermore, sophisticated mathematical programming techniques can be used to efficiently produce packings that are guaranteed to be very close to a true jamming point, including for strict jamming. Molecular dynamics on the other hand relies on the natural dynamics in the system, which slows down dramatically in the jamming limit and may get trapped in packings that are not truly jammed. Additionally, MD is not able to generate truly strictly jammed packings. Algorithms based on contact networks are important because they guarantee the production of an ideal (within the numerical tolerances used) collectively or strictly jammed packing. Furthermore, with the help of better theoretical understanding of the conditions that a given contact network be realizable as a jammed packing, there is the potential of developing algorithms that rigorously enumerate all jammed packings by enumerating contact networks. Such algorithms will surely be prohibitive for very large packings. However, designing rigorous algorithms is important even if only applicable to small systems.

Improved packing-generation algorithms are needed in order to identify MRJ packings for a variety of systems. Understanding the behavior of the packing algorithms, and in particular, the influence of the various control parameters on the density and amount of ordering in the produced packings, is necessary in order to avoid heuristics or guesses for what the "correct" parameters are. In two dimensions, it is not yet firmly established whether there really are no disordered non-crystalline strictly jammed packings, or whether crystallization is observed in packing algorithms simply because it is favored by the (thermo)dynamics of the algorithms commonly in use. Designing algorithms aimed at generating packings other than MRJ packings, is also important. In particular, it is important to generate lower- and higher-density disordered jammed structures. The LS algorithm does not produce jammed packings with densities below that of the MRJ state because such low-density jammed packings are thermodynamically-unfavored. In Chapter 13 we constructed disordered jammed packings of binary hard disks spanning the density range from the most disordered to the most ordered. Can we do the same in three dimensions for monodisperse spheres?

Finally, adding friction to packing-generation algorithms is an important task for future research.

14.1.3 Better Understanding of Jamming

In Chapter 4 we demonstrated the importance and usefulness of including the lattice vectors as degrees of freedom when discussing jamming. In particular, we related lattice deformations to macroscopic strain and demonstrated that the Lagrange multipliers corresponding to the strain degrees of freedom can be interpreted as macroscopic stress. It is important to explicitly derive second-order derivatives of overlap between particles with respect to strain degrees of freedom, in order to extend the second-order conditions for jamming from Chapter 5 to strict jamming, and also express the elastic moduli in terms of the forces and stiffnesses of the interparticle contacts. Several other calculations carried out in this work should be extended to second order, and in particular, the sensitivity of the jamming density with respect to changes of the particle shapes or other perturbations should be calculated to second order.

Deformation of the lattice vectors can be seen as a modification of the boundary conditions to allow for macroscopic degrees of freedom. Additionally, modifications of standard periodic boundary conditions to include non-periodic motions is important. A periodic packing can be considered to be a small part of an infinite repetitive packing, and looking for unjamming motions with periodic boundary conditions corresponds to looking for unjamming motions of the infinite packing that have a single Bloch wave (Fourier) component. However, the infinite system may also possess unjamming motions (modes) that have larger wave-lengths, and Bloch wave (Fourier) decomposition is a natural way to extend the concept of unjamming motions to infinite (repetitive) systems. The difficulty with using standard Bloch wave (Fourier) decompositions is that impenetrability constraints are inequalities, rather than equalities, which cannot easily be decomposed into inequalities for each of the wavevectors using standard orthogonality relations. However, both the self stresses and the first-order flexes in a packing are solutions to systems of equalities and they can be decomposed into Fourier components.

For standard periodic boundary conditions, image particles move exactly the same way as their originals. In particular, for nonspherical particles, the orientation of image particles is identical to the orientation of the corresponding original particle. Extensions of periodic boundary conditions to break strict periodicity would enable one to calculate the stresses, strains, and stiffnesses appearing in micropolar elasticity theories, such as Cosserat Elasticity.

Finally, future work should consider the mathematics of jamming for packings of hard particles that are convex, but not necessarily smooth or strictly convex. In particular, particles with sharp corners and/or flat edges (e.g., squares, cubes, etc.) are of interest. We also believe that jamming in frictional packings, even for the case of spheres, is not well-understood. Mathematical programming, and in particular (linear) complementarity models, can be used to rigorously study mechanical equilibrium and stability in frictional packings.

14.1.4 Understanding (Dis)Order

In this work we expended considerable effort in understanding and defining rigorously the concept of jamming. However, an equally important task is to better quantify the concept of randomness, and in particular, to design order metrics for jammed packings of hard particles from "first principles". Only once this is done can the study of "random jammed" packings be freed of heuristics and assumptions.

In Chapter 13 we proposed a way to quantify disorder for jammed hard disk packings through the entropy of a partitioning of the sites of a triangular lattice into "large" and "small" ones. In particular, more ordered packings correspond to partitionings drawn from an ensemble of smaller entropy (degeneracy). The concept of entropy as an information measure is well-understood in the context of discrete systems such as lattices, and extending this to constrained but continuous systems such as jammed particle packings is a challenging task. Nonetheless, it offers hope of discovering a universal order metric, or in other words, understanding the very concept of order mathematically. One possible approach is to consider the ensemble of all jammed packings for a finite number of particles, give a unit weight to every packing, and consider the distribution of macroscopic properties such as density in that ensemble. It is not however obvious that equal weight should be given to every packing, and a different weighting may be appropriate in different contexts.

Understanding order will also help understand the geometric structure of MRJ packings. Why are disordered sphere packings prepared by a variety of algorithms at a density around $\phi \approx 0.64$, exhibiting very similar structure? In Chapter 13 we proposed that disordered binary disk packings may be considered as continuous deformations of the monodisperse triangular crystal. In Chapter 11 we suggested that disordered

hard sphere packings in dimension higher than two may be deformations (perturbations) of the crystal, however, an explicit construction of such a deformation is lacking. If disordered packings could indeed be described as perturbations of certain regular and easy-to-characterize packings, such as crystal packings or polytetrahedral packings in curved spaces, the structure of disordered packings could be quantified from "first principles".

14.1.5 Dynamics in Nearly Jammed Packings

Our work has made it clear that the study of nearly jammed packings requires the use of mathematical programming techniques. Unjamming motions at very high densities are collective and involve careful correlated motion of most if not all of the particles in the packing, and such collective motions cannot be efficiently found using Monte Carlo or Molecular Dynamics methods because of the long separation of time scales between local rattling motion and collective rearrangements. We designed rigorous algorithms to test for jamming in ideal packings of hard spheres and ellipsoids. It remains a challenge for the future to develop quadratic programming techniques for studying collective rearrangements in non-ideal (hypostatic) jammed packings of hard nonspherical particles, as discussed in Chapter 5. Even for spheres, a quantitative understanding of the performance of the randomized linear programming algorithm from Chapter 4 is lacking, for example, the probability of detecting unjamming motions when such motions exist has not been estimated for non-ideal packings. This probability should be estimated analytically for isostatic sphere packings by using the (simplex) jamming polytope picture we discussed extensively in this work.

In this work we formalized the notion of jamming as a permanent arrest of the packing structure. However, this rigorous formalism only applies to finite ideal packings. Understanding the nature of unjamming motions in large non-ideal packings quantitatively is needed in order to give insight into the dynamics of large nearly jammed packings. In particular, to understand why large nearly jammed packings remain within a jamming basin for long periods of time even though there are unjamming motions leading to different (possibly thermodynamically favored) basins. Understanding the dynamics near the jamming point is essential to the study of the kinetic glass transition. This work has identified two avenues of research that would be useful in this respect. For small systems it should be possible to rigorously identify all of the jammed states as well as the transition states and paths between the different jamming basins. This would enable a prediction of transition rates between different jammed packings, and comparison with numerical transition rates would test whether the assumptions underlying the transition rate theory (for example, ergodicity, separation of positions and velocities in phase space, etc.) are correct. Calculation of such transition rates at different densities could elucidate the origin and nature of the kinetic glass transition.

