



SAPIENZA
UNIVERSITÀ DI ROMA

Numerical study of the microscopic structure of jammed systems: from inferring their dynamics to finite size scaling

Scuola di dottorato Vito Volterra

Dottorato di Ricerca in Fisica – XXXIII Ciclo

Candidate

Rafael Alberto Díaz Hernández Rojas

ID number 1798865

Thesis Advisors

Prof. Federico Ricci Tersenghi

Prof. Giorgio Parisi

February 2021

Thesis defended on 5 July 2021
in front of a Board of Examiners composed by:
Andrea Pelissetto (chairman)
Andrea Gambassi
Giuseppe Gonnella

Numerical study of the microscopic structure of jammed systems: from inferring their dynamics to finite size scaling

Ph.D. thesis. Sapienza – University of Rome

© year 2021 Rafael Alberto Díaz Hernández Rojas. All rights reserved

This thesis has been typeset by L^AT_EX and the Sapthesis class.

Author's email: rafael.diazhernandezrojas@uniroma1.it

Abstract

In this thesis I present most of the results obtained during my PhD, where I worked on different subjects regarding jamming in systems of frictionless spheres. In particular, I focused on microscopic properties of jammed packings, such as the distribution of contact forces and interparticle gaps, as well as the single particle dynamics that occur near the jamming point. Several of these results have already been presented in Refs. [1, 2], but here I include a more detailed analysis of some of them. Besides, all the results of the second chapter are new, even if related to such works.

The thesis is structured as follows. Chapter 1 provides a quick introduction to the phenomenology of glassy systems, with an entire section dedicated to hard spheres (HS) as a minimalistic model of a glass former. I also include a brief survey of a recently developed mean-field theory, capable of providing an exact description of liquids, glasses and jammed systems in infinite dimensions. This theory serves as a framework for explaining many of the features of real glasses as “blurred” or imperfect analogies of the sharp transitions predicted in $d = \infty$. Moreover, its predictions about the distributions of contact forces and gaps seem to remain valid in low dimensions, *i.e.* $d = 2, 3$. Given that here I only consider tridimensional system, its relevance is obvious. But not only, the physical picture based on the meta-basin structure of the free energy landscape will be a guiding principle for many of the topics in later chapters. The final section of this first chapter is devoted to a general discussion about the jamming transition. Naturally, special attention is given to the microscopic features of jammed packings. Thus, I explain in detail some of the most important results derived in previous works. In particular, I give a careful description of the network of contact forces, showing that it is entirely determined by the particles’ position, and that it contains important information about the packings stability. Similarly, I reproduce the proof that the exponents of the distribution of forces and gaps are connected through stability inequalities. Because such bounds are saturated, critical jammed packings are only marginally stable. This last feature is rationalized in terms of the constant density of states for vanishingly small frequencies. I hope readers will benefit from having all this material gathered in a single source.

Chapter 2 contains a detailed description of the iterative Linear Programming (iLP) algorithm we developed to generate jammed packings. This method has been previously introduced in Refs. [1, 3], but here I present a complete explanation and proofs of several of its features. Moreover, I carry out a detailed characterization of it, although restricted to $3d$ systems. Additionally, in this chapter I review the Lubachevsky–Stillinger compression protocol, based on event-driven molecular dynamics simulations. This method is able to efficiently compress configurations of HS up to very high pressures, and at the same time allows to probe their glassy phase. Complementing our iLP crunching algorithm with this procedure results in a reliable and reasonably fast method to generate jammed packings of HS. Even more, I show that we can apply the same technique to produce jammed configurations of the Mari–Kurchan (MK) model.

The dynamics of particles near the jamming point is explored in Chapter 3.

Given that few works have addressed this issue before us, the first step was to characterise the statistical properties of the trajectories of individual particles in this dynamical regime. Then, the idea was to investigate if the information of the network of contacts can be used to make statistical inferences of the particles' motion close to their jamming point. We found that by considering only the contact *vectors* (*i.e.* ignoring the magnitude of the forces) we were able to construct a couple of structural variables that correlate well with the dynamical features of the particles. More precisely, the vectorial sum of contact vectors is a good descriptor of the preferential directions in single-particle trajectories; while the sum of dot products between such vectors can be used as a predictor of a particle's mobility. The correlations thus obtained are significantly high, although rather short-lived. Importantly, our method proved to be superior than a normal modes approach, which fails to capture these dynamical features at the level of individual particles. I should also mention that most of the results discussed in this chapter have appeared in Ref. [1].

Finally, Chapter 4 deals with a thorough analysis of the critical distributions of contact forces and interparticle gaps in configurations of HS, soft spheres, as well as in the MK model in $d = 3$. The main results can also be found in Ref. [2], where polydisperse disks and near crystals with an FCC structure were also considered. The purpose is to verify the expected power-laws using an analysis based on the finite size effects of such critical distributions. As mentioned above, current numerical estimations of the power-law exponents suggest that jammed packings saturate the stability criteria, so carefully validating their value is clearly important. Moreover, theoretical predictions of these exponents imply that jamming defines an universality class, to which a broad range of constraint-satisfaction problems belong. However, the calculations involved, being based on mean-field theory, are only exact in the $d \rightarrow \infty$ limit. And yet, available numerical results indicate that such universal criticality holds in low dimensions as well. Proving that the same scalings occur in finite dimensions is thus of major theoretical interest, as it would represent the most precise prediction of the so called replica method in realistic materials models. Our approach uses the scaling collapse of the distributions obtained from systems of different sizes and it amply confirms the predicted values in both models. However, it also reveals a striking difference in the size effects on the distribution of forces and gaps, namely, that such effects are negligible in the former but very pronounced in the latter. We rationalize this feature in terms of two correlation lengths that determine how fast the thermodynamic limit behaviour is reached.

I would be very glad if new graduate students entering the intriguing (and complex) world of jamming should find this thesis useful. With this in mind, I took advantage of the fact that after Chap. 1 all chapters are independent from each other, and thus I tried to make them as self-contained as possible. (Or, in any case, I refer to the relevant sections that provide the necessary context for the topics discussed.) For the same reasons, I decided to give independent conclusions in each chapter, rather than providing general ones at the end. However, I hope that it is sufficiently clear that the unifying thread of all the results presented here are the peculiar properties in the microscopic structure of jammed systems.

Acknowledgments

I've seen somewhere the advice that acknowledgements should be kept concise, or even be avoided all together, lest a work be considered unprofessional. I would hope that the last 15 months suffice to change the mind of anyone who still holds that view. At least for me, taking any opportunity at hand to express my deepest gratitude for the immense fortune I've had is now more important than ever. Or, in any case, let the following lines be a firm testimony that I rather be considered unprofessional than ungrateful.

First of all, I sincerely want to thank my supervisors, Prof. Federico Ricci Tersenghi and Prof. Giorgio Parisi, for guiding me through my PhD. It's been an honour to study with scientists of their stature. Working with them has been a most rewarding experience for several reasons, but mainly due to their genuine interest in discovering and understanding every possible aspect of a problem, and also because they transmit their (incredibly vast) knowledge in a very unassuming way. All of this combined with the friendly and cooperative environment of the Chimera group. And even if doing a PhD usually entails becoming overspecialized in a single topic, I am ready to admit at any given time that my major lesson has been that I know nothing but " $1/N^3$ contributions" of all the things they have to teach.

I also want to thank the Physics Department at Sapienza and the PhD coordination for the facilities and travel funds, as well as to the ERC for providing me with the grant that allowed me to pursue my PhD. Additionally, I want to express my gratitude to Profs. Walter Kob and Eric Corwin for reviewing this work as external referees. Their positive appraisal of my thesis has been truly encouraging.

The last chapter of this thesis presents the results obtained in a joint work with an outstanding group of colleagues: Patrick Charbonneau, Eric Corwin, Cameron Dennis and Harukuni Ikeda. I want to warmly thank them all for sharing their data with us and for all the things I learnt during this great collaboration.

I also benefited amply from many interactions with other PhD students: Angelo Cavaliere, Claudia Artiaco, Paolo Baldan; and other researchers: Patrick Charbonneau, Olivier Dauchot, Andrea Liu, Georgios Tsekenis, Francesco Zamponi. From informal talks to concrete suggestions and comments, all of them have undoubtedly helped me and influenced my understanding of the topics of this thesis. Thank you all for your good will in discussing physics with me!

The research I had the opportunity to carry out was in itself a terrific experience, but another very fulfilling aspect of my stay in Rome was the amazing persons I had pleasure to meet. In particular, the other PhD students and expats: Gui, Ilaria, Javi and Masry. But also the people from the Stanza 117/118 NEF: Angelo, Fabrizio, Giovanni, José, Silvia. Grazie mille, ragazzi! Tutti voi avete reso Roma una città indimenticabile per me.

Obviamente no puedo dejar de agradecer profundamente a toda la gente que siempre ha estado "ahí" para mí, cuando los he necesitado y cuando no. Primero, a mis padres, Antonieta y Rafael, mi abuela, mis tíos Roberto y Claudia, mis primos

Emiliano y Constanza, además de mi tía Lina. Ustedes son la evidencia de que hay mucho más cariño que océano. Por supuesto otra mención especial se merecen mis amigos: Andrea, Bofo, Motita, Karina, Bego, Tony, Eve, Ángela y Humberto. Y claro, los mexas que uno (re)encuentra de este lado del mundo y que hacen que el Atlántico se sienta menos ancho: Tania y Santi; Álvaro y Judith; David; Pato y Gabi.

Y a Viri, que ha estado aquí, en los momentos más difíciles y en los más memorables. Además de tu paciencia para escucharme (usualmente a deshoras) cada vez que pienso en voz alta porque algo no cuadra en mis simulaciones, análisis de datos, etc. Pensar que “todo fue casualidad; la misma beca, la misma universidad” me hace suponer que “debo caerle (muy) bien a alguien allá arriba”. Estos años a tu lado, viviendo con Megara en nuestra pseudo embajada mexicana y armados de mole, un molcajete y mayonesa McCormick (rigurosamente), me han transformado en un tipo capaz de cocinar algo más que quesadillas, en alguien más enamorado que el pinguino inamorado, y sobre todo en una mejor persona. Y yo que pensé que el amor traía chueco el tino...

A todos ustedes, de verdad, infinitas gracias. De nada servirían los logros que he tenido si no pudiera compartirlos con ustedes.

Contents

Contents

vii

1	Introduction: Why do we care about jammed systems?	1
1.1	Phenomenology of glassy systems: the departure point towards jamming	3
1.1.1	Local structure and its connection with dynamics	13
1.2	Hard sphere systems: From liquid to jamming with a glassy interlude	18
1.2.1	Basic Statistical Mechanics of liquids	20
1.2.2	Equation of state of HS liquids	24
1.2.3	Glass phase and Gardner transition in HS systems	28
1.3	Mean-field theory: The $d \rightarrow \infty$ limit	31
1.3.1	Liquids in high dimensions	32
1.3.2	Dynamics and the glass transition	35
1.3.3	Thermodynamics of mean-field glasses	37
1.3.4	Gardner Transition and fractal free energy landscape	39
1.4	The jamming transition	42
1.4.1	All the roads lead to jamming	43
1.4.2	Criticality approaching the jamming transition	46
1.4.3	Network of contacts and isostaticity	49
1.4.4	Criticality of the microscopic structure <i>at</i> the jamming point	55
1.4.5	Marginal stability	56
1.4.5.1	Density of states at jamming	60
2	Linear Programming algorithm to generate jammed configurations	63
2.1	Jamming as a Linear Programming problem	64
2.1.1	Some results of Convex Optimization	68
2.1.2	Stability and network of contacts from iLP	71
2.1.3	Details of the algorithm and some examples	73
2.2	Approaching jamming using molecular dynamics	76
2.2.1	Basics of event-driven MD	78
2.2.2	Compression protocol	82
2.3	Characterization of MD+iLP algorithm	86
2.3.1	Influence of the initial condition's parameters	87
2.3.2	The MD+iLP route to jamming	90
2.3.3	Complexity of iLP: Scaling with size	98
2.4	Conclusions and future work	105

3	Inferring the particle-wise dynamics near the jamming point	107
3.1	Statics: Physical quantities at the jamming point	109
3.1.1	Critical properties of the jammed configurations	110
3.1.2	Structural variables of interest	112
3.2	Dynamical variables of interest	116
3.3	Results with MD simulations	118
3.3.1	Further details of the MD simulations	118
3.3.2	Particle trajectories and their statistics	119
3.3.3	Correlation with structure	123
3.4	Results with MC simulations	128
3.5	Contact vectors <i>vs</i> contact forces	132
3.6	Further statistical properties	134
3.6.1	Predicting the mean displacement and mobility	134
3.6.2	“Universal” decorrelation rate	136
3.7	Comparison with normal modes and other works	139
3.8	Conclusions	143
3.8.1	Future work	147
	Appendices	148
3.A	Reduced analysis on an independent configuration	148
4	Finite size scaling of microscopic structural variables	151
4.1	Introduction	151
4.2	Methods for generating jammed packings	154
4.3	Methods for studying finite size effects	156
4.4	Finite size effects in spheres systems	159
4.4.1	Some structural properties	159
4.4.2	Finite size scaling of the critical distributions	161
4.4.2.1	Distribution of extended forces	161
4.4.2.2	Distribution of gaps	164
4.4.2.3	Distribution of localized forces	167
4.5	Finite size effects in the MK model	168
4.5.1	Distribution of contact vectors	168
4.5.2	Finite size scaling of the critical distributions	171
4.5.2.1	Distribution of extended forces	171
4.5.2.2	Distribution of gaps	172
4.5.2.3	Distribution of localized forces	176
4.6	Conclusions	178
	Appendices	180
4.A	Equivalence of cumulative distributions	180
	Bibliography	183
	List of Figures	199

Chapter 1

Introduction: Why do we care about jammed systems?

Broadly speaking, a system is said to be jammed when all its degrees of freedom are blocked (see Fig. 1.1 and Figs. 2.3, 3.1 below). In particular, in disordered systems such as glasses and amorphous solids, this is usually caused by geometric frustration. But in contrast to their lack of dynamics, research on jammed systems has been far from still. In this introductory chapter, I will give a general, and definitely non-exhaustive, overview of the myriad of fascinating phenomena that occur when a system approaches and *reaches* its jamming point. Yet, an early warning is in place: the results presented in this thesis were obtained by working on configurations of frictionless spherical particles and, therefore, most of what will be discussed is restricted to this kind of systems. Despite such simplification, I will argue that these packings present features that are far from trivial. I will be mainly concerned with the *microscopic* structural properties of jammed configurations, specifically (i) the small gaps formed between particles in near contact, and (ii) the network of contact forces. As I will explain below, it has been known for some time now that, among other things, their distributions exhibit non-trivial scalings at the jamming point. The main goal of this thesis is to investigate what other information they contain. It should be clear that such investigation is far from over, but I hope that the next chapters will contribute to expand our knowledge in that direction; or at least to better demarcate our (my) ignorance.

Let me now sketch the path I will follow. In this first chapter I will discuss many of the ideas needed to better understand the main results of this work. Thus, Sec. 1.1 is a very quick and superficial description of the phenomenology of super-cooled liquids and glasses. This might seem unnecessary given that jamming is not exclusive of glasses. However, it will prove useful for appreciating that the glass and the jamming transitions are *different* transitions. We now understand their differences clearly enough, but few years ago the situation was much blurrier. Furthermore, in the algorithm described in Chp. 2 we mimic the process of taking a liquid deep in its glass phase, and only then compress it further until the jamming point is reached. Having a clear idea of the physics that takes place in these regimes is thus important. Additionally, the discussion of Sec. 1.1.1 contains many of the necessary background to the results of Chp. 3. In Sec. 1.2 I will show that many of

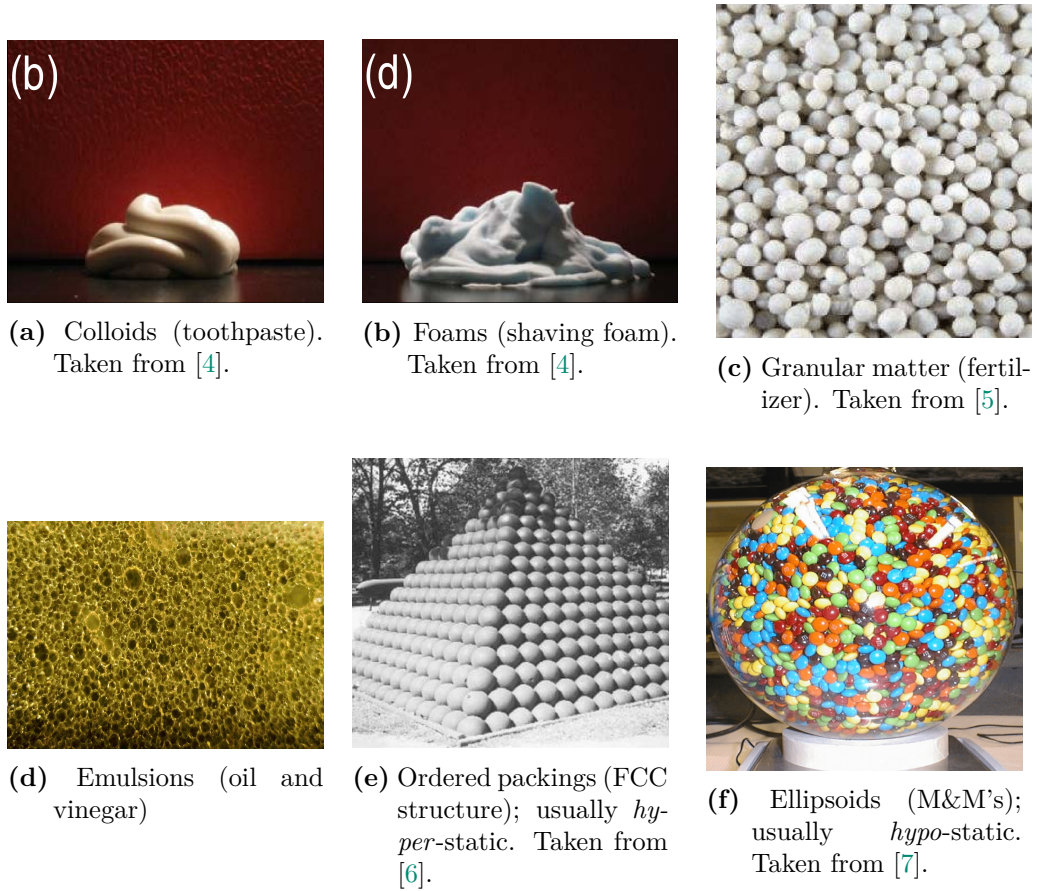


Figure 1.1. Some examples of systems that exhibit a jammed regime. However, the results in this thesis do not concern the last two, mainly due to their ordering and shape, but also because I will only focus on *isostatic* systems; see Sec. 1.4.

the peculiar features of glass formers can be reproduced using a very simple model: a hard-sphere (HS) fluid. The resemblance in their behaviour then justifies the usage of this minimal model for studying glassy and jammed systems. Hence, in this second section, after a very basic introduction to the statistical mechanics of liquids, I analyse the equations of state for HS in liquid and glass phase, as well as other important phenomena, like the Gardner transition. Moreover, a recently developed mean-field theory –exact, in the infinite dimensional limit– uses both HS and soft spheres (SS) as archetypical models. The theory is powerful enough to consistently describe the behaviour of glass-forming systems from the point they fall out of equilibrium all the way down to jamming. A brief outline of this theory and some of its many predictions are given in of Sec. 1.3. Finally, in Sec. 1.4 I discuss in detail the jamming transition. From the approach I follow, jamming is an (out of equilibrium) phase transition that marks the end of the glass in HS and SS systems. Surprisingly, this phase transition is shared by several other systems (Sec. 1.4.1) such as grains, foams, and colloids as depicted in Fig. 1.1. All of them present common properties and can thus be studied using similar tools (Secs. 1.4.2 and 1.4.3). The last parts of the chapter (Secs. 1.4.4 and 1.4.5) are devoted to analyse the properties of the

contact forces and interparticle gaps. These microscopic structural variables have a central role for the results of this work.

1.1 Phenomenology of glassy systems: the departure point towards jamming

Glasses are very common in our daily lives and, yet, incredibly hard to explain satisfactorily from a physical point of view. And every year that passes by without a complete theory of the *glass transition* (defined soon below), helps to proof the now famous quote by Philip W. Anderson^[8], “The deepest and most interesting unsolved problem in solid state theory is probably the theory of the nature of glass and the glass transition.” As with any other interesting problem, the lack of a complete theory of the glass phase is not caused by a shortage of efforts. Indeed, the current literature on glassy systems is immense, as a (very personally biased) sample can show: [5, 9–22]. The contents of this section strongly follows Refs. [5, 9, 10], which provide very amenable reviews of the topic of glasses and supercooled liquids.

In very simple terms, what makes glasses so hard to understand is that they behave, in many aspects, just as (non-crystallized) solids, but structurally they are much more similar to liquids. In fact, it is “relatively simple” to produce a glass by cooling a liquid fast enough in order to avoid crystallization. (See however [9, Sec. 2] for a detailed account of the obstacles that can occur in practice when cooling a liquid.) As temperature decreases below the liquid’s melting point (T_m), the viscosity (η) increases drastically and, at some point, the liquid simply stops flowing and behaves, mechanically, as a solid. For instance, it develops a finite shear modulus that does not decay on time. Such point identifies the *glass transition* of the liquid, and the temperature at which this happens is the so called *glass transition temperature*, T_g . Pragmatically, a rule of thumb to define T_g is the temperature such that^[9,11]

$$\eta(T_g) \sim 10^{13} \text{ Poise} = 10^{12} \text{ Pa s} . \quad (1.1)$$

In comparison, the viscosity of water at room temperature is minute: $\eta_{\text{water}} \approx 10^{-3} \text{ Pa s}$.

Yet, analysing the microscopic structure of the systems slightly above and slightly below T_g reveals that little, if anything, has changed: glasses are, essentially, solidified non-crystalline liquids. This process is depicted in the volume-temperature plane of Fig. 1.2. We can see that at for $T > T_m$ the liquid is the only stable phase, but for temperatures lower than T_m the crystal becomes the equilibrium branch. Nevertheless, by applying a cooling rate fast enough to avoid crystal nucleation, but small enough to let the system relax, the liquid phase can be sustained in the so called supercooled phase, identified by the shaded region. In other words, we have to cool the system in such a way that it remains in *metastable* equilibrium. As shown in the insets, there is a clear difference in the structure (*i.e.* the arrangement of particles) between the crystalline and liquid phase. However, when reaching T_g there does not seem to be any structural change in the configuration, but the physical properties become very different. This figure also shows a very peculiar feature of the glass transition: the exact location of T_g actually depends on the protocol




Figure 1.2. Main: Glass – supercooled liquid phase diagram. For temperatures above (resp. below) the melting temperature (T_m), the liquid (resp. crystal) is the equilibrium state. When the temperature is not too small compared to T_m , crystallization can be avoided by fast cooling, allowing the liquid to remain in a metastable state: the supercooled liquid. In such metastable phase (shaded region) the relaxation time grows drastically. If T is further decreased, the supercooled liquid is bound to become a glass, characterized by an “immense” relaxation time and viscosity, at the glass transition temperature (T_g). Interestingly, different cooling protocols locate T_g at different values (curves labelled as “1” and “2”); see text for a discussion. Insets: Snapshots of the configurations obtained through molecular dynamics simulations of a crystal (right), supercooled liquid (centre), and glass (left). Notably, there is little structural difference between the supercooled liquid and glass configurations, but dynamically they are completely different as discussed throughout this section. Figure taken from Ref. [23].

used to reach the glass state. That is, if a given cooling rate results in the glass state identified by the branch “1”, a slower protocol would result in the branch identified by “2”. Such behaviour immediately begs the question: what kind of transition does T_g signals, if its value depends on the protocol employed? Even more, the definition of T_g in the expression (1.1) is arbitrary. Could we have chosen a different value for the threshold of η instead of 10^{13} Poise?