For large systems, rigorous and full enumeration of the transition states and paths will not be possible. However, a better understanding of the shape of the jamming polytope and its elongated directions would give insight into the dynamics near the jamming point. In this respect, it is believed that the multitude of near contacts in disordered jammed packings is an essential ingredient to the (mechanical) stability of jammed packings. Isostatic packings, as we find disordered sphere packings to be, are marginally rigid, that is, the removal of even a single contact destroys the jamming property. Furthermore, the extreme protuberance of the simplex jamming polytope makes some of the vertices of this polytope become unjamming motions at densities very close to the jamming density. Are these unjamming motions blocked by some of the multitude of near contacts present in disordered packings? An analysis of unjamming motions and dynamics in nearly jammed packings must focus on non-ideal packings and take into account the near contacts in addition to the true contacts.

14.1.6 Thermodynamics of Nearly Jammed Packings

More detailed analysis of the jamming basin $\mathcal{J}_{\Delta R}$ for isostatic sphere packings near the jamming point, going beyond the simplex picture, is needed in order to gain better insight into the thermodynamics of nearly jammed packings. In particular, it is important to develop an equation of state for nearly jammed packings that includes higher-order terms than the free-volume pressure. This will require considering both the influence of the multitude of near contacts and the influence of the curvature of the faces of the jamming "polytope".

Knowing the pressure as a function of jamming gap gives the bulk modulus of the packing, and in

particular, in the jamming limit $B \sim p^2$. Calculating the shear modulus of nearly jammed packings is an important open problem. Calculations using the simplex polytope picture show that the shear modulus scales with pressure with a power lower than that of the bulk modulus, that is, that $G/B \to 0$. Some recent theoretical work on soft spheres has suggested that $G \sim p^{3/2}$ [173]. Establishing such a result for hard sphere packings is a challenging but important problem. Characterizing the actual non-anharmonic elastic response of nearly jammed packings is even more important, since the very definition of elastic moduli relies on assumptions about harmonic macroscopic elastic behavior. Such assumptions rigorously break down in the jamming limit, since the mechanical behavior of jammed packings cannot be described by linear elasticity [111]. However, experiments and simulations have shown that sufficiently far from the jamming point there is some merit to considering the macroscopic response of dense liquids/glasses through linear elasticity.

Formally, the single-cell occupancy (SOC) constraints described in Chapter 6 give a rigorous definition of the free energy of a nearly jammed packing. However, it is not clear whether the free energy or its derivatives (such as pressure or elastic moduli) can be calculated analytically for SOC solids and used to approximate, in a controlled manner, the properties of the unconstrained system. The algorithms we designed in this work can be used to probe these issues numerically, and the jamming polytope picture extended to SOC-constrained systems.

14.1.7 Further Applications

The algorithms and theories we designed can be applied to the study of a wide range of systems. The BCMD algorithm we developed should be applied to the determination of free energies of solid phases and thus the phase diagrams for various particle shapes, and in particular ellipsoids and superellipsoids (especially shapes approaching rectangles, cubes, cylinders, etc.). The identification of the solid phases, i.e., the densest packings, for different particle shapes can be carried out using molecular dynamics as we did in Chapter 8 for ellipsoids.

Statistical characterization of random ellipse and random ellipsoid packings, similar to the one we carried out for spheres, is still lacking. In this work we focused on nearly spherical ellipsoids, however, it is also important to understand the behavior of ellipsoids with large aspect ratios. Studying other particle shapes, and in particular superellipsoids, may give surprising new insights about packings. One still-unresolved question of interest involves shape dependence of the packing-fraction ratio for amorphous *vs.* crystalline arrangements of identical hard particles. It will be enlightening eventually to establish for various particle shapes: (a) what particle shapes yield the attainable lower and upper limits for this ratio, (b) whether the upper limit can equal or exceed unity, and (c) whether there are useful correlations between this ratio and features of the equilibrium phase diagrams or glassy behavior.

Bibliography

[1] A. Mehta, editor. *Granular Matter*. Springer-Verlag, New York, 1994.

[2] H. Hinrichsen and D. E. Wolf, editors. *The Physics of Granular Media*. Wiley-VCH, Berlin, 2004.

[3] H. A. Makse, J. Brujic, and S. F. Edwards. *The Physics of Granular Media*, chapter Statistical Mechanics of Jammed Matter, pages 45–86. John Wiley & Sons, 2004.

[4] R. Zallen. *The Physics of Amorphous Solids*. Wiley, New York, 1983.

[5] J. P. Hansen and I. R. McDonald. *Theory of Simple Liquids*. Academic Press, New York, 1986.

[6] S. Torquato. *Random Heterogeneous Materials: Microstructure and Macroscopic Properties*. Springer-Verlag, New York, 2002.

[7] L. Fejes Tóth. *Regular figures*. Pergamon Press, 1964.

[8] T. Aste and D. Weaire. *The Pursuit of Perfect Packing*. IOP Publishing, 2000.

[9] J. H. Conway and N. J. A. Sloane. *Sphere Packings, Lattices, and Groups*. Springer-Verlag, New York, 3rd edition, 1999.

[10] C. Zong. *Sphere Packings*. Springer-Verlag, New York, 1999.

[11] A. Donev, S. Torquato, and F. H. Stillinger. Neighbor List Collision-Driven Molecular Dynamics Simulation for Nonspherical Particles: I. Algorithmic Details II. Applications to Ellipses and Ellipsoids. *J. Comp. Phys.*, 202(2):737–764, 765–793, 2005.

[12] B. D. Lubachevsky and F. H. Stillinger. Geometric Properties of Random Disk Packings. *J. Stat. Phys.*, 60:561–583, 1990.

[13] B. D. Lubachevsky, F. H. Stillinger, and E. N. Pinson. Disks vs. Spheres: Contrasting Properties of Random Packings. *J. Stat. Phys.*, 64:501–525, 1991.

[14] J. D. Sherwood. Packing of spheroids in three-dimensional space by random sequential addition. *J. Phys. A*, 30:L839–L843, 1997.

[15] S. Torquato and F. H. Stillinger. Multiplicity of Generation, Selection, and Classification Procedures for Jammed Hard-Particle Packings. *J. Phys. Chem. B*, 105:11849–11853, 2001.

[16] J. Brujic, S. F. Edwards, I. Hopkinson, and H. A. Makse. Measuring the distribution of interdroplet forces in a compressed emulsion system. *Physica A*, 327:201–212, 2003.

[17] P. Richard, L. Oger, J. P. Troadec, and A. Gervois. Tessalation of binary assemblies of spheres. *Physica A*, 259:205–221, 1998.

[18] V. A. Luchnikov, N. N. Medvedev, L. Oger, and J. P. Troadec. Voronoi-delaunay analysis of voids in systems of nonspherical particles. *Phys. Rev. E*, 59:7205–7212, 1999.

[19] G. T. Nolan and P. E. Kavanagh. Computer simulation of random packing of hard spheres. *Powder Tech.*, 72:149–155, 1992.

[20] C. H. Bennett. Serially Depositied Amorphous Aggregates of Hard Spheres. *J. Appl. Phys.*, 32(6):2727–2734, 1972.

[21] W. S. Jodrey and E. M. Tory. Computer simulation of close random packing of equal spheres. *Phys. Rev. A*, 32(4):2347–2351, 1985.

[22] E. L. Hinrichsen, J. Feder, and T. Jossang. Random packing of disks in two dimensions. *Phys. Rev. A*, 41(8):4199–4209, 1990.

[23] M. Bargiel and J. Moscinski. C-language program for the irregular close packing of hard spheres. *Comp. Phys. Comm.*, 64:183–192, 1991.

[24] A. Bezrukov, M. Bargiel, and D. Stoyan. Statistical analysis of simulated random packings of spheres. *Part. Part. Syst. Charact.*, 19:111–118, 2002.

[25] A. S. Clarke and J. D. Wiley. Numerical simulation of a dense random packing of a binary mixture of hard spheres. *Phys. Rev. B.*, 35(14):7350–7356, 1987.

[26] P. Richard, A. Gervois, L. Oger, and J.-P. Troadec. Order and disorder in hard-sphere packings. *Europhys. Lett.*, 48(4):415–420, 1999.

[27] S. R. Williams and A. P. Philipse. Random Packings of Spheres and Spherocylinders Simulated by Mechanical Contraction. *Phys. Rev. E*, 67:051301, 2003.

[28] A. Zinchenko. Algorithm for Random Close Packing of Spheres with Periodic Boundary Conditions. *J. Comp. Phys.*, 114:298–307, 1994.

[29] F. H. Stillinger and T. A. Weber. Inherent structure of liquids in the hard-sphere limit. *J. Chem. Phys.*, 83(9):4767–4775, 1985.

[30] D. A. Kottwitz. The densest packing of equal circles on a sphere. *Acta Cryst.*, 47:158–165, 1991.

[31] C. S. O'Hern, S. A. Langer, A. J. Liu, and S. R. Nagel. Random Packings of Frictionless Particles. *Phys. Rev. Lett.*, 88(7):075507–1, 2002.

[32] L. E. Silbert, D. Ertas, G. S. Grest, T. C. Halsey, and D. Levine. Geometry of frictionless and frictional sphere packings. *Phys. Rev. E*, 65:031304–1, 2002.

[33] H. A. Makse, D. L. Johnson, and L. M. Schwartz. Packing of compressible granular materials. *Phys. Rev. Lett.*, 84(18):4160–4163, 2000.

[34] B. J. Buchalter and R. M. Bradley. Orientational order in amorphous packings of ellipsoids. *Europhys. Lett.*, 26(3):159–164, 1994.

[35] D. Coelho, J.-F. Thovert, and P. M. Adler. Geometrical and transport properties of random packings of spheres and aspherical particles. *Phys. Rev. E*, 55(2):1959–1978, 1997.

[36] R. Jullien, P. Jund, D. Caprion, and D. Quitmann. Computer investigation of long-range correlations and local order in random packings of spheres. *Phys. Rev. E*, 54:6035–6041, 1996.

[37] B. D. Lubachevsky. How to Simulate Billiards and Similar Systems. *J. Comp. Phys.*, 94:255–283, 1991.

[38] B. J. Buchalter and R. M. Bradley. Orientational order in random packings of ellipses. *Phys. Rev. A*, 46(6):3046–3056, 1992.

[39] T. C. Hales. A proof of the Kepler conjecture. *Ann. Math.*, 2006.

[40] G. D. Scott and D. M. Kilgour. The density of random close packing of spheres. *J. Phys. D*, 2:863–866, 1969.

[41] J. G. Berryman. Random Close Packing of Hard Spheres and Disks. *Phys. Rev. A*, 27(2):1053–1061, 1983. See also references contained therein.

[42] J. D. Bernal. *Liquids: Structure, Properties, Solid Interactions (T. J. Hughel, editor)*, pages 25–50. Elsevier, New York, 1965.

[43] M. P. Allen, G. T. Evans, D. Frenkel, and B. M. Mulder. Hard convex body fluids. *Advances in Chemical Physics*, 86:1–166, 1993.

[44] D. Frenkel and A. J. C. Ladd. New Monte Carlo method to compute the free energy of arbitrary solids. Application to the fcc and hcp phases of hard spheres. *J. Chem. Phys.*, 81:3188–3193, October 1984.

[45] S.-C. Mau and D. A. Huse. Stacking entropy of hard-sphere crystals. *Phys. Rev. E*, 59:4396–4401, 1999.

[46] H. Koch, C. Radin, and L. Sadun. Most stable structure for hard spheres. *Phys. Rev. E*, 72:016708, 2005.

[47] J. M. Kosterlitz and D. J. Thouless. Ordering, metastability and phase transitions in two-dimensional systems. *J. Phys.*, C6:1181–1203, 1973.

[48] B. I. Halperin and D. R. Nelson. Dislocation-mediated melting in two dimensions. *Phys. Rev. B*, 19(5):2457–2484, 1979.

[49] A. P. Young. Melting and the vector Coulomb gas in two dimensions. *Phys. Rev. B*, 19:1855–1866, 1979.

[50] R. J. Speedy. Glass transition in hard disc mixtures. *J. Chem. Phys.*, 110:4559–4565, 1999.

[51] D. N. Perera and P. Harrowell. Stability and structure of a supercooled liquid mixture in two dimensions. *Phys. Rev. E*, 59:5721, 1999.

[52] Z. W. Salsburg and W. W. Wood. Equation of state of classical hard spheres at high density. *J. Chem. Phys.*, 37:798, 1962.

[53] R. J. Speedy. Pressure of the Metastable Hard-Sphere Fluid. *J. Phys. Condens. Matter*, 9:8591–8599, 1997.

[54] R. J. Speedy. Pressure and entropy of hard-sphere crystals. *J. Phys. Condens. Matter*, 10:4387–4391, 1998.

[55] J. K. Percus and G. J. Yevick. Analysis of Classical Statistical Mechanics by Means of Collective

Coordinates. *Phys. Rev.*, 110:1–13, 1958.

[56] A. R. Kansal, S. Torquato, and F. H. Stillinger. Diversity of Order and Densities in Jammed Hard-Particle Packings. *Phys. Rev. E*, 66:041109, 2002.

[57] A. Donev, F. H. Stillinger, and S. Torquato. Do Binary Hard Disks Exhibit an Ideal Glass Transition? *Phys. Rev. Lett.*, 96(22):225502, 2006.

[58] S. Torquato, T. M. Truskett, and P. G. Debenedetti. Is random close packing of spheres well defined? *Phys. Rev. Lett.*, 84:2064–2067, 2000.

[59] F. H. Stillinger, H. Sakai, and S. Torquato. Lattice-based random jammed configurations for hard particles. *Phys. Rev. E*, 67:031107, 2003.

[60] A. R. Kansal, S. Torquato, and F. H. Stillinger. Computer generation of dense polydisperse sphere packings. *J. Chem. Phys.*, 117(18):8212–8218, 2002.

[61] C. N. Likos and C. L. Henley. Complex alloy phases for binary hard-disc mixtures. *Phil. Mag. B*, 68(1):85–113, 1993.

[62] K. L. Kristiansen, A. Wouterse, and A. Philipse. Simulation of random packing of binary sphere mixtures by mechanical contraction. *Physica A*, 358:249–262, 2005.

[63] A. Yang, C. T. Miller, and L. D. Turcoliver. Simulation of correlated and uncorrelated packing of random size spheres. *Phys. Rev. E*, 53(2):1516–1524, 1996.

[64] D. He, N. N. Ekere, and L. Cai. Computer simulation of random packing of unequal particles. *Phys. Rev. E*, 60:7098, 1999.

[65] E. Santiso and E. A. Muller. Dense packing of binary and polydisperse hard spheres. *Mol. Phys.*, 100(9):2461–2469, 2002.

[66] N. Xu, J. Blawzdziewicz, and C.S. O'Hern. Reexamination of Random Close Packing: Ways to pack frictionless disks. *Phys. Rev. E*, 71:061306, 2005.

[67] O. Ouche, F. H. Stillinger, and S. Torquato. Concerning Maximal Packing Arrangements of Binary Disk Mixtures. *Physica A*, 342:428, 2004.

[68] N. Hunt, R. Jardine, and P. Bartlett. Superlattice formation in mixtures of hard-sphere colloids. *Phys. Rev. E*, 62(1):900–913, 2000.

[69] S.-E. Phan, J. Zhu, W. B. Russel, and P. M. Chaikin. Effects of polydispersity on hard sphere crystals. *J. Chem. Phys.*, 108:9789, 1998.

[70] P. M. Chaikin. *Soft and Fragile Matter, Nonequilibrium Dynamics, Metastability and Flow (M. E. Cates and M. R. Evans, editors)*, chapter Thermodynamics and Hydrodynamics of Hard Spheres; the role of gravity, pages 315–348. Institute of Physics Publishing, London, 2000.

[71] A. Donev, S. Torquato, F. H. Stillinger, and R. Connelly. A Linear Programming Algorithm to Test for Jamming in Hard-Sphere Packings. *J. Comp. Phys.*, 197(1):139–166, 2004.

[72] A. Donev, S. Torquato, F. H. Stillinger, and R. Connelly. Jamming in Hard Sphere and Disk Packings. *J. App. Phys.*, 95(3):989, 2004.

[73] R. Connelly and W. Whiteley. Second-Order Rigidity and Prestress Stability for Tensegrity Frameworks. *SIAM Journal of Discrete Mathematics*, 9(3):453–491, 1996.