From these considerations it should be clear that the glass transition is not a usual thermodynamic transition, while it is also true that the definition of T_g entails

a certain degree of arbitrariness. To better understand the physics of supercooled liquids and glasses, it is useful to recall the relation between the viscosity and relaxation time τ_α . For instance, from the Maxwell model of elasticity, we have that^[5,9] $\eta = G_\infty \tau_\alpha$, where G_∞ is the instantaneous shear modulus. τ_α determines the typical time scale on which density fluctuations in a system relax. Hence, it is natural that as the system becomes more viscous, fluctuations take longer to relax in the system. In several systems it has been found empirically that as T_g is approached, the increase in relaxation time follows an exponential behaviour, captured by the so call Arrhenius relation:

$$\tau_\alpha = \tau_0 \exp\left(\frac{E}{k_B T}\right), \quad (1.2)$$

where E is an activation energy, and k_B is the Boltzmann constant, which from now on I will assume equal to 1. Even more, there are various other materials whose viscosity increases at an even faster rate. Their sluggish behaviour is well captured by the Vogel–Fulcher–Tamman (VFT) law:

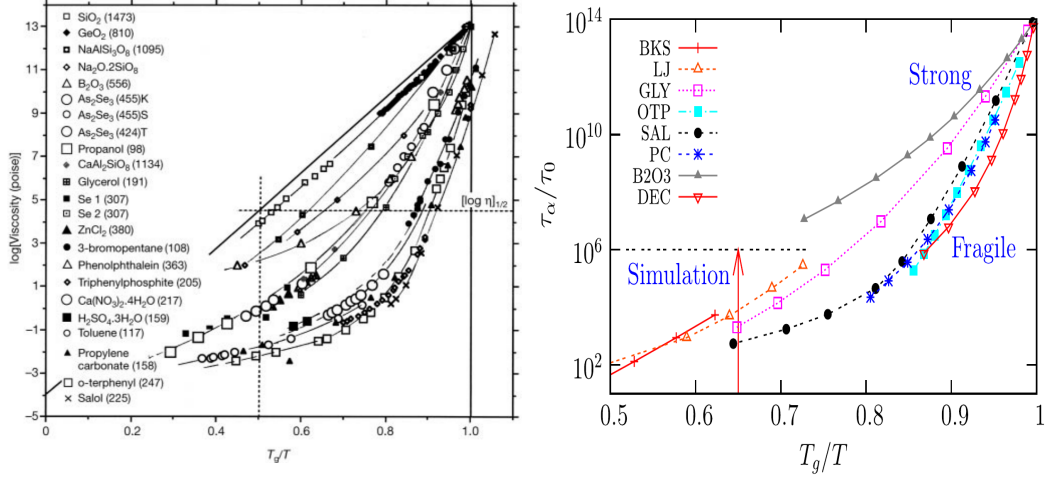
$$\tau_\alpha = \tau_0 \exp\left(\frac{D}{T/T_0 - 1}\right). \quad (1.3)$$

These two different types of growth are used to classify glass formers as *strong* or *fragile*, depending on whether they follow Eq. (1.2) or (1.3), respectively. It is a rather unfortunate terminology, since it is not related to the rigidity of glasses themselves. In any case, plotting^[24] the logarithm of relaxation time as a function of T_g/T results in the curves shown in Fig. 1.3 for several materials. Strong glass formers (such as SiO_2) follow straight lines, while fragile ones (*e.g.* o-terphenyl) first exhibit a relatively slow increase of τ_α , followed by a sharp super-exponential growth as $T \rightarrow T_g$. Note that data in Fig. 1.3a correspond to experiments where η was measured, while Fig. 1.3b contains data of τ_α both from experiments and from simulations. Thus, putting together these two figures validates the relation $\eta \propto \tau_\alpha$ mentioned above. In fact, an alternative definition for T_g is such that $\tau_\alpha(T_g) \sim 10^2 - 10^3$ s.

This last statement seems to reinforce the idea that T_g is a very arbitrary definition, because setting 100 s as the limit for τ_α apparently indicates that we are not patient enough to perform longer experiments. However, if considered carefully, the exponential and super-exponential growths shown in Fig. 1.3 indicate that any reasonable change in the threshold value of τ_α or η would have a very small effect in T_g ¹. In specific, at the melting point typical values of the viscosity are^[9] $\eta(T_m) \sim 10^{-2} - 10^{-3}$ Poise, which means that η increases by 15 orders of magnitude in the range $T_m \gtrsim T \gtrsim T_g$. Importantly, usual values of the glass transition temperature are around^[10] $T_g \approx 2T_m/3$. Therefore, it is clear that even if T_g is not a unique nor precise temperature, it identifies the region where characteristic time scales change dramatically.

To continue, let me address the physical meaning of the different terms that appear in the VFT law (1.3). In general, D and T_0 can be used just as fitting parameters. So, for instance, the smaller the value of D the more fragile the respective

¹Indeed, it can be shown^[9] that if we increase the threshold by a certain factor, its effect in shifting T_g to lower values is exponentially damped, at least.



(a) Angell plot of the logarithm of the viscosity as a function of (inverse) temperature, using experimental data from several materials. Taken from [25].

(b) Angell plot of the relaxation time as a function of (inverse) temperature. Points show both experimental and numerical data. Taken from [5].

Figure 1.3. The so called *Angell* (or sometimes *Arrhenius*) plots obtained by analysing the dependence of $\log \eta$ or $\log \tau_\alpha$ as a function of T_g/T . They allow to determine whether a given material is a strong or fragile glass former if its corresponding curve is a straight line or not, respectively.

glass. Similarly, T_0 indicates a temperature at which a divergence in the relaxation time is expected. Despite the fact that it plays the role of a fitting parameter, several studies suggest that it might be a relevant thermodynamic quantity[25]. This connection stems from the fact that its value is very close[5,26] to the Kauzmann temperature, T_K , which identifies a temperature at which the entropy of the glass (*i.e.* a state with a liquid-like structure) becomes *smaller* than the one of the crystal. A very counter-intuitive feature! Theoretically, the (supercooled) liquid's excess entropy with respect to the crystal, $S_{exc}(T) = S_{liq}(T) - S_{crys}(T)$ is considered and values of T_K are obtained by extrapolating the excess entropy until it vanishes, *i.e.* the condition $S_{exc}(T_K) = 0$ identifies T_K . In practice however, this point can never be reached since $T_K < T_g$, implying that the glass transition always intervenes, making it impossible for the system to equilibrate and thus to measure S_{exc} . Nonetheless, the excess entropy is a very important quantity because it is closely related with the configurational entropy, which serves as a measure of the amount of local minima present in energy landscape of the system; see Fig. 1.6. The so-called “Kauzmann paradox”[27] and the configurational entropy are very interesting topics that would deserve a section of their own, but length considerations dictate that I should omit them in this thesis. Fortunately, a good review is available in Ref. [28].

Now, given that the glass transition also marks an increase in the relaxation times of the system, it is only natural to investigate the effects on dynamical properties, instead of the static ones, as a system approaches T_g . For instance, Fig. 1.4 depicts the structure factor, $S(q)$, of a Lennard–Jones (LJ) fluid at different temperatures. $S(q)$ is defined in Eq. (1.18) below, but essentially it is the Fourier

transform of the radial distribution function, $g(r)$ (see Eq. (1.17)) and thus it provides information about spatial ordering present in the configuration. Data of the relaxation time (denoted in Fig. 1.4 by τ) at each temperatures are also included. Note that $S(q)$ is practically unchanged while τ increases by more than 3 orders of magnitude. This demonstrates quantitatively that, structurally, glasses are very similar to liquids.




Figure 1.4. (Partial) structure factor (S) of a Lennard-Jones mixture as a function of the wave-vectors, q , and for different temperatures. $S(q)$ is defined in Eq. (1.18) below. The relaxation time at each temperature (τ) is also indicated. Taken from [29].

In contrast, let us consider the mean square displacement (MSD),

$$\langle \Delta(t) \rangle = \left\langle \frac{1}{N} \sum_{i=1}^N |\mathbf{r}_i(t) - \mathbf{r}_i(0)|^2 \right\rangle, \quad (1.4)$$

where $\langle \bullet \rangle$ denotes thermal average. This variable will prove to be essential for many of the topics considered here, as shown later. In any case, in Fig. 1.5a the behaviour of the MSD is reported, considering again a LJ mixture at different temperatures. We can see that at high temperatures (curves to the left), the initial ballistic behaviour –characterized by $\Delta(t) \simeq t^2$ – quickly leads to a diffusive one –where $\Delta(t) \simeq t$. This latter regime is characteristic of a liquid in equilibrium, or even metastable equilibrium, as in the supercooled phase. On the other hand, as T decreases a plateau begins to form, separating these two different regimes of Δ . Importantly this figure also shows that the lower the temperature the longer the plateau, but its height remains roughly constant. Other dynamical variables present very similar features. An important case are dynamical correlation functions, such as the intermediate scattering function $F(q, t)$. Just as $S(q)$ provides information about the *static* density correlations at a length scale $\sim 1/q$, $F(q, t)$ measures the

correlations of two points separated by the same length, but at different times. Its temperature dependence is shown in Fig. 1.5b, once again for a LJ mixture. Note the marked slow down in the relaxation towards the (metastable) equilibrium state, identified by the lack of correlations. In fact, what these curves suggest is that as $T \rightarrow T_g$ relaxation bifurcates into two different processes, each of which has its own time scale.

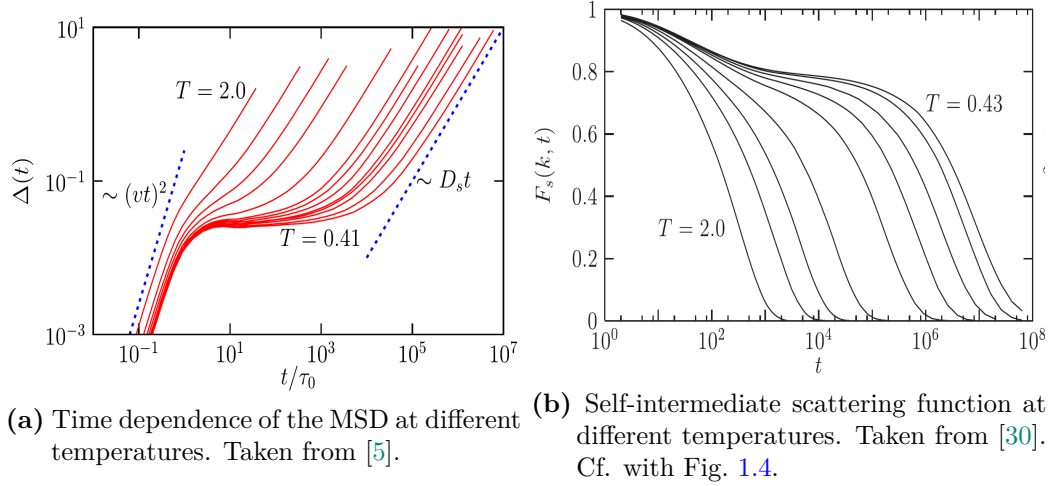


Figure 1.5. As $T \rightarrow T_g$, dynamical variables develop a plateau that separates their characteristic behaviours. Both plots show that the plateau lengthens as the temperature decreases. These curves are characteristic of the two-steps relaxation process, as discussed in the main text.

The potential energy landscape^[21] (PEL) picture, first introduced by Goldstein^[31], provides a fruitful scheme to understand these and several other properties of glassy systems. As sketched in Fig. 1.6, the PEL is a very rough surface in the configuration space and its several local minima identify possible metastable states in which a liquid can be trapped forming a glass if the temperature decreases below T_g . From this perspective, when the system's temperature is high enough, it is able to explore the full landscape and therefore equilibrium can always be attained. As T decreases, the system remains trapped for some relatively short time in a basin before it can escape and thus explore a larger region of the phase space. Then, as T is further decreased the time spent trapped in a basin increases exponentially until, at $T \lesssim T_g$, the system remains trapped in one of these local minima. At this point, ergodicity is broken and thus equilibration is impossible. Several clarifications are in place. First, in all this argument I blatantly ignored the crystalline state. So staying in equilibrium should be understood with respect to the supercooled liquid state. It is equally important to stress that the mechanism by which a system escapes one basin and jumps into another, when $T > T_g$, is by *very localized rearrangements of particles*. In other words, transitions between different local minima are accomplished by changing a sub-extensive number of coordinates^[9]. Hence, what in configuration space looks as jumping an energy barrier, in real space corresponds to small displacements of few nearby particles. Consequently, the vast majority of the system is virtually unaffected by the rearrangement of such cluster, and serves as a

rather constant background. Therefore, changing T does not impact the structure of the PEL itself, but instead influences how the system *samples it*^[10].




Figure 1.6. Sketch of the potential energy landscape of a glass former. Different basins define metastable states in which the system can remain trapped for sufficiently low temperatures. Its very rough structure implies that a vast amount of (nearby) minima exist, separated by relatively small energy barriers. Of course, the global minimum corresponds to a crystalline state, and yet, the sampling is dominated by the local minima due to their abundance. The transition to an ideal glass state is not discussed in this work. Taken from [10].

With this in mind, let us see how the PEL helps to explain the plateaus of the MSD and $F(q, t)$ in Fig. 1.5. The idea is that the initial and fast process (corresponding to the length of the plateaus) can be identified with relaxation *within* a basin, as the system gets temporarily trapped, while the second and slower relaxation corresponds to the system escaping such basin and exploring the rest of the configuration space. These two different processes are customarily called β and α relaxations, respectively. From this stand point we can better understand that the lower the T the longer the plateau –since the system takes longer to escape– but its height is roughly constant –because the process leading to a basin trapping is unaltered by T . Thus, the behaviour of Δ shown in Fig. 1.5a can be explained in the following manner. After the initial ballistic regime has ended, the system has reached a minimum and remains trapped in a situation where the main source of relaxation are vibrations around such minimum, identified with the β process. This coincides with the point where particles are *caged* by their neighbours and remain mostly arrested, leading to the constant value of Δ in the plateau. And recall that because T does not change the structure of the landscape, different minima are reached at roughly the same value of Δ . Then, when enough time has passed, the

system is able to leave the basin, with a concomitant relaxation process mostly driven by kinetic contributions, characteristic of the α branch. As expected, this also signals the point where Δ detaches from its plateau value and the diffusive regime takes over². Naturally, when $T < T_g$ the α relaxation never takes place and Δ saturates at its constant plateau value even for very long times. See a similar description of the different regimes of $F(q, t)$ in [29, Fig. 3]. Therefore, the PEL offers a very intuitive and powerful perspective for the study of glassy systems: several of its predictions have been verified experimentally^[10], while it also provides crucial information of the decorrelation rate of $F(q, t)$.