[74] A. Donev, R. Connelly, F. H. Stillinger, and S. Torquato. Hypostatic Jammed Packings of Hard Ellipsoids. In preparation, 2006.

[75] A. Donev, S. Torquato, and F. H. Stillinger. Pair Correlation Function Characteristics of Nearly Jammed Disordered and Ordered Hard-Sphere Packings. *Phys. Rev. E*, 71:011105, 2005.

[76] A. Donev, S. Torquato, and F. H. Stillinger. Unexpected Density Fluctuations in Jammed Disordered Sphere Packings. *Phys. Rev. Lett.*, 95(9):090604, 2005.

[77] M. Skoge, A. Donev, F. H. Stillinger, and S. Torquato. Packing Hyperspheres in High-Dimensional Euclidean Spaces. Submitted to *Phys. Rev. E*, 2006.

[78] A. Donev, I. Cisse, D. Sachs, E. A. Variano, F. H. Stillinger, R. Connelly, S. Torquato, and P. M. Chaikin. Improving the Density of Jammed Disordered Packings using Ellipsoids. *Science*, 303:990–993, 2004.

[79] W. Man, A. Donev, F. H. Stillinger, M. Sullivan, William B. Russel, D. Heeger, S. Inati, S. Torquato, and P. M. Chaikin. Experiments on Random Packing of Ellipsoids. *Phys. Rev. Lett.*, 94:198001, 2005.

[80] A. Donev, P. M. Chaikin, F. H. Stillinger, and S. Torquato. Unusually Dense Crystal Packings of Ellipsoids. *Phys. Rev. Lett.*, 92:255506, 2004.

[81] A. Donev, J. Burton, F. H. Stillinger, and S. Torquato. Tetratic Order in the Phase Behavior of a Hard-Rectangle System. *Phys. Rev. B*, 73:054109, 2006.

[82] A. Donev, F. H. Stillinger, and S. Torquato. Calculating the Free Energy of Nearly Jammed Hard-Particle Packings Using Molecular Dynamics. Submitted to *J. Comp. Phys*, 2006.

[83] A. Donev, F. H. Stillinger, and S. Torquato. Configurational Entropy of Binary Hard-Disk Glasses. In preparation, 2006.

[84] H. Lütkepohl. *Handbook of Matrices*. John Wiley and Sons, 1997.

[85] Ronald Goldman. *Cross Product in Four Dimensions and Beyond*, volume II of *Matrix Gems*, chapter II.8, pages 84–88. Academic Press, 1992.

[86] M. P. Allen and D. J. Tildesley. *Computer Simulations of Liquids*. Oxford Science Publications, 1987.

[87] David Eberly. Rotation representations and performance issues. Magic Software, Inc. `http://www.magic-software.com`.

[88] H. Schaub, P. Tsiotras, and J. L. Junkins. Principal Rotation Representations of NxN Orthogonal Matrices. *Int. J. Eng. Sci.*, 33(15):2277–2295, 1995.

[89] J. W. Perram and M. S. Wertheim. Statistical Mechanics of Hard Ellipsoids. I. Overlap Algorithm and the Contact Function. *J. Comp. Phys.*, 58:409–416, 1985.

[90] J. W. Perram, J. Rasmussen, E. Praestgaard, and J. L. Lebowitz. Ellipsoid contact potential: Theory and relation to overlap potentials. *Phys. Rev. E*, 54(6):6565–6572, 1996.

[91] L. F. Shampine, I. Gladwell, and R. W. Brankin. Reliable Solution of Special Event Location Problems for ODEs. *TOMS*, 17(1):11–25, 1991.

[92] J. Vieillard-Baron. Phase transitions of the classical hard-ellipse system. *J. Chem. Phys.*, 56:4729–4744, 1972.

[93] W. Wang, J. Wang, and M.-S. Kim. An Algebraic Condition for the Separation of Two Ellipsoids. *Comp. Aid. Geom. Design*, 18:531–539, 2001.

[94] E. Rimon and S. P. Boyd. Obstacle collision detection using best ellipsoid fit. *J. Intelligent and Robotic Systems*, 18:105–126, 1997.

[95] X. Lin and T. T. Ng. Contact detection algorithms for three-dimensional ellipsoids in discrete element modeling. *Int. J. Num. Anal. Meth. Geomech.*, 19:653–659, 1995.

[96] David Eberly. I. distance between ellipses in 2d II. distance between two ellipses in 3d. Magic Software, Inc. `http://www.magic-software.com`.

[97] David Eberly. I. intersection of ellipses II. intersection of ellipsoids. Magic Software, Inc. `http://www.magic-software.com`.

[98] A. Lin and S. P. Han. On the distance between two ellipsoids. *SIAM J. Optim.*, 13:298–308, 2002.

[99] and S. N. Yaliraki L. Paramonov. The directional contact distance of two ellipsoids: Coarse-grained potentials for anisotropic interactions. *J. Chem. Phys*, 123:194111, 2005.

[100] A. H. Barr. Superquadrics and Angle Preserving Transformations. *IEEE Comp. Graph. and App.*, 1(1):11–23, 1981.

[101] A. Jaklič, A. Leonardis, and F. Solina. *Segmentation and Recovery of Superquadrics*, volume 20 of *Computational imaging and vision*. Kluwer, Dordrecth, 2000.

[102] B. J. Alder and T. E. Wainwright. Studies in molecular dynamics. I. General method. *J. Chem. Phys.*, 31:459, 1959.

[103] Y. A. Houndonougbo, B. B. Laird, and B. J. Leimkuhler. Molecular dynamics algorithms for mixed hard-core/continuous potentials. *Mol. Phys.*, 98:309–316, 1999.

[104] H. Sigurgeirsson, A. Stuart, and W.-L. Wan. Algorithms for Particle-Field Simulations with Collisions. *J. Comp. Phys.*, 172:766–807, 2001.

[105] D. Frenkel and J. F. Maguire. Molecular Dynamics Study of the Dynamical Properties of an Assembly of Infinitely Thin Hard Rods. *Mol. Phys.*, 49(3):503–541, 1983.

[106] M. P. Allen, D. Frenkel, and J. Talbot. Molecular Dynamics Simulation Using Hard Particles. *Comp. Phys. Rep.*, 9:301–353, 1989.

[107] M. Parinello and A. Rahman. Polymorphic transitions in single crystals: A new molecular dynamics method. *J. Appl. Phys.*, 52(12):7182–7190, 1981.

[108] D. W. Rebertus and K. M. Sando. Molecular Dynamics Simulation of a Fluid of Hard Spherocylinders. *J. Chem. Phys.*, 67(6):2585–2590, 1977.

[109] B. D. Lubachevsky. Simulating colliding rigid disks in parallel using bounded lag without time warp. *Distributed Simulation, SCS Simulation Series*, 22(1):194–202, 1990.

[110] S. Miller and S. Luding. Event-driven Molecular Dynamics in Parallel. *J. Comput. Phys.*, 193(1):306–316, 2004.

[111] S. Torquato, A. Donev, and F. H. Stillinger. Breakdown of Elasticity Theory for Jammed Hard-Particle Packings: Conical Nonlinear Constitutive Theory. *Int. J. Solids Structures*, 40(25):7143 – 7153, 2003.

[112] G. Germano and F. Schmid. Simulation of nematic-isotropic phase coexistence in liquid crystals under shear. In H. Rollnik and D. Wolf, editors, *NIC Symposium 2004, NIC Series Vol. 20*, page 311. Forschungszentrum Jülich, 2004.

[113] A. W. Lees and S. F. Edwards. The Computer Study of Transport Processes Under Extreme Conditions. *J. Phys. C.*, 5:1921–1929, 1972.

[114] M. Yoneya K. M. Aoki and H. Yokoyama. Molecular dynamic simulation methods for anisotropic liquids. *J. Chem. Phys.*, 120(12):5576–5582, 2004.

[115] S.C. Harvey, R. K.-Z. Tan, and T. E. Cheatham III. The flying ice cube: Velocity rescaling in molecular dynamics leads to violation of energy equipartition. *J. Comp. Chem.*, 19:726–740, 1998.

[116] D. Frenkel and B. Smit. *Understanding Molecular Simulation*. Academic Press, 2002.