Another important theoretical framework is the so called *Mode Coupling Theory* (MCT). It is a very interesting and powerful scheme, but it certainly lies beyond the scope of this thesis. However, many good accounts are available, for instance: [13, 18, 20, 29]. Most of the phenomenology it predicts resembles closely the one obtained through the mean-field theory discussed in Sec. 1.3.2. Here, I will only mention that, a major success of MCT is being able to identify a “transition temperature”, T_{MCT} , at which the relaxation time diverges. Moreover, it successfully predicts the two steps relaxation of $F(q, t)$ by showing that its form is given by^[18]

$$F(q, t) \sim f + At^{-a}; \quad \text{and} \quad F(q, t) \sim f - Bt^{-b}. \quad (1.5)$$

Where the first relation holds as $F(q, t)$ approaches its plateau, and the second one as $F(q, t)$ decays from it. Sufficiently close to (but above) T_{MCT} , the two exponents a and b are related through

$$\frac{\Gamma(1-a)^2}{\Gamma(1-2a)} = \frac{\Gamma(1+b)^2}{\Gamma(1+2b)}. \quad (1.6)$$

This is an important relation to which I will return in a later section. The characteristic behaviour for $F(q, t)$ is sketched in Fig. 1.7. It also illustrates the connection of $F(q, t)$ with several important dynamical regimes discussed throughout this section. Notice the close resemblance with the intermediate scattering function of Fig. 1.5b obtained through simulations.

A final feature I would like to consider are the *dynamical heterogeneities*^[17] present in supercooled liquids and glasses alike. As the name suggests, it has been found that the dynamics in glassy systems are non-homogenous in space nor time. The first case is illustrated in Fig. 1.8a, where particles are coloured according to their displacement at a given time, in such a way that highly mobile (mostly fixed) ones are shown in red (blue) tones. Heterogeneity is visually clear from the figure: particles able to travel a (relatively) large distance form clusters –wherein all particles have similar mobility– that are surrounded by large portions of mostly immobile particles. Nevertheless, at later times the picture can be very different because arrested particles can become mobile. Conversely, those whose motion was initially rather unimpeded can afterwards remain blocked by their neighbours. Intuitively, this is the main component of the temporal heterogeneity in particles’ dynamics. Fig. 1.8b illustrates this idea more clearly by plotting the square displacement of individual particles. We can see that the particles’ trajectories consist

²It should be mention that the last part of the β relaxation can coincide with the initial part of the α one, so there is no clear cut distinction between these two processes^[29].




Figure 1.7. Typical form of the intermediate scattering function, Eq. (1.5), as derived from MCT. Each regime is identified with a characteristic dynamical regime, as illustrated by the insets. See also Figs. 1.5 and 1.9. Taken from [18].

mostly of vibrational motion for long intervals, interrupted by large, spontaneous jumps. Moreover, not all the jumps occur at the same time, signalling the temporally heterogeneous character of the dynamics. This also means that the appearance and disappearance of domains of particles with similar mobility is not a synchronized process, but occurs randomly. It is worth noting that the abrupt jumps are not reflected in the MSD, as shown by the smooth straight red line in the same figure.

As peculiar as dynamical heterogeneity can be, its relevance is mainly due to its role in the study of dynamical correlation functions and dynamical susceptibilities. To understand the connection, consider a set of particles separated by a distance r that are part of a high mobility cluster. If at later times the cluster transforms into a domain of mostly blocked particles, then the *fluctuations* of their mobilities would be correlated in time, and on a scale $\sim r$. However, to properly identify this sort of correlations *four point* correlation functions ($G_4(\mathbf{r}, t)$) and susceptibilities ($\chi_4(t)$) are required. From the (very important) relation $\chi_4(t) = \int d\mathbf{r} G_4(\mathbf{r}, t)$, it is easy to see that a high value of the susceptibility signals mobility correlations on a long scale. The typical behaviour of $\chi_4(t)$ at different temperatures (indicated by different curves) in a supercooled LJ mixture is shown in Fig. 1.9. For a fixed temperature, χ_4 increases monotonically until it reaches a maximum χ_4^* at $t = t^*$, and then decays. This behaviour actually makes sense: considering that the initial regime of the dynamics is ballistic, the whole system lacks correlation, leading to small values of the susceptibility. As time passes, clusters of particles with

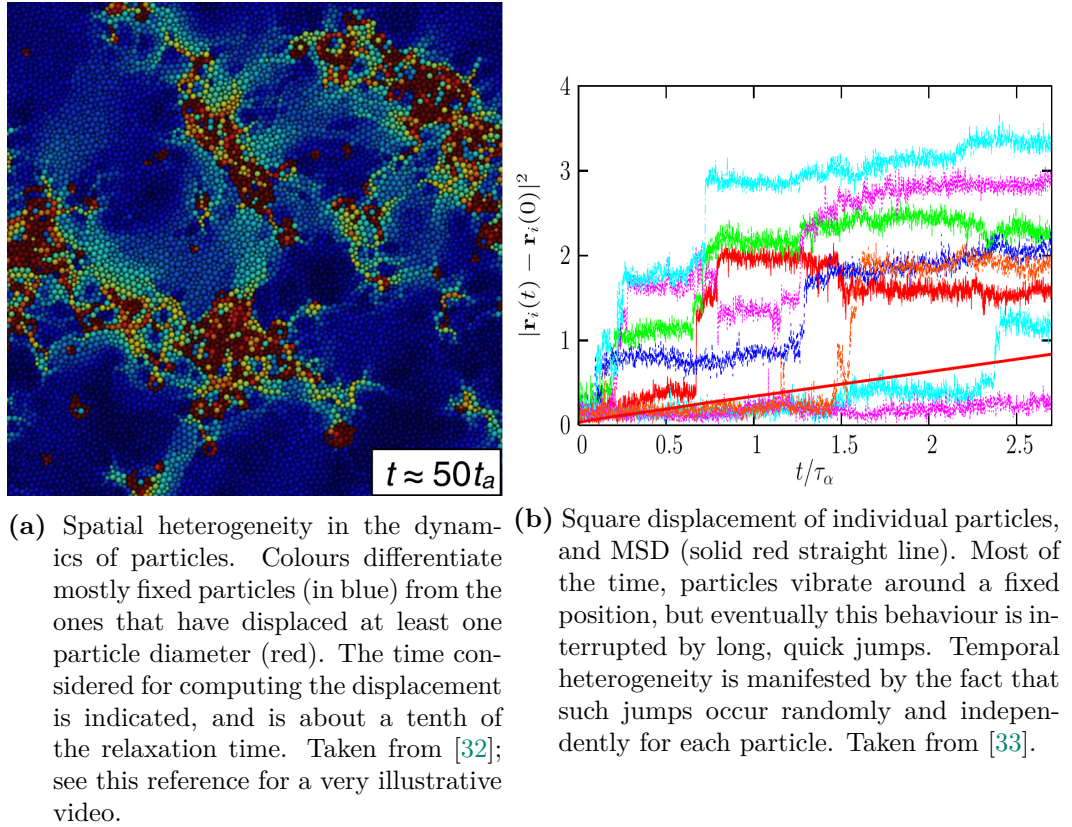


Figure 1.8. Dynamical heterogeneity in glassy systems. Glass formers usually show both spatial (panel a) and temporal (panel b) heterogeneity. It is a salient feature of this type of systems, but its connection with local structure is far from understood. See Sec. 1.1.1.

similar mobility start to form and thus correlations begin to appear; the larger the clusters the higher the correlation. Yet, these clusters cannot be maintained forever, since the system, being in the supercooled phase, is bound to eventually relax to equilibrium. When this happens, the correlations are lost and consequently χ_4 decreases. From this argument, it is natural to expect that t^* should be similar to the relaxation time, and this is indeed what happens^[5,9]. Moreover, as the temperature decreases, both t^* and χ_4^* constantly increase, signalling the presence of higher, longer lived cluster correlations. Interestingly, Mode Coupling Theory predicts that χ_4 should actually diverge at T_{MCT} . Unfortunately, this temperature is rather high^[9,18] when compared to T_g . At any rate, χ_4 plays the role of the usual “two point” susceptibility of critical phenomena, with its characteristic divergent behaviour at a phase transition. Within this framework, we are also able to identify the corresponding order parameter: the dynamical correlation function³! What is more, a very similar scenario is found in a famous spin glass mean-field model: the p -spin model^[34–38]. There, the corresponding χ_4 diverges as the system approaches its *dynamical* transition at T_d , the model’s analogue of T_{MCT} .

³Actually, a more precise statement is that the order parameter should be the *long time limit* of such correlation function, see [9, Sec. 7.8]

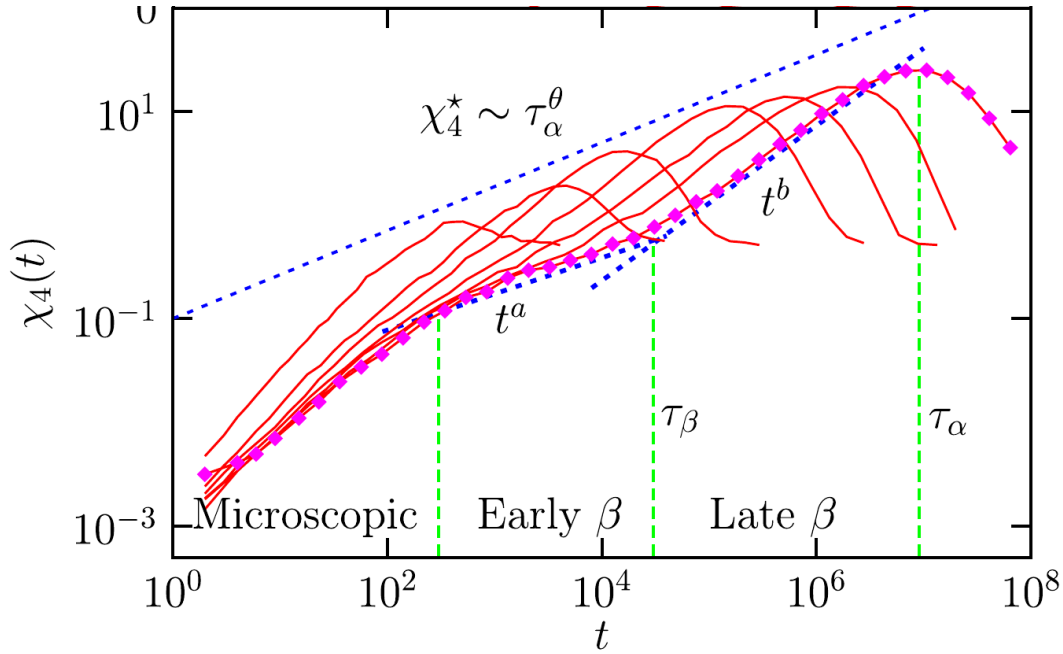


Figure 1.9. Four point susceptibility as a function of time in a LJ supercooled liquid. It is a very useful for studying dynamical heterogeneities in supercooled liquids. Different curves show that as temperature decreases (from left to right), mobility fluctuations grow in the system, as indicated by the larger values of χ_4 . In particular, its maximum, closely related to the relaxation time τ_α , is shifted to larger times. The connection with the two steps relaxation, depicted in Fig. 1.7, is also indicated. Taken from [5].

These results support the hypothesis that by studying the behaviour of χ_4 it is possible to detect the onset of “glassy behaviour”, characterized by the non-trivial dynamical correlations, whose length scales increase as the temperature is lowered. However, the analogies have to be taken with care since no phase transition occurs in real systems, as evinced by the fact that χ_4^* just keeps on growing even *below* T_{MCT} . Nevertheless, these theoretical considerations suggest that T_{MCT} might be a temperature more physically meaningful than T_g , and that the putative divergence of χ_4 is avoided possibly due to barrier crossing^[9]. Furthermore, just as in the usual case of critical phenomena, a correlation length^[39] (ξ_4) can be found that is expected to diverge at T_{MCT} , but that in supercooled liquids indicates the length scale of mobility correlations; thus it is essentially a measure of the size of the clusters of particles’ with the same mobility. Moreover, as discussed in Sec. 1.3.2, mean-field theory provides an exact description that resembles closely the phenomenology of MCT and qualitatively agrees with the results observed empirically.

1.1.1 Local structure and its connection with dynamics

Now that we have a clearer picture of the intriguing dynamical phenomena that ensue near the glass transition, it is worth bringing back the main obstacle for a complete physical description of it: finding a connection between the complex dynamical features of supercooled liquids and glasses and their local structure. In

other words, so far I have described the peculiar dynamical properties of supercooled liquids as well as some tools to unveil and analyse them. Yet, I have not mention how such properties can be explained in terms of structural (and hence *static*) information of the system. The problem is notably intricate, so an entire sub-section seemed (just) enough to give an account of such pressing question. Nevertheless, I should anticipate that, maybe unsurprisingly, there is a lack of consensus on which are the correct variables to consider nor on their relative degree of success.