[117] A. Ferguson, B. Fisher, and B. Chakraborty. Impulse distributions in dense granular flows: Signatures of large-scale spatial structures. *Europhys. Lett.*, 66(2):277–283, 2004.

[118] M. Marin and P. Cordero. Hashing-Cell Combination for Boundless Space Event-Driven Molecular Dynamics. pages 315–318, Krakow, Poland, September 1996. 8th Joint EPS-APS International Conference on Physics Computing (Eds. P. Borcherds and M. Bubak), World Scientific Publishing Company.

[119] H. Bekker. Unification of Box Shapes in Molecular Simulations. *J. Comp. Chem.*, 18(15):1930–1942, 1997.

[120] M. Isobe. Simple and Efficient Algorithm for Large Scale Molecular Dynamics Simulation in Hard Disk Systems. *Int. J. Mod. Phys. C*, 10(7):1281–1293, 1999.

[121] L. G. Casado, I. Garcia, and Ya. D. Sergeyev. Interval Algorithms for Finding The Minimal Root in a Set of Multiextremal One-Dimenisional Nondiferentiable Functions. *SIAM J. Sci. Comput.*, 24(2):359–376, 2002.

[122] P. M. Hubbard. Approximating Polyhedra with Spheres for Time-Critical Collision Detection. *ACM Trans. Graph.*, 15(3):179–210, 1996.

[123] M. H. Overmars and A. F. van der Stappen. Range Searching and Point Location among Fat Objects. *J. Algorithms*, 21(3):629–656, 1996.

[124] J. J. Erpenbeck and W. W. Wood. *Statistical mechanics B: Modern theoretical chemistry (B. J. Berne, editor)*, volume 6, chapter Molecular dynamics techniques for hard core systems, pages 1–40. Institute of Physics Publishing, London, 1977.

[125] D. C. Rapaport. The Event Scheduling Problem in Molecular Dynamics Simulation. *J. Comp. Phys.*, 34:184, 1980.

[126] A. T. Krantz. Analysis of an Efficient Algorithm for the Hard-Sphere Problem. *TOMACS*, 6(3):185–209, 1996.

[127] K. Shida and Y. Anzai. Reduction of the event list for molecular dynamic simulation. *Comp. Phys. Comm.*, 69:317–329, 1992.

[128] M. Marin, D. Risso, and P. Cordero. Efficient Algorithms for Many-Body Hard Particle Molecular Dynamics. *J. Comp. Phys.*, 109:306–317, 1993.

[129] B. C. Vemuri, Y. Cao, and L. Chen. Fast Collision Detection Algorithms with Applications to Particle Flow. *Computer Graphics Forum*, 17(2):121–134, 1998.

[130] P. Daponte, D. Grimaldi, A. Molinaro, and Ya. D. Sergeyev. An Algorithm for Finding the Zero Crossing of Time Signals with Lipschitzean Derivatives. *Measurement*, 16:37–49, 1995.

[131] L. T. Watson, M. Sosonkina, R. C. Melville, A. P. Morgan, and H. F. Walker. Algorithm 777: HOMPACK90: A Suite of Fortran 90 Codes for Globally Convergent Homotopy Algorithms. *TOMS*, 23(4):514–549, 1997.

[132] L. F. Shampine. *Numerical Solution of Ordinary Differential Equations*. Chapman and Hall, 1994.

[133] S. Gottschalk. *Collision Queries using Oriented Bounding Boxes*. PhD thesis, UNC Chapel Hill, Department of Computer Science, 2000.

[134] David Eberly. I. dynamic collision detection using oriented bounding boxes II. intersection of objects with linear and angular velocities using oriented bounding boxes. Magic Software, Inc. `http://www.magic-software.com`.

[135] A. Donev. `http://atom.princeton.edu/donev/Packing`, homepage for the sphere packing project, with useful supplementary materials.

[136] R. Connelly, K. Bezdek, and A. Bezdek. Finite and Uniform Stability of Sphere Packings. *Discrete and Computational Geometry*, 20:111–130, 1998.

[137] M. F. Thorpe and P. M. Duxbury, editors. *Rigidity Theory and Applications*. Fundamental Materials Research. Kluwer/Plenum, 1999.

[138] R. Connelly. Rigid Circle And Sphere Packings, *Part I. Finite Packings*. *Structural Topology*, 14:43–60, 1988. See also Ref. [142].

[139] W. Uhler and R. Schilling. A local description of stable $2D$ random packings. *J. Phys. C*, 18:L979–L983, 1985.

[140] D. E. G. Williams. Stability in 2D Random Packing. *J. Phys. C*, 18:L181–L184, 1985.

[141] E. A. J. F. Peters, M. Kollmann, Th. M. A. O. M. Barenbrug, and A. P. Philipse. Caging of a d-dimensional sphere and its relevance for the random dense sphere packing. *Phys. Rev. E*, 63:021404, 2001.

[142] R. Connelly. Rigid Circle And Sphere Packings *Part II. Infinite Packings*. *Structural Topology*, 16:57–76, 1991. Second part of Ref. [138].

[143] J. N. Roux. Geometric Origin of Mechanical Properties of Granular Materials. *Phys. Rev. E*, 61(6):6802–6836, 2000.

[144] S. F. Edwards and D. V. Grinev. The Tensorial Formulation of Volume Function for Packings of Particles. *Chem. Eng. Sci.*, 56:5451–5455, 2001.

[145] F. H. Stillinger and Z. W. Salsburg. Limiting Polytope Geometry for Rigid Rods, Disks, and Spheres. *J. Stat. Physics*, 1(1):179, 1969.

[146] R. Vanderbei. *Linear Programming:* Foundations and Extensions. Kluwer, 1997.

[147] A. Donev and S. Torquato. Energy-Efficient Actuation in Infinite Lattice Structures. *J. Mech. Phys. Solids*, 51(8):1459–1475, 2003.

[148] H. P. F. Swinnerton-Dyer. Extremal lattices of convex bodies. *Proc. Cambridge Philos. Soc.*, 49:161–162, 1953.

[149] M. Lätzel, S. Luding, and H. J. Herrmann. From discontinuous models towards a continuum description. In P. A. Vermeer, S. Diebels, W. Ehlers, H. J. Herrmann, S. Luding, and E. Ramm, editors, *Continuous and Discontinuous Modelling of Cohesive Frictional Materials*, pages 215–230, Berlin, 2001. Springer.

[150] N. P. Kruyt and L. Rothenburg. Micromechanical definition of strain tensor for Granular Materials. *ASME Journal of Applied Mechanics*, 118:706–711, 1996.

[151] L. Rothenburg and A. P. S. Selvadurai. A micromechanical definition of the Cauchy stress tensor for particulate media. In A. P. S. Selvadurai, editor, *Mechanics of Structured Media*, pages 469–486. Elsevier, Amsterdam, 1981.

[152] M. Madadi, O. Tsoungui, M. Lätzel, and S. Luding. On the fabric tensor of polydisperse granular media in 2D. *Int. J. Sol. Struct.*, 41(9-10):2563–2580, 2004.

[153] WWW. The general purpose interior-point `LOQO` optimization library is not public-domain, but can be tried at `http://orfe.princeton.edu/~loqo/`. The public-domain `PCx` library implements interior point linear programming algorithms and can be found at `http://www-fp.mcs.anl.gov/otc/Tools/PCx/`.

[154] T. J. Quickenden and G. K. Tan. Random Packing in Two Dimensions and the Structure of Monolayers. *J. Colloid Interface Sci.*, 48:382, 1974.

[155] G. Y. Onoda and E. G. Liniger. Random loose packings of uniform spheres and the dilatancy onset. *Phys. Rev. Lett.*, 64:2727–2730, 1990.

[156] K. J. Dong, R. Y. Yang, R. P. Zou, and A. B. Yu. Role of interparticle forces in the formation of random loose packing. *Phys. Rev. Lett.*, 96:145505, 2006.

[157] C. H. Liu, S. R. Nagel, and D. A. S. Nonlinear dynamics: Jamming is not just cool any more. *Nature*, 396:21–22, 1998.

[158] H. A. Makse and J. Kurchan. Testing the thermodynamic approach to granular matter with a numerical model of a decisive experiment. *Nature*, 415:614–617, 2002.