Ever since the discovery of dynamical heterogeneities people have been looking for connections with different structural variables^[5,17,40,41]. Most of the works proceed in a similar fashion: the mobility of individual particles is measured by their squared displacement ($|\delta \mathbf{r}_i(t)|^2 = |\mathbf{r}_i(t) - \mathbf{r}_i(0)|^2$), and then interactions with their local environment are used to construct physical observables, expecting to find significant correlations between them. Of course, the crux of these studies consists in where such locality ends and which are the relevant physical quantities. In any case, with this approach, the influence of thermal fluctuations, shear deformations, etc. on particles rearrangements and plastic dislocations are studied. The numerous set of structural quantities that have been proposed is a proof of how difficult the problem is. Relevant (and personally biased) examples include: point-to-set correlations^[42–45] (where a length scale over which particles are pinned by their neighbours is identified; see Fig. 1.10a), local thermal energy^[46] (where the thermal energy of each interacting pair is used to predict plastic instabilities related to anharmonicity; see Fig. 1.10b), or plain geometrical parameters like bond orientations between nearby particles^[47–49] (Fig. 1.10c). This list is far from exhaustive, but an insightful review can be found in [50]. A fruitful approach uses the vibrational modes of the system^[51–54], and relates the density of states (DOS) with the collective dynamics of the system. This link is justified because the Fourier transform of the velocity autocorrelation function is closely related to the DOS^[52]. This relation, however, relies on the validity of an harmonic approximation for the interaction energy of the system, a scenario that in several cases might not be true. Furthermore, the vibrational spectra is also a relevant property of jammed and nearly jammed systems^[55–61], so it will have an important role in this thesis; see Secs. 1.4.3, 1.4.5.1 and Sec. 3.7 in Chapter 3.

An alternative and rather more intuitive approach to the same problem can be adopted by using the so called “isoconfigurational ensemble”^[62–65] (ICE) where the dynamics of a system is studied by simulating several trajectories, departing from the same initial configuration (hence the name). For instance, if molecular dynamics (MD) simulations are employed, the particles’ initial position is fixed, while the velocities are randomly assigned at each run according to the Maxwell-Boltzmann distribution with a given temperature. Clearly, this method provides a way of sampling the space of all possible trajectories of the configuration and, in doing so, it allows to identify how the local environment of a particle influences its average mobility. Its only basic assumptions are that (i) the thermal noise is washed out when a large enough number of MD realizations are used; and (ii) that trajectories are self-averaging. The main advantage of the ICE is that regions where clusters of particles undergo major rearrangements during their dynamics can be easily identified. Moreover, such regions can be related to structural properties such as the Debye–Waller factor or the local energy density^[66,67]; see Fig. 1.11. Recently,

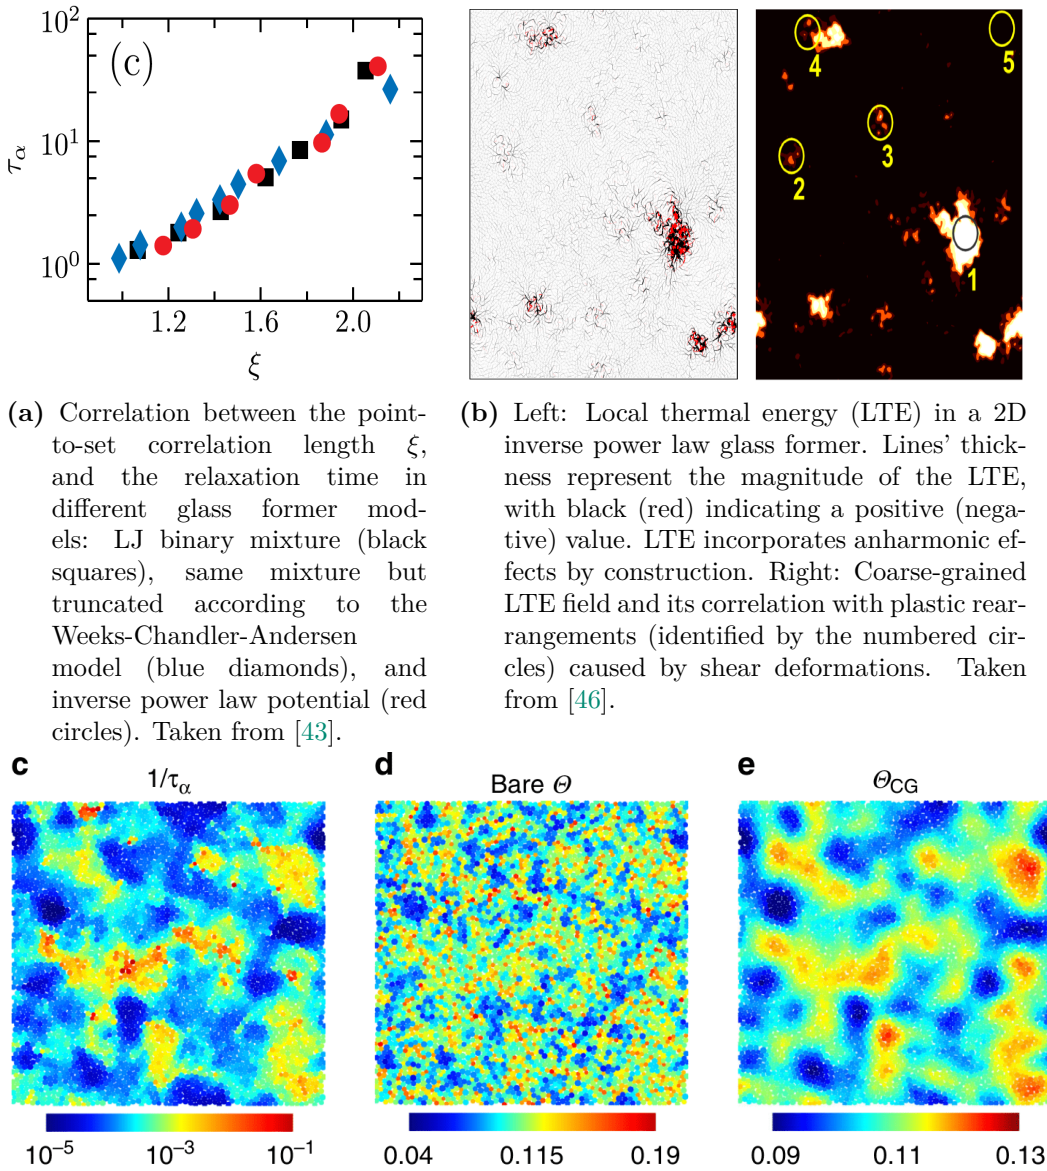


Figure 1.10. Different methods for connecting the local structure with dynamical features or plastic rearrangements in a variety of models.

information-theoretic methods have been used to assess more accurately the role of these and other structural quantities^[65], as well as how reproducible the observed trajectories are. However, it has also been pointed out that providing a link between dynamical and static properties using exclusively the particles' mobility might be inadequate^[68]. Moreover, the main problem of the ICE approach is that even if a

structural variable shows a high correlation with the dynamics in a given system, it may lead to poor predictions in a different model^[69]. In other words, sampling from the ICE leads to structure-based inferences whose quality might be very sensitive to type of glass former model employed. These features will be discussed in more detail in Chp. 3, where I will show that the ICE can also be used to study the dynamics of amorphous solids near their jamming point and, notably, overcome several of the issues just mentioned.

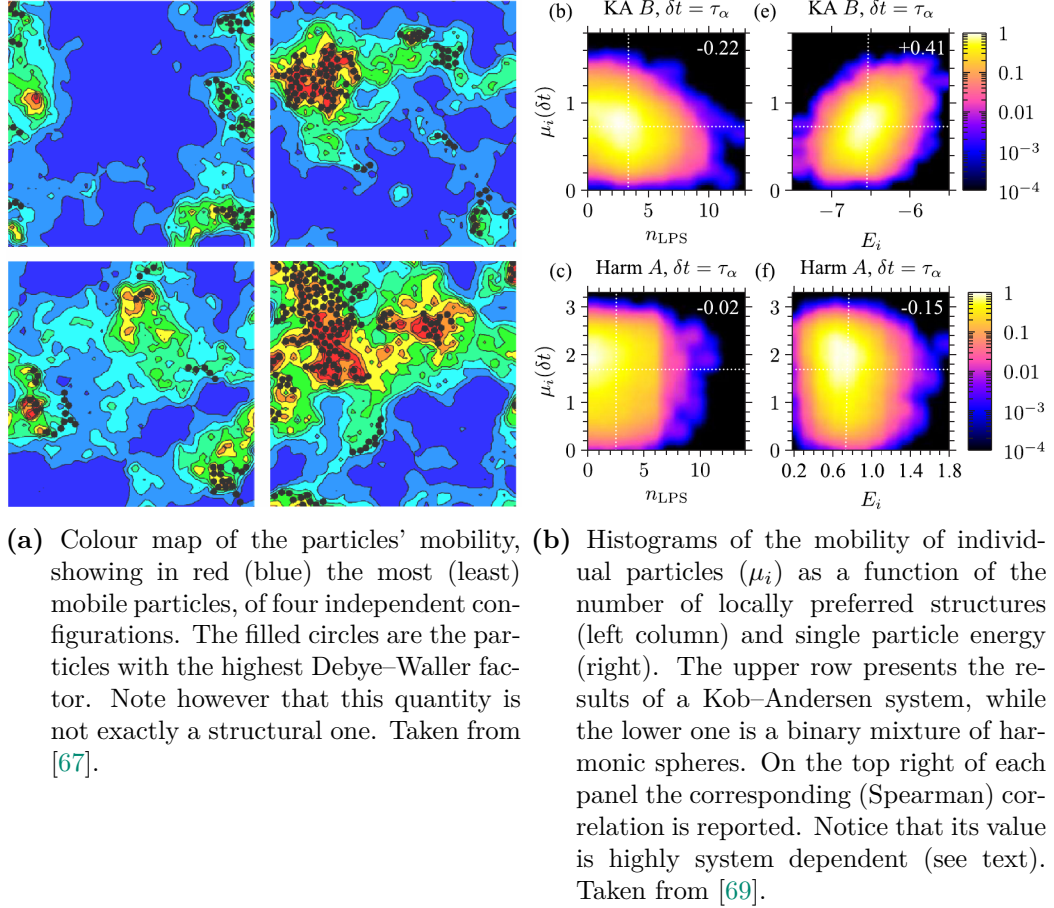


Figure 1.11. The usual approach in the isoconfigurational ensemble is to compute the particles' propensity, defined as the mobility averaged over several trajectories, and test whether it is correlated with physically sound (but arbitrary) structural variables.

On the other hand, a major step forward has been recently achieved by employing Machine Learning methods^[70,71]. In these works, the information of a particle's local structure is encoded in a new variable termed “softness” that is then related to the displacement of such particle in a given time interval. A Support Vector Machine is used to find a hyperplane (in features space) dividing the mostly movable particles from the mostly arrested ones. Then, the softness is computed as the signed distance of each particle's features to such an hyperplane, thus “soft” particles are prone to be displaced by a significant amount while “hard” ones will remain mostly fixed (see Fig. 1.12a). It has been shown that softness is strongly

correlated with physical quantities such as local energy and coordination number^[72], as well as being a useful variable for modelling the Arrhenius behaviour observed in supercooled liquids^[71,73,74]. Yet, it performs better at identifying mobile particles than these other structural variables. Furthermore, new results using Graph Neural Networks^[75] have shown to outperform several physically motivated variables in inferring (and predicting) dynamical properties of glassy systems (Fig. 1.12b). However, these methods are based on parametrizing a particle's local environment in terms of (artificial) feature functions, yielding a synthetic representation of the structure. Therefore, even though Machine or Deep Learning techniques yield high quality predictions about which particles are the likeliest to be highly mobile, their associated structural variables lack a clear physical meaning. In short, even if Machine Learning methods can be used to construct good predictors of particles' mobility, they fail to provide an answer about which are the real physical variables that determine such mobility.

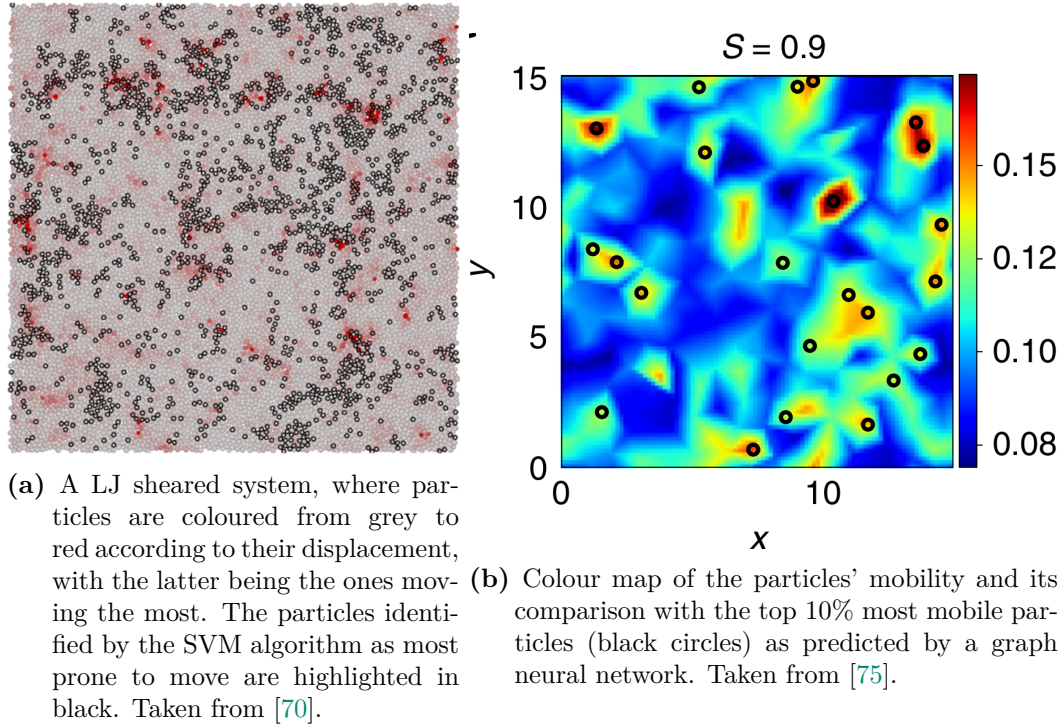


Figure 1.12. Machine Learning methods yield high quality predictions of particles rearrangements, but they rely on artificial representations of the particles' local environment.

As a final remark, note that for several of the methods described here, additionally to the structural variable considered, a coarse-graining procedure is needed to obtain better inferences; *e.g.* Figs. 1.10b, 1.10c, 1.11a, 1.12b. Likewise, in the case of the SVM, even though predictions for individual particles can be obtained, the features functions are constructed essentially as weighted radial distribution functions. Hence, structural information of a small region is compressed into a single number, and can thus be consider as another instance of coarse-graining. In contrast, in Chp. 3 I will show that the method we developed for inferring the statistics of the

particles' dynamics works using only information of individual and well identified particles.

To close this introductory section I should emphasise that there are many interesting topics that I have left completely uncover such as: ageing^[19], spin glasses^[76–78], complexity and configurational entropy^[28], the ideal glass transition^[79], kinetically constrained models^[80], and a long etcetera.

1.2 Hard sphere systems: From liquid to jamming with a glassy interlude

In this section, I will quickly review some of the basic properties of simple hard sphere (HS) fluids, based mostly on Refs. [16, 81, 82]. That is, systems composed of frictionless and infinitely rigid spherical particles. Hence, no interaction is present between the spheres except when they come into contact with one another and an elastic collision occurs. Notwithstanding its simplicity, HS liquids display a non-trivial behaviour giving rise to a liquid-solid phase transition. This is summarized in Fig. 1.13 in the pressure-density plane for a monodisperse configuration. The solid line represents the stable equilibrium branch, showing that the liquid freezes at a density⁴ $\varphi_f \approx 0.494$, while the solid melts at $\varphi_m \approx 0.545$. In other words, for densities larger than φ_m an equilibrated HS system is in a crystalline phase, which ends when the spheres occupy the maximum possible volume. As conjectured by Kepler and proved much later^[83], this happens when the spheres are packed following a face-centred cubic (FCC) structure and each particle is in contact with its closest neighbours. The associated density is $\varphi_{FCC} = \frac{\pi}{3\sqrt{2}} \approx 0.740$, which thus acts as an upper limit for the density in monodisperse HS systems. Fig. 1.13 also shows (dashed curve) the metastable liquid branch, which has been found to finish at the so called “random close packing” density, $\varphi_{RCP} \approx 0.64$. I will discuss at length what happens *near* and *at* this point in the following sections and chapters, but for the time being let me just mention that configurations with φ_{FCC} and φ_{RCP} are two examples of jammed packings. On the other hand, the presence of the metastable branch for $\varphi > \varphi_f$ suggests that HS liquids could also be used as glass-formers. And indeed a glass transition, taking place around^[82,84] $0.58 \lesssim \varphi_g \lesssim 0.62$, has been found. Here, two points are worth considering. First, note that for HS fluids the melting, glass, and other transitions are specified in terms of the density instead of the temperature, in contrast with the systems considered in the previous section. The reason is that in HS systems, the interaction energy is either zero or infinity, and therefore the temperature only acts as a scaling parameter of the different thermodynamic variables; see *e.g.* Eqs. (1.25), (1.26), (1.30) below. I will argue in the rest of this section that HS liquids exhibit many of the complex phenomenology described above for usual supercooled liquids and glasses. But to make a proper comparison, it is useful to map the role of specific volume and temperature (in standard glass formers) to the pressure and density (in HS systems).

The second point worth mentioning is the rather unexpected crystallization at high densities in HS liquids. Indeed, how come there is a phase transition if no

⁴Throughout this thesis, the terms density and packing fraction will be used interchangeably and will be denoted by φ .

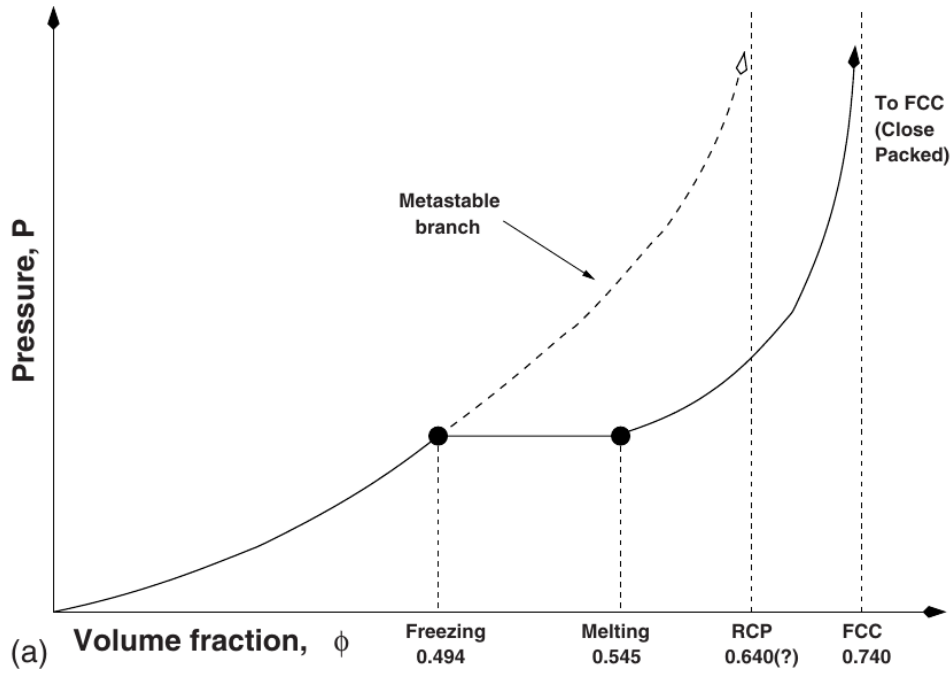


Figure 1.13. Phase diagram of monodisperse HS fluids. The solid black line is the equilibrium branch, while the dashed one indicates the metastable liquid (and glass) phase, beyond the freezing point (first dotted vertical line). The second vertical line marks the melting point, while the last two indicate the maximal possible density of the glass and solid phases, respectively. As discussed in the text, the former is associated to the so called *random close packing*, and the latter to an FCC crystal. Taken from [84].

attractive force is present? Moreover, because the potential energy of HS fluids is always zero, the energetic contribution to their free energy is only due to a trivial kinetic term. This means that for large packing fractions the entropy of the crystal becomes *larger* than the liquid's one. This very counterintuitive result brings us back to the Kauzmann paradox mentioned above. But, mainly, it shows that there is no reason *a priori* to assume that the ordered, crystalline phase has a lower entropy than the inherently disordered liquid^[5]. This sort of “entropic ordering” can be understood by considering that the entropy has a contribution associated with the (logarithm of the) accessible volume per particle^[85]. The idea is that, while arranging the particles following a regular structure certainly reduces the entropy due to orientational degrees of freedom, it may also increase the part related to translational degrees of freedom. For high enough densities, the gain in the latter exceeds the loss caused by a reduced number of possible orientations and therefore (at least partial) crystallization occurs. In other words, the only way in which particles can benefit from a larger accessible volume, and thus allow the system to further increase its density, is by forming crystalline domains. Historically, the case of HS is very important because the first “empirical” signature of the freezing transition was obtained in the debut of molecular dynamics^[86] and Monte Carlo^[87] simulations as computational techniques to study the thermodynamics of many body systems. The case of hard disks followed soon afterwards^[88]. As will be discussed in Chapter

2, the simple potential of HS systems makes their dynamics trivial, rendering them specially amenable for the computational capabilities available in the late 1950's. The details of the story are fascinating –see *e.g.* [89–91]– starring an episode where several scientific personalities pondered about the freezing phase transition and “decided” *by voting* whether it really occurred or not. I cannot help finding Sisyphean that, so many years later, we are still studying HS systems, with ever faster and more powerful computers and techniques, and yet marvelling by the many striking physics hiding underneath.

1.2.1 Basic Statistical Mechanics of liquids

Before proceeding in the analysis of HS systems, it is convenient to develop some general results of liquids theory, which is that I will do next. Later, in Sec. 1.2.2, I will particularize such results to HS liquids and show that many simplifications are possible due to their peculiar potential, see Eq. (1.23). Additionally, this section is important because I will introduce some notation that will be used throughout this work. First of all, N will denote the number of particles (or system's size), \mathbf{r}_i will be used to indicate the centre's position of the i -th particle, while $\vec{\mathbf{r}} = \{\mathbf{r}_i\}_{i=1}^N$ will be the $3N$ -dimensional vector that specifies the location of the system in configuration space. Analogously, let m_i and \mathbf{p}_i be the mass and momentum vector of the i -th particle and $\vec{\mathbf{p}}$ the $3N$ -dimensional vector of momenta. Clearly, the same notation is applicable when dealing with systems in dimensions higher than $d = 3$, and I will do so when needed. Hence, the state of the system in phase space is specified in general by $\vec{\mathbf{z}} \equiv (\vec{\mathbf{r}}, \vec{\mathbf{p}})$.

To continue, I will assume that particles interact via a (sufficiently well-behaved)⁵ potential function, $u(r)$, that only depends on the distance between pairs of particles, $r_{ij} \equiv |\mathbf{r}_i - \mathbf{r}_j|$. We can thus define the Hamiltonian as

$$\mathcal{H}(\vec{\mathbf{z}}) = \sum_{i=1}^N \frac{|\mathbf{p}_i|^2}{2m_i} + U(\vec{\mathbf{r}}) = \sum_{i=1}^N \frac{|\mathbf{p}_i|^2}{2m_i} + \frac{1}{2} \sum_{i \neq j}^{1,N} u(r_{ij}). \quad (1.7)$$

I will focus on deriving the equation of state (EOS) of the fluid in the canonical ensemble, at a fixed temperature $T = 1/\beta$ and volume V . The partition function is given by

$$Q_N(\beta, V) = \frac{1}{N! h^{3N}} \int d\vec{\mathbf{z}} e^{-\beta \mathcal{H}(\vec{\mathbf{z}})} = \frac{Z_N(\beta, V)}{N! \Lambda^{3N}}, \quad (1.8)$$

where h is the Planck's constant and $\Lambda = \sqrt{\beta \hbar^2 / m}$ is the De Broglie wavelength, that appears when the Gaussian integral of each component of $\vec{\mathbf{p}}$ is performed. On the other hand, Z_N defines the configurational integral

$$Z_N(\beta, V) = \int d\vec{\mathbf{r}} \exp(-\beta U(\vec{\mathbf{r}})). \quad (1.9)$$

Note that the dependence of Z_N on V is implicitly contained in the integration limits of $\int d\vec{\mathbf{r}}$.

⁵What is meant by “sufficiently well behaved”, of course depends on which properties we are interested in. For instance, it is clear that the HS interaction, Eq. (1.23), defines a potential that is *not* continuous, yet it is suitable to model the properties of HS systems in the thermodynamic limit. For a more detailed discussion see [16, Sec. 2.1.1] and references therein.