[159] S. F. Edwards. *Granular Matter (A. Mehta, editor)*, chapter The Role of Entropy in the Specification of a Powder, pages 121–140. Springer-Verlag, New York, 1994.

[160] S. Alexander. Amorphous Solids: Their Structure, Lattice Dynamics and Elasticity. *Phys. Rep.*, 296:65–236, 1998.

[161] C. F. Moukarzel. Isostatic phase transition and instability in stiff granular materials. *Phys. Rev. Lett.*, 81:1634–1638, 1998.

[162] A. V. Tkachenko and T. A. Witten. Stress propagation through frictionless granular material. *Phys. Rev. E*, 60:687–696, 1999.

[163] S. F. Edwards. The equations of stress in a granular material. *Physica A*, 249:226–231, 1998.

[164] C. S. O'Hern, E. Silbert, A. J. Liu, and S. R. Nagel. Jamming at zero temperature and zero applied stress: The epitome of disorder. *Phys. Rev. E.*, 68:011306, 2003.

[165] R. Connelly. Rigidity and Energy. *Invent. Math.*, 66:11–33, 1982.

[166] 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.

[167] J. M Rallison and S. E. Harding. Excluded volume for pairs of triaxial ellipsoids at dominant Brownian motion. *J. Colloid. Interface Sci.*, 103:284–289, 1985.

[168] G. S. Singh and B. Kumar. Geometry of hard ellipsoidal fluids and their virial coefficients. *J. Chem. Phys.*, 105:2429–2435, 1996.

[169] C. S. O'Hern, S. A. Langer, A. J. Liu, and S. R. Nagel. Force Distributions near Jamming and Glass Transitions. *Phys. Rev. Lett.*, 86:111–114, 2001.

[170] T. M. Truskett, S. Torquato, and P. G. Debenedetti. Towards a quantification of disorder in materials: Distinguishing equilibrium and glassy sphere packings. *Phys. Rev. E*, 62(1):993–1001, 2000.

[171] C. F. Moukarzel. Isostaticity in granular matter. *Granular Matter*, 3:41–52, 2001.

[172] M. Todd. Semidefinite Optimization. *Acta Numerica*, 10:515–560, 2001.

[173] M. Wyart. On the rigidity of amorphous solids. arXiv:cond-mat/0512155, 2005.

[174] C. Gotsman and S. Toledo. On the computation of null spaces of sparse rectangular matrices. 2005.

[175] S. M. Rump. Verification of Positive Definiteness. *BIT Num. Math.*, 43(1):001–018, 2003.

[176] A. Donev, S. Torquato, F. H. Stillinger, and R. Connelly. Comment on "Jamming at zero temperature and zero applied stress: The epitome of disorder". *Phys. Rev. E*, 70:043301, 2004.

[177] M. Wyart, L. E. Silbert, S. R. Nagel, and T. A. Witten. Effects of compression on the vibrational modes of marginally jammed solids. *Phys. Rev. E*, 73:051306, 2005.

[178] J. M. Rickman and R. LeSar. Free-energy calculations in materials research. *Ann. Rev. of Mat. Res.*, 32:195–217, 2002.

[179] L. V. Woodcock. Computation of the free energy of alternative crystal structures of hard spheres. *Faraday Discuss.*, 106:325–338, 1997.

[180] F. Sciortino, W. Kob, and P. Tartaglia. Inherent Structure Entropy of Supercooled Liquids. *Phys. Rev. Lett.*, 83:3214–3217, 1999.

[181] S. Sastry. Evaluation of the configurational entropy of a model liquid from computer simulations. *J. Phys.: Condens. Matter*, 12(29):6515–6523, 2000.

[182] Q. Yan, T. S. Jain, and J. J. de Pablo. Density-of-States Monte Carlo Simulation of a Binary Glass. *Phys. Rev. Lett.*, 92:235701, 2004.

[183] L. Angelani, G. Foffi, and F. Sciortino. Configurational entropy of hard spheres. arXiv:cond-mat/0506447, 2005.

[184] R. J. Speedy. The Entropy of a Glass. *Mol. Phys.*, 80(5):1105–1120, 1993.

[185] M. Simonovits. How to Compute the Volume in High Dimensions. *Math. Program. Ser. B*, 97:337–374, 2003.

[186] B. Bueler, A. Enge, and K. Fukuda. *Polytopes-Combinatorics and Computation*, chapter Exact volume computation for polytopes: A practical study. Birhauser Verlag, Basel, 2000.

[187] A. Donev, F. H. Stillinger, and S. Torquato. Free Energy of Nearly Isostatic Hard-Sphere Packings. In preparation, 2006.

[188] L. E. Silbert, A. J. Liu, and S. R. Nagel. Vibrations and diverging length scales near the unjamming transition. *Phys. Rev. Lett.*, 95:098301, 2005.

[189] A. D. Bruce, A. N. Jackson, G. J. Ackland, and N. B. Wilding. Lattice-switch Monte Carlo method. *Phys. Rev. E*, 61:906, 2000.

[190] N. Simányi. *Encyclopedia of Mathematical Sciences*, volume 101 of *Mathematical Physics II*, chapter Hard Ball Systems and Semi-Dispersive Billiards, pages 51–88. Springer Verlag, 2000.

[191] L. Lovasz and S. Vempala. Simulated Annealing in Convex Bodies and an $O(n^4)$ Volume Algorithm. MSR-TR 2003-31, Microsoft Research, 2003.

[192] J. M. Polson, E. Trizac, S.Pronk, and D.Frenkel. Finite-size corrections to the free energies of crystalline solids. *J. Chem. Phys.*, 112:5339–5342, 2000.

[193] N. F. Carnahan and K. E. Starling. Equation of State for Nonattracting Rigid Spheres. *J. Chem. Phys.*, 51:635–636, 1969.

[194] W. G. Hoover and F. H. Ree. Melting Transition and Communal Entropy for Hard Spheres. *J. Chem. Phys.*, 49(8):3609, 1968.

[195] J. B. Knight, C. G. Fandrich, C. N. Lau, H. M. Jaeger, and S. R. Nagel. Density Relaxation in a Vibrated Granular Material. *Phys. Rev. E*, 51:3957, 1995.

[196] D. L. Blair, N. W. Mueggenburg, A. H. Marshall, H. M. Jaeger, and S. R. Nagel. Force distributions in three-dimensional granular assemblies: Effects of packing order and interparticle friction. *Phys. Rev. E*, 63:041304, 2001.

[197] J. L. Finney. Random packings and the structure of simple liquids. I. The geometry of random close packing. *Proc. R. Soc. Lond. A*, 319:479–493, 1970.

[198] G. T. Seidler, G. Martinez, L. H. Seeley, K. H. Kim, E. A. Behne, S. Zaranek, B. D. Chapman, S. M. Heald, and D. L. Brewe. Granule-by-granule reconstruction of a sandpile from x-ray microtomography data. *Phys. Rev. E*, 62:8175–8181, 2000.

[199] T. Aste, M. Saadatfar, A. Sakellariou, and T. J. Senden. Investigating the geometrical structure of disordered sphere packings. *Physica A*, 339:16–23, 2004.

[200] V. Narayan, N. Menon, and S. Ramaswamy. Nonequilibrium steady states in a vibrated-rod monolayer: tetratic, nematic and smectic correlations. `cond-mat/0510082`, 2005.

[201] J. D. Bernal and J. Mason. Co-ordination of Randomly Packed Spheres. *Nature*, 385:910–911, 1990.

[202] S. Torquato and F. H. Stillinger. Local Density Fluctuations, Hyperuniform Systems, and Order Metrics. *Phys. Rev. E*, 68:041113, 2003. ibid. 069901.

[203] A. Bezdek and W. Kuperberg. *Applied geometry and discrete mathematics: DIMACS Ser. Discrete Math. Theoret. Comput. Sci. 4 (P. Gritzmann and B. Sturmfels, editors)*, chapter Packing Euclidean space with congruent cylinders and with congruent ellipsoids, pages 71–80. Amer. Math. Soc., Providence, RI, 1991.

[204] J. Pach and P. K. Agarwal. *Combinatorial Geometry*. Wiley-Interscience, 1st edition, 1995.