From the partition function (1.8), we obtain the Helmholtz free energy:

$$F(N, V, T) = -T \log Q_N(T, V) = 3TN \log \Lambda - T \log \left(\frac{Z_N(\beta, V)}{N!} \right). \quad (1.10)$$

Obviously, if for a given potential $u(r)$ we were able to compute the integral (1.9) exactly and thus obtain a close form for F , all the thermodynamic variables would follow easily. For instance, the average energy and pressure are obtained simply by differentiating the free energy:

$$\langle E \rangle = \frac{\partial(\beta F)}{\partial \beta}, \quad (1.11a)$$

$$P = -\frac{\partial F}{\partial V}. \quad (1.11b)$$

Unfortunately, no closed expression is known for Z_N for the HS potential, let alone a general one as I am considering for the moment. Nevertheless, it is still possible to obtain some important results and thus gain some insight of the physical properties such as the variables just introduced. To do so, let us first note that in the thermodynamic limit, $N \rightarrow \infty$ and $V \rightarrow \infty$, the (intensive) free energy becomes a function of the temperature and the *number density*, $\rho \equiv N/V$. Thus, letting

$$f(\rho, T) \equiv \lim_{N, V \rightarrow \infty} \frac{F(N, V, T)}{N},$$

the free energy can be conveniently decomposed into two independent parts: (i) a term related to the energy of the system as if it was an ideal gas; and (ii) an “excess” contribution taking into account the interaction between particles. That is^[16,81],

$$f(\rho, T) = f^{(\text{id})}(\rho, T) + f^{(\text{ex})}(\rho, T) \quad (1.12a)$$

where

$$f^{(\text{id})}(\rho, T) = 3T \log \Lambda - T(1 - \log \rho) \quad (1.12b)$$

$$f^{(\text{ex})}(\rho, T) = -T \lim_{N, V \rightarrow \infty, \rho = N/V} \frac{1}{N} \log \left(\frac{Z_N}{N!} \right). \quad (1.12c)$$

Note that this implies that in the thermodynamic limit the pressure is simply given by^[16]

$$P = \rho^2 \frac{\partial f(\rho, T)}{\partial \rho}. \quad (1.13)$$

This seems a rather unimpressive rewriting of Eq. (1.11b), but it makes clear that, for a constant T value, the pressure can be written only in terms of ρ , for instance, as a power series:

$$\beta P = \rho + \sum_{n=2}^N B_n(T) \rho^n. \quad (1.14)$$

This is the famous *Virial Expansion*. The first term is easily obtained from the contribution of $f^{(\text{id})}$ to the total free energy. Therefore, terms of order 2 and higher

in ρ are due to “deviations” from the ideal gas caused by the interaction between particles.

As I will argue next, working in terms of densities brings several other advantages; see also Sec. 1.3.2. So let me spend some time working out some of the most important results. First of all, given the Hamiltonian of Eq. (1.7), we can compute the probability of finding a given configuration $\vec{\mathbf{r}}$ as

$$\varrho(\vec{\mathbf{r}}) = \frac{1}{N!h^{3N}Q_N} \int d\vec{\mathbf{p}} e^{-\beta\mathcal{H}(\vec{\mathbf{r}},\vec{\mathbf{p}})} = \frac{1}{Z_N} e^{-\beta U(\vec{\mathbf{r}})}.$$

From this distribution and the indistinguishability of particles it follows that the so called *n-particle density*, is given by

$$\varrho_N^{(n)}(\vec{\mathbf{r}}^{(n)}) = \frac{N!}{(N-n)!Z_N} \int d\vec{\mathbf{r}}^{(N-n)} \exp(-\beta U(\vec{\mathbf{r}})), \quad (1.15)$$

where $\vec{\mathbf{r}}^{(n)}$ is the $3n$ -dimensional vector obtained from considering the components of $\{\mathbf{r}_i\}_{i=1}^n$, while the integral of $d\vec{\mathbf{r}}^{(N-n)}$ is performed over the remaining coordinates, *i.e.* $\{\mathbf{r}_i\}_{i=n+1}^N$. This type of functions determine the probability of finding n particles in the volume element $d\vec{\mathbf{r}}^{(n)}$ independently of the positions of the rest of the particles and the configuration’s momenta. From the definition of Z_N , Eq. (1.9), it is easy to obtain that

$$\int d\vec{\mathbf{r}}^{(n)} \varrho_N^{(n)}(\vec{\mathbf{r}}^{(n)}) = \frac{N!}{(N-n)!},$$

and therefore, for $n = 1$,

$$\int d\mathbf{r} \varrho_N^{(1)}(\mathbf{r}) = N.$$

It then follows that, for a uniform system the single-particle density coincides with the number density, $\varrho_N^{(1)}(\mathbf{r}) = \rho$.

A closely related set of functions are the *n-particle distributions*, $g_N^{(n)}(\vec{\mathbf{r}}^{(n)})$, defined in terms of the n-particle density as

$$g_N^{(n)}(\vec{\mathbf{r}}^{(n)}) = \frac{\varrho_N^{(n)}(\vec{\mathbf{r}}^{(n)})}{\prod_{i=1}^n \varrho_N^{(1)}(\mathbf{r}_i)}. \quad (1.16)$$

I will only be interested in homogeneous systems and thus, this last equation reduces to $\rho^n g_N^{(n)}(\vec{\mathbf{r}}^{(n)}) = \varrho_N^{(n)}(\vec{\mathbf{r}}^{(n)})$. By far, the most important of these functions is the pair distribution function, $g_N^{(2)}(\mathbf{r}_1, \mathbf{r}_2)$. It is amply used to distinguish between gas, liquid and solid phases, but it provides a lot of information about the structure of a system as I will show next. For isotropic systems, as I will consider here, $g_N^{(2)}(\mathbf{r}_1, \mathbf{r}_2) = g_N^{(2)}(\mathbf{r}_1 - \mathbf{r}_2)$, and it is called *radial distribution function* (RDF). Furthermore, for central potentials the RDF only depends on the norm of its argument and will be simply denoted $g(r)$. From the identity,

$$\langle \delta(\mathbf{r} - \mathbf{r}_1) \rangle = \frac{1}{Z_N} \int d\vec{\mathbf{r}} \delta(\mathbf{r} - \mathbf{r}_1) e^{-\beta U(\vec{\mathbf{r}})} = \frac{1}{Z_N} \int d\vec{\mathbf{r}}^{(N-1)} \exp(-\beta U(\mathbf{r}, \mathbf{r}_2, \dots, \mathbf{r}_N))$$

it follows easily^[81] that $\varrho_N^{(1)}(\mathbf{r}) = \langle \delta(\mathbf{r} - \mathbf{r}_i) \rangle$. An analogous relation for $\varrho_N^{(2)}(\mathbf{r}_1, \mathbf{r}_2)$ leads to a convenient representation for $g(\mathbf{r})$ in homogeneous systems:

$$g(\mathbf{r}) = \frac{1}{N\rho} \left\langle \sum_{i \neq j}^{1,N} \delta(\mathbf{r} - \mathbf{r}_j + \mathbf{r}_i) \right\rangle \implies g(r) = \frac{1}{4\pi N r^2 \rho} \left\langle \sum_{i \neq j}^{1,N} \delta(r - r_{ij}) \right\rangle; \quad (1.17)$$

where the rightmost equation follows if the system is isotropic.

The RDF allows to compute several structural properties of interest. For instance, defining the *total correlation function*, $h(\mathbf{r}) = g(\mathbf{r}) - 1$, the static structure factor can be obtained by computing its Fourier transform:

$$S(\mathbf{q}) \equiv 1 + \rho \int d\mathbf{r} h(\mathbf{r}) e^{-i\mathbf{q} \cdot \mathbf{r}} = 1 + \rho h(\mathbf{q}). \quad (1.18)$$

As mentioned in the previous section, $S(\mathbf{q})$ measures the density fluctuation in a scale of order $\sim 1/|\mathbf{q}|$ and is a very relevant quantity since it can be measured accurately in experiments. Moreover, it can be directly linked to response functions, such as the isothermal compressibility, $\kappa \equiv \left(-V \frac{\partial P}{\partial V}\right)^{-1}$ ^[82]:

$$\frac{\rho\kappa}{\beta} = 1 + 4\pi\rho \int dr r^2 h(r) = S(0). \quad (1.19)$$

Even more importantly for our purposes, $g(r)$ is related to the average energy and pressure through the following equations:

$$\langle E \rangle = \frac{3}{2}NT + 2\pi N\rho \int dr u(r)g(r), \quad (1.20a)$$

$$\frac{\beta P}{\rho} = 1 - \frac{2\pi\beta\rho}{3} \int dr r^3 u'(r)g(r). \quad (1.20b)$$

Comparing expression (1.20b) with the virial expansion, Eq. (1.14), suggests at a first sight that only the ρ^2 term is needed in order to compute the pressure of the liquid. Unfortunately, this is not the case because $g(r)$ itself depends on the density and on $\varrho_N(2)$. Thus, ρ is implicitly contained in the integrand of Eq. (1.20b).

Finally, another important quantity is the *direct correlation function*, $c(r)$, defined implicitly through the integral equation^[81,82]

$$h(r) = c(r) + \rho \int d\mathbf{r}' c(|\mathbf{r} - \mathbf{r}'|)h(r'), \quad (1.21)$$

called Ornstein–Zernike relation. Intuitively, it states that the total correlation between a pair of particles is made of two contributions: first, the direct correlation between both particles; and, in second place, an indirect contribution caused by correlations with the rest of the particles as intermediaries. Note however that, even though this equation defines $c(r)$, it is not possible to obtain a closed expression for it. Nevertheless, the Fourier transforms of $h(r)$ and $c(r)$ are related through simple, algebraic expressions:

$$\hat{h}(q) = \frac{\hat{c}(q)}{1 - \rho\hat{c}(q)}.$$

Whence another expression for the compressibility is readily obtained,

$$\frac{\rho\kappa}{\beta} = \frac{1}{1 - \hat{c}(0)}. \quad (1.22)$$

The set of equations (1.19), (1.20), (1.22), together with the Ornstein–Zernike relation Eq. 1.21, are important because they provide direct routes to compute the EOS, *provided we knew the exact form of* $g(r)$ –or $h(r)$, $c(r)$ or even their Fourier transforms for that matter. Unfortunately, except for the ideal gas, this is never the case and several approximations are needed to make any progress. Nevertheless, in HS systems the situation is more tractable and many important results are readily available as discussed next.

1.2.2 Equation of state of HS liquids

The HS interaction is modelled through a contact, pair-wise potential of the following form:

$$u_{HS}(x; a) = \begin{cases} 0 & x > a, \\ \infty & x \leq a. \end{cases} \quad (1.23)$$

Naturally, if $\vec{\sigma} = \{\sigma_i\}_{i=1}^N$ is the set of diameters, $a = \frac{\sigma_i + \sigma_j}{2}$ is a parameter equal to the sum of a pair’s radii. However, throughout this part I will assume that all particles are identical (thus $\sigma_i = \sigma$) and henceforth omit reference to the parameter a unless it is necessary. Next, given that a particle’s volume is equal to $v = \pi\sigma^3/6$ and that spheres never overlap, the system’s density or packing fraction is simply

$$\varphi = Nv/V = \rho v, \quad (1.24)$$

where V is the system’s volume (usually a cubic box of size L), while ρ is the number density as above.

Next, notice that for u_{HS} the integral of Eq. (1.20a) identically vanishes, whence $\langle E \rangle = \frac{3}{2}NT$. This makes sense since HS liquids do not have any potential energy and therefore $\langle E \rangle$ is purely kinetic. Similarly, Eq. (1.20b) can be considerably simplified. Introducing the *cavity distribution function*^[81] $y(r) = g(r)e^{\beta u_{HS}(r)}$, and a function $w(r) = e^{-\beta u_{HS}(r)}$ the integrand can be written as

$$u'_{HS}(r)g(r) = -\frac{1}{\beta}w'(r)y(r) = -\frac{1}{\beta}y(r)\delta(r - \sigma).$$

The last equation follows from the fact that $w(r) = 1$ when $r > \sigma$ and vanishes otherwise; *i.e.* it is the step function $w(r) = \Theta(r - \sigma)$. It is important to mention that $y(r)$ is a continuous function even if the potential is not^[81,82], and therefore the integral is well defined. The δ -function makes the integration trivial thus obtaining

$$\frac{\beta P}{\rho} = 1 + 4\varphi g^+; \quad g^+ \equiv \lim_{r \rightarrow \sigma^+} g(r). \quad (1.25)$$

This means that the full equation of state for HS systems can be derived solely from the value of $g(r)$ at the contact distance. Although such value is not known exactly, many good approximations have been devised^[82] with ever greater accuracy.

Before continuing, it is worth recalling that in HS liquids temperature only acts as a scale parameter, fixing the average kinetic energy. For instance, consider the configurational integral:

$$Z_N = \int d\mathbf{r} e^{-\beta U_{HS}(\mathbf{r})} = \int d\mathbf{r} \prod_{i<j}^{1,N} w(r_{ij}) = \int d\mathbf{r} \prod_{i<j}^{1,N} \Theta(r_{ij} - \sigma).$$

This last expression clearly shows that its value is independent of T . For convenience I will henceforth adopt a “temperature free” description for HS systems. Therefore, many of the quantities I will describe next are assumed to be scaled by T and ρ in order to make them dimensionless. The most relevant example is the *reduced pressure*,

$$p \equiv \frac{\beta P}{\rho}, \quad (1.26)$$

which will appear often.