[205] J. M. Wills. An ellipsoid packing in e^3 of unexpected high density. *Mathematika*, 38:318–320, 1991.

[206] Y. Song and E. A. Mason. Equation of State for a Fluid of Hard Convex Bodies in Any Number of Dimensions. *Phys. Rev. A*, 41(6):3121–3124, 1990.

[207] J. R. Morris and X. Song. The Melting Lines of Model Systems Calculated from Coexistence Simulations. *J. Chem. Phys.*, 116(21):9352–9358, 2002.

[208] T. A. Schilling and D. Frenkel. Nucleation in suspensions of anisotropic colloids. *Comp. Phys. Comm.*, 169(1-3):117–121, 2005.

[209] S. F. Edwards and D. V. Grinev. Statistical Mechanics of Stress Transmission in Disordered Granular Arrays. *Phys. Rev. Lett.*, 82:5397–5400, 1999.

[210] L. E. Silbert, G. S. Grest, and J. W. Landry. Statistics of the contact network in frictional and frictionless granular packings. *Phys. Rev. E*, 66:061303, 2002.

[211] R. J. Speedy. Quench Rate Independence of the Hard Sphere Glass Transition. *Mol. Phys.*, 83(3):591–597, 1994.

[212] L. E. Silbert, A. J. Liu, , and S. R. Nagel. Structural signatures of the unjamming transition at zero temperature. *Phys. Rev. E*, 73:041304, 2006.

[213] A. J. Liu, S. R. Nagel, D. A. Schecter, S. N. Coppersmith, S. Majumdar, O. Narayan, and T. A. Witten. Force fluctuations in bead packs. *Science*, 260:513, 1995.

[214] P. Metzger. Granular contact force density of states and entropy in a modified Edwards ensemble. *Phys. Rev. E*, 70:051303, 2004.

[215] A. S. Clarke and H. Jonsson. Structural Changes Accompanying Densification of Random Hard-Sphere Packings. *Phys. Rev. E*, 47(6):3975–3984, 1993.

[216] C. Rascon, L. Mederos, and G. Navascues. Theoretical Approach to the Correlations of a Classical Crystal. *Phys. Rev. E*, 54:1261–1264, 1996.

[217] M. D. Rintoul and S. Torquato. Metastability and Crystallization in Hard-Sphere Systems. *Phys. Rev. Lett.*, 77:4198, 1996.

[218] K. Kendall, C. Stainton amd F. van Swol, and L. V. Woodcock. The Computer Study of Transport Processes Under Extreme Conditions. *Int. J. Thermophysics*, 23(1):175–186, 2002.

[219] J. H. Snoeijer, M. van Hecke, E. Somfai, and W. van Saarloos. Force and Weight Distributions in Granular Media: Effects of Contact Geometry. *Phys. Rev. E*, 67:030302(R), 2003.

[220] A. Gabrielli, M. Joyce, and F. S. Labini. Glass-like universe: Real-space correlation properties of standard cosmological models. *Phys. Rev. D*, 65:083523, 2002.

[221] J. J. Salacuse, A. R. Denton, and P. A. Egelstaff. Finite-Size Effects in Molecular Dynamics Simulations: Static Structure Factor and Compressibility. I. Theoretical Method II. Application to a Model Krypton Fluid. *Phys. Rev. E*, 53:2382–2401, 1996.

[222] D. Potts, G. Steidl, and M. Tasche. *Modern Sampling Theory: Mathematics and Applications*, chapter 12. Fast Fourier transforms for nonequispaced data: A tutorial, pages 249–274. Applied and Numerical Harmonic Analysis Series. Birkhauser Boston, 2001.

[223] L. Greengard and J.-Y. Lee. Accelerating the Nonuniform Fast Fourier Transform. *SIAM Review*, 46:443, 2004.

[224] A. Baumketner and Y. Hiwatari. Finite-size dependence of the bridge function extracted from molecular dynamics simulations. *Phys. Rev. E*, 63:061201, 2001.

[225] M. Dijkstra and R. Evans. A Simulation Study of the Decay of the Pair Correlation Function in Simple Liquids. *J. Chem. Phys.*, 112:1449, 2000.

[226] F. L. Roman, J. A. White, and S. Velasco. Fluctuations in an Equilibrium Hard-Disk Fluid: Explicit Size Effects. *J. Chem. Phys.*, 107:4635–4641, 1997.

[227] F. L. Roman, J. A. White, A. Gonzalez, and S. Velasco. Fluctuations in a Small Hard-Disk System: Implicit Finite Size Effects. *J. Chem. Phys.*, 110:9821–9824, 1999.

[228] P. Perry and G. J. Throop. Decay of Pair Correlations in Hard Sphere Fluids. *J. Chem. Phys.*, 57:1827–1829, 1972.

[229] S. Torquato and F. H. Stillinger. Controlling the Short-Range Order and Packing Densities of Many-Particle Systems. *J. Phys. Chem. B*, 106:8354–8359, 2002.

[230] P. M. Chaikin and T. C. Lubensky. *Principles of condensed matter physics*. Cambridge University Press, Cambridge, United Kingdom, 1995.

[231] M. Luban and A. Baram. Third and fourth virial coefficients of hard hyperspheres of arbitrary dimensionality. *J. Chem. Phys.*, 76:3233, 1982.

[232] N. Clisby and B. McCoy. Analytic Calculation of B4 for Hard Spheres in Even Dimensions. *J. Stat. Phys.*, 114:1343–1361, 2004.

[233] N. Clisby and B. McCoy. Ninth and Tenth Order Virial Coefficients for Hard Spheres in d-Dimensions. *J. Stat. Phys.*, 122:15–57, 2006.

[234] M. Bishop, A. Masters, and A. Y. Vlasov. Higher virial coefficients of four and five dimensional hard hyperspheres. *J. Chem. Phys.*, 121:6884, 2004.

[235] M. Luban and J. P. J. Michels. Equation of state of hard d-dimensional hyperspheres. *Phys. Rev. A*, 41:6796, 1990.

[236] H. L. Frisch and J. K. Percus. High dimensionality as an organizing device for classical fluids. *Phys. Rev. E*, 60:2942, 1999.

[237] G. Parisi and F. Slanina. Toy model for the mean-field theory of hard-sphere liquids. *Phys. Rev. E*, 62:6554, 2000.

[238] R. Finken, M. Schmidt, and H. Lowen. Freezing transition of hard hyperspheres. *Phys. Rev. E*, 65:016108, 2001.

[239] A. P. Philipse. Caging effects in amorphous hard-sphere solids. *Colloids and Surfaces A*, 213:167–173, 2003.

[240] S. Torquato and F. H. Stillinger. New conjectural lower bounds on the optimal density of sphere packings. *Experimental Math.*, in press.

[241] S. Torquato and F. H. Stillinger. Exactly Solvable Disordered Hard-Sphere Packing Model in Arbitrary-Dimensional Euclidean Spaces. *Phys. Rev. E*, 73:031106, 2006.

[242] M. Bishop, P. Whitlock, and D. Klein. The structure of hyperspherical fluids in various dimensions. *J. Chem. Phys.*, 122:074508, 2005.

[243] J. P. J. Michels and N. J. Trappaniers. Dynamical computer simulations on hard hyperspheres in four- and five-dimensional space. *Phys. Lett. A*, 104:425, 1984.

[244] M. Ross and B. J. Alder. Melting Curve at High Pressure. *Phys. Rev. Lett.*, 16:1077, 1966.

[245] W. B. Streett, H. J. Raveche, and R. D. Mountain. Monte Carlo studies of the fluid-solid phase transition in the Lennard-Jones system. *J. Chem. Phys.*, 61:1960, 1974.

[246] S. Torquato, O. Ouche, and F. H. Stillinger. 2006.

[247] K. W. Wojciechowski and D. Frenkel. Tetratic phase in the planar hard square system? *Comp. Methods in Science and Tech.*, 10(2):235–255, 2004.

[248] K. W. Wojciechowski, D. Frenkel, and A.C. Brańka. Nonperiodic Solid Phase in a Two-Dimensional Hard-Dimer System. *Phys. Rev. Lett.*, 66(24):3168–3171, 1991.

[249] K. W. Wojciechowski. Monte Carlo simulations of highly anisotropic two-dimensional hard dumbbell-shaped molecules: Nonperiodic phase between fluid and dense solid. *Phys. Rev. B*, 46:26–39, 1992.

[250] J.-L. Barrat and J.-P. Hansen. *Basic Concepts for Simple and Complex Liquids*. Cambridge University Press, 2003.

[251] P. J. Camp and M. P. Allen. Phase diagram of the hard biaxial ellipsoid fluid. *J. Chem. Phys.*, 106:6681–6688, 1997.

[252] B. R. Acharya, A. Primak, and S. Kumar. Biaxial Nematic Phase in Bent-core Thermotropic Mesogens. *Phys. Rev. Lett.*, 92:145506, 2004.

[253] R. Zwanzig. First-Order Phase Transition in a Gas of Long Thin Rods. *J. Chem. Phys.*, 39(7):1714, 1963.

[254] K. Shundyak and R. van Roij. Isotropic-nematic transition in hard-rod fluids: Relation between continuous and restricted-orientation models. *Phys. Rev. E*, 69:041703, 2004.

[255] Y. Martínez-Ratón, E. Velasco, and L. Mederos. Effect of particle geometry on phase transitions in two-dimensional liquid crystals. *J. of Chem. Phys.*, 122(064903), 2005.

[256] H. Schlacken, H.-J. Mogel, and P. Schiller. Orientational transitions of two-dimensional hard rod fluids. *Mol. Phys.*, 93(5):777–787, 1998.

[257] P. W. Kasteleyn. Statistics of dimers on a lattice 1. Number of dimer arrangements on a quadratic lattice. *Physica*, 27(12):1209–1211, 1961.

[258] H. N. V. Temperley and M. E. Fisher. Dimer Problem in Statistical Mechanics: An Exact Result. *Phil. Mag.*, 6(68):1061–1063, 1961.

[259] T. A. Weber and F. H. Stillinger. Melting of square crystals in two dimensions. *Phys. Rev. E*, 48:4351–4358, 1993.

[260] T. Schilling, S. Pronk, B. Mulder, and D. Frenkel. Monte Carlo study of hard pentagons. *Phys. Rev. E*, 71:036138, 2005.

[261] T. Boublík. Two-dimensional convex particle liquid. *Mol. Phys.*, 29(2):421–428, 1975.

[262] R. W. Kenyon, J. G. Propp, and D. B. Wilson. Trees and Matchings. *Electronic J. of Combinatorics*, 7(1):R25, 2000.

[263] H. Cohn, R. Kenyon, and J. Propp. A variational principle for domino tilings. *J. Am. Math. Soc.*, 14:297–346, 2001.

[264] R. J. Low. Measuring order and biaxiality. *Eur. J. Phys.*, 23(2):111–117, 2002.

[265] D. Frenkel and R. Eppenga. Evidence for algebraic orientational order in a two-dimensional hard-core nematic. *Phys. Rev. A*, 31(3):1776–1787, 1985.

[266] M. A. Bates and D. Frenkel. Phase behavior of two-dimensional hard rod fluids. *J. of Chem. Phys.*, 112(22):10034–10041, 2000.

[267] H. Weber, D. Marx, and K. Binder. Melting transition in two dimensions: A finite-size scaling analysis of bond-orientational order in hard disks. *Phys. Rev. B*, 51:14636–14651, 1995.

[268] C. Richard, M. Höffe, J. Hermisson, and M. Baake. Random tilings: concepts and examples. *J. Phys. A: Math. Gen.*, 31:6385–6408, 1998.

[269] K. W. Wojciechowski. Solid phases of two-dimensional hard dumb-bells in the free volume approximation: Crystal-aperiodic solid phase transition. *Phys. Lett. A*, 122(6):377–380, 1987.

[270] R. Kenyon. Local statistics of lattice dimers. *Annales de Inst. H. Poincaré, Probabilités et Statistiques*, 33:591–618, 1997.

[271] J. A. Cuesta and D. Frenkel. Monte Carlo simulation of two-dimensional hard ellipses. *Phys. Rev. A*, 42:2126–2136, 1990.

[272] G. Adam and J. H. Gibbs. On the Temperature Dependence of Cooperative Relaxation Properties in Glass-Forming Liquids. *J. Chem. Phys.*, 43:139–146, 1965.

[273] F. Sciortino. Potential energy landscape description of supercooled liquids and glasses. *J. Stat. Mech.*, 2005:P05015, 2005.

[274] R. J. Speedy. The Hard Sphere Glass Transition. *Mol. Phys.*, 95(2):169–178, 1998.

[275] B. Coluzzi, G. Parisi, and P. Verrocchio. Thermodynamical Liquid-Glass Transition in a Lennard-Jones Binary Mixture. *Phys. Rev. Lett.*, 84:306–309, 2000.

[276] G. Parisi and F. Zamponi. The Ideal Glass Transition of Hard Spheres. *J. Chem. Phys.*, 123:144501, 2005.

[277] S. Luding. Global equation of state of two-dimensional hard sphere systems. *Phys. Rev. E*, 63:042201, 2001.

[278] Tom Kennedy. A densest compact planar packing with two sizes of discs, 2004.

[279] X. C. Zeng, D. W. Oxtoby, and Y. Rosenfeld. Freezing of a two-dimensional binary hard-disk liquid:

A density-functional approach. *Phys. Rev. A*, 43:2064–2067, 1991.

[280] R. J. Wheatley. Phase diagrams hard disc mixtures. *Mol. Phys.*, 93:965–969, 1998.

[281] S. R. Williams, I. K. Snook, and W. van Megen. Molecular dynamics study of the stability of the hard sphere glass. *Phys. Rev. E*, 64:021506, 2001.

[282] S. Luding and A. Santos. Molecular dynamics and theory for the contact values of the radial distribution functions of hard-disk fluid mixtures. *J. Chem. Phys.*, 121:8458–8465, 2004.

[283] Werner Krauth. Absence of Thermodynamic Phase Transition in a Model Glass Former. *Nature*, 405:550, 2000.

[284] P. G. Debenedetti, M. S. Shell, and F. H. Stillinger. Model Energy Landscapes. *J. Phys. Chem. B*, 107:14434–14442, 2003.

[285] Y. Brumer and D. R. Reichman. Numerical Investigation of the Entropy Crisis in Model Glass Formers. *J. Phys. Chem. B*, 108:6832 –6837, 2004.

[286] T. S. Grigera and G. Parisi. Fast Monte Carlo algorithm for supercooled soft spheres. *Phys. Rev. E*, 63:045102R, 2001.

[287] K. K. Bhattacharya and J. P. Sethna. Multicanonical methods, molecular dynamics, and Monte Carlo methods: Comparison for Lennard-Jones glasses. *Phys. Rev. E*, 57:2553–2562, 1998.

[288] T. S. Grigera, A. Cavagna, I. Giardina, and G. Parisi. Geometric Approach to the Dynamic Glass Transition. *Phys. Rev. Lett.*, 88:055502, 2002.

[289] M. S. Shell, P. G. Debenedetti, and A. Z. Panagiotopoulos. Saddles in the Energy Landscape: Extensivity and Thermodynamic Formalism. *Phys. Rev. Lett.*, 92:035506, 2004.

[290] P. G. Debenedetti and F. H. Stillinger. Supercooled Liquids and the Glass Transition. *Nature*, 410:259–267, 2001.

[291] R. J. Speedy. Configurational entropy and diffusion in a hard disc fluid. *J. Chem. Phys*, 114(20):9069–9074, 2001.

[292] S. Marčelja. Entropy of phase-separated structures. *Physica A*, 231:168–177, 1996.

[293] B. Kozintsev and B. Kedem. Generation of "Similar" Images From a Given Discrete Image. *J. Comp. Graph. Stat*, 9:286–302, 2000. Software for `gaussian` package is at `http://www.math.umd.edu/~bnk/CLIP/clip.gauss.htm`.

[294] B. Matern. *Spatial Variation*, volume 36 of *Lecture Notes in Statistics*. Springer Verlag, Berlin, 1960.

[295] P. Attard. *In Statistical Physics on the Eve of the Twenty-First Century*, chapter Markov Superposition Expansion for the Entropy and Correlation Functions in Two and Three Dimensions. World Scientific, Singapore, 1999.