Now, as mentioned above, were we to know exactly either $g(r)$ or $c(r)$ ⁶ the EOS would follow easily. But these two correlation functions are actually coupled through the Ornstein–Zernike relation. To overcome this additional obstacle and obtain a closed integral equation approximations are unavoidable. The idea is to write $c(r)$ as a functional of $h(r)$, $\mathcal{C}[h(r)]$. One of the most famous of such expressions is the Percus–Yevick (PY) closure relation^[82],

$$\mathcal{C}[h(r)] = (1 - h(r))(1 - e^{\beta v(r)}). \quad (1.27)$$

Importantly, for HS systems the integral equation derived from the PY closure can be solved exactly and, for instance, obtained a simple expression for g^+ :

$$g^+ = \frac{1 + \frac{1}{2}\varphi}{(1 - \varphi)^2}.$$

Plugging it into Eq. (1.25) we obtain a (long sought) EOS:

$$p_{PY} = \frac{1 + 2\varphi + 3\varphi^3}{(1 - \varphi)^2}. \quad (1.28)$$

Unfortunately, if we proceeded by first computing the compressibility, say, via Eq. (1.19) or (1.22), a *different* EOS is obtained^[82]:

$$p_{PY,\kappa} = \frac{1 + \varphi + \varphi^2}{(1 - \varphi)^3}. \quad (1.29)$$

This rather disappointing result is termed *thermodynamic consistency problem*. It is related to the different *routes*^[82] in which an EOS can be obtained. In general, Eqs. (1.20) are called the energy and virial routes, while Eq. (1.19) corresponds

⁶The path I have been following in fact suggests that $c(r)$ is a secondary function, defined in terms of $h(r) = g(r) - 1$, and therefore implying that the RDF is fundamental function to consider. However, as argued in [81, Chp. 3] $c(r)$ is more properly defined as a functional derivative of F with respect of $\varrho_N^{(1)}(\mathbf{r})$, and thus independent of $g(r)$. Similarly, it can also be shown that the RDF is obtained from the functional derivative of F with respect to v .

to the compressibility route. Even when there are other types, these are the most common ones. What causes the thermodynamic consistency problem is that once a given approximation for the RDF (or the other related quantities such as $S(q)$) is chosen, the free energy obtained when tracing back each of these routes is different. This is evinced, for instance, by noting that according to Eq. (1.13) we can recover f by integrating with respect to φ the EOS for the pressure. But given that the available EOS's are different (cf. Eqs. (1.28), (1.29), or (1.31) below), we would obtain a *different* expression for the free energy depending on which expression for p we choose.

On the other hand, let us consider the EOS resulting from the virial expansion, Eq. (1.14). As with other quantities in HS systems, the virial coefficients $\{B_n\}_{n=2}^{\infty}$ are independent of the temperature. Moreover, values of B_2 , B_3 , and B_4 are known exactly in $d = 3$ and higher dimensions^[16], while numerical estimates are known up to B_{12} for spherical particles^[82]. When the first few terms are included, the following series is obtained for the reduced pressure^[81]:

$$p = 1 + 4\varphi + 10\varphi^2 + 18.365\varphi^3 + 28.224\varphi^4 + 39.82\varphi^5 + 53.34\varphi^6 \dots \quad (1.30)$$

Surprisingly, a very good approximation is obtained if the virial coefficients are restricted to integer values. This was first proposed by Carnahan and Starling (CS) in [92], by noting that the first 2 coefficients are integers, while the next few ones can be well approximated by $B_n = n^3 + 3n$. Plugging this expression into Eq. (1.14) results in a power series of φ that can be easily summed, resulting in a fairly simple formula:

$$p_{CS} = \frac{1 + \varphi + \varphi^2 - \varphi^3}{(1 - \varphi)^3}. \quad (1.31)$$

Although p_{CS} can also be obtained from a combination of p_{PY} and $p_{PY,\kappa}$, its merit relies on its notable accuracy (compared with numerical simulations) in a wide range of densities. Although it tends to underestimate the pressure, it does so by a small amount (about 3%)^[81]. And while some other rather simple approximations for the EOS exist^[93], the CS expression is accurate enough for the topics that I will consider.

Once an EOS for HS systems is available, we can explore how the liquid branch continues beyond the freezing and melting points, thus forming a “super-compressed” liquid. In Fig. 1.14 I present numerical results of such experiment, in the $(\varphi, 1/p)$ plane. The simulations were done using an event-driven molecular dynamics algorithm^[94] with a Lubachevsky–Stillinger compression protocol⁷. Different curves correspond to different compression rates (κ) and the CS EOS is also included for comparison (black dashed line). First of all, note that the curve corresponding to the fastest compression is always off the EOS line, indicating that the system is never allowed to relax and remains out of equilibrium. However, by reducing the compression rate the configuration is able to remain in the liquid branch well beyond the freezing density, until at higher densities it eventually falls out of equilibrium. In turn, when κ is further decreased the pressure follows very closely the EOS until, around φ_m , the configuration suddenly crystallizes (partially). This is signalled by a sudden decrease in p (seen as a “jump” in the pink and brown curves

⁷This algorithm will be discussed in much more detail in the next chapter.

of Fig. 1.14), caused by the particles rearranging themselves in a regular structure, increasing their free volume as discussed above.

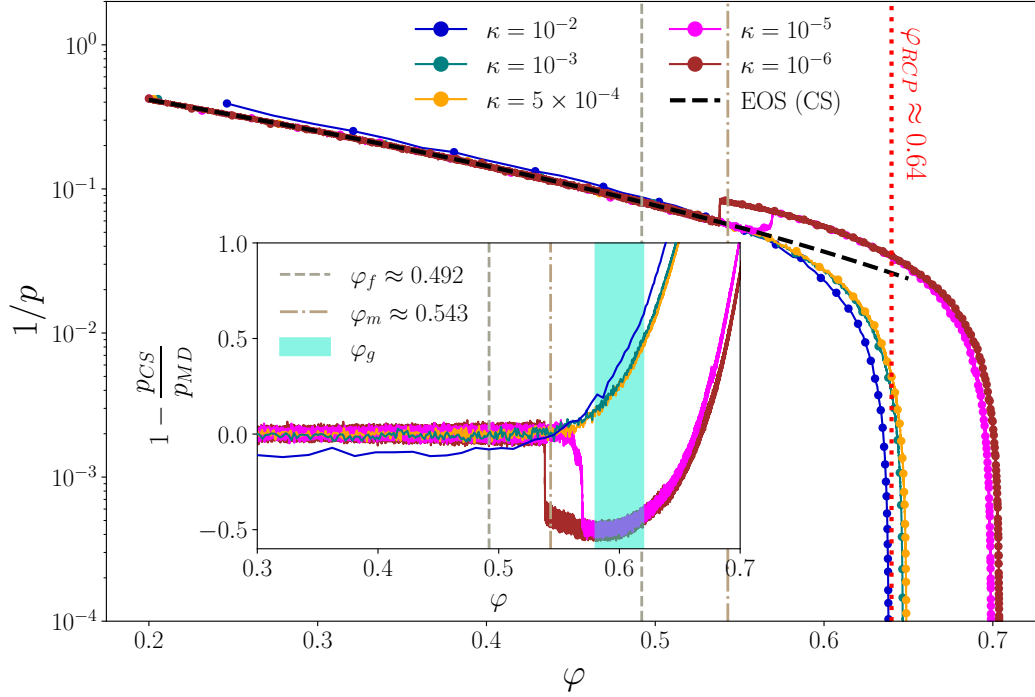


Figure 1.14. Main: Inverse (reduced) pressure as a function the packing fraction. Markers correspond to estimations of p through molecular dynamics simulations^[94] with different compression rates (κ), as indicated by the labels. The system closely follows the liquid CS EOS (black, dashed line) even beyond the freezing point, φ_f . Inset: Accuracy of the CS EOS. It also shows that the system is unable to remain in the liquid branch much beyond the melting density φ_m . The region identified with the glass transition is shaded in light blue. See text for a more detailed discussion.

The most salient feature of these results is that they qualitatively reproduce the phenomenology discussed in Sec. 1.1 for other glass-former models. For instance, compare Figs. 1.2 and 1.14: it is clear that by rotating the former by 90° the curves of both figures have the same qualitative behaviour. Moreover, notice that in all the cases where the system is able to remain in the liquid phase, all curves approach the $p \rightarrow \infty$ limit around the same packing fraction, namely, $\varphi_{RCP} \simeq 0.64$. Such value is evidently different from the one attained by systems with partial crystallization in the same high p limit, as can be seen from the gap in the two groups of curves as $1/p \rightarrow 0$. This gap also mimics the analogous volume difference between a glass and a crystal in the cases mentioned in the previous section.

Results of Fig. 1.14 also show that monodisperse HS systems in $3d$ have a high propensity to crystallize. That is the reason why it is very hard to extend the liquid branch beyond φ_m , even by small amount. Although the same happens with other types of algorithms^[95], in this case it is essentially due to the fact that I have employed a linear compression protocol. Consider, for instance, the inset of Fig. 1.14, which shows the difference between the pressures obtained according

to the EOS from Eq. (1.31) and the numerically from molecular dynamics. We can observe that even if κ is chosen so that crystallization is avoided, the liquid surpasses coexistence region $[\varphi_f, \varphi_m]$ just slightly. This is, in fact, a general feature of protocols with constant compression (or cooling). As explained in detail in [9, Sec. 3], using constant cooling rates it is impossible to avoid both the formation of crystals and prolonging the metastable liquid phase until the glass transition is reached. Analogously for HS systems, by using a constant compression rate the liquid is bound either to form crystalline domains or fall out of equilibrium soon after φ_m , and certainly before the glass transition density. To overcome this difficulty, a common technique is to use a polydisperse mixture (so crystallization is suppressed) and thus allow the liquid to continue for considerably higher densities⁸. As we will see next, polydispersity and a clever type of Monte Carlo algorithms can be used to explore what happens deep in the super-compressed and glassy phases. But in any case, all the considerations so far show that HS systems can be used as a minimal model to study the physics of glassy systems.

Let me close this subsection with a remark about the validity of the extrapolations of the liquid's EOS. Note that the EOS we considered here, Eqs. (1.28)-(1.31), does not contain any singularity when $\varphi < 1$. Such singularity would identify the end of the liquid branch, so this property is obviously unphysical, given that all the region $\varphi > \varphi_{FCC} \simeq 0.74$ is inaccessible. More realistic relations can be obtained from approximated EOS's that are constructed explicitly to match the observed behaviour in the high density regime. A myriad of proposals have been put forward^[96], with different degrees of success, and with p developing a singularity at some density $\varphi \in [\varphi_{RCP}, \varphi_{FCC}]$. It is widely believed that such divergence is reached at either of the interval's extremes, but it has not been established if $p = \infty$ at the random close packing or FCC densities^[82,96]. However, in a recent work^[97] the authors showed that polydisperse liquids can be thermalised even when $\varphi \gtrsim \varphi_{RCP}$.

1.2.3 Glass phase and Gardner transition in HS systems

Let us now analyse the deeply compressed regime of HS systems. From now on, I will assume that no order is present in the configuration, and therefore the $p \rightarrow \infty$ divergence occurs at $\varphi_{RCP} \approx 0.64$, or equivalently, that the accessible jammed states correspond to randomly arranged configurations. In other words, I will henceforth assume that the only possible jamming density is $\varphi_J = \varphi_{RCP}$.

To begin this part, let us consider the glass EOS^[94,98,99] in d dimensions:

$$p(\varphi; \varphi_J) = \frac{d}{1 - \varphi/\varphi_J}. \quad (1.32)$$

This equation can be derived from the free-volume approximation of the partition function for HS, which becomes exact in the $N \rightarrow \infty$ limit, and the condition of forming a stable packing when particles are in contact^[98]. Importantly, this EOS depends explicitly on the jamming packing fraction. But this is natural given that

⁸In higher dimensions such a trick is unnecessary because the crystalline state can be effectively avoided with particles of the same size, despite ordered configurations having a smaller free energy^[16].

φ_J marks the point where the free volume per particle vanishes. However, I want to stress that Eq. (1.32) is *not* a thermodynamic EOS, *i.e.* it has no free energy associated with it, nor does it incorporate any reference to the equilibrium liquid state. Yet, it describes quite accurately the values of p and φ found in numerical simulations as I will show here. I should also mention that no exact thermodynamic EOS is known for HS in any finite d , but in Sec. 1.3.3 I will describe how such relation is constructed within mean-field theory.




Figure 1.15. Main: Liquid (black, dashed) and glass (red) EOS in a HS system and its comparison with molecular dynamics simulations. Numerical results are obtained using two different compression protocols as indicated by the legend in the plot. Inset: Divergence of pressure, as expected from the glass EOS of Eq. 1.32. Note however that this equation does not inform us how to obtain the density at the singularity, φ_J .

In any case, a relevant property that is readily derived from Eq. (1.32) is that $p \sim (\varphi_J - \varphi)^{-1}$; this is an important scaling that will show up again in Secs. 1.4.2 and 2.2. The accuracy of Eq. (1.32) is put to test in Fig. 1.15, where the pressure is computed numerically using the same protocol^[94] than in the previous figure. In order to closely follow the equilibrium branch for each value of φ , I used two different compression rates. The system is initialized as a low density liquid and is compressed relatively fast up to $p = 500$. In this way, a glass is formed by suppressing eventual relaxation to the crystalline phase. Then, to simulate a well thermalised glass, a small value of κ is used to further compressed the system, approaching the $p \rightarrow \infty$ limit. The excellent agreement between numerical estimations of p and the those from Eq. (1.32) is outstanding, specially considering its simple form. Nonetheless, it is important to mention that the value of φ_J was obtained after bringing the configuration to its jamming point using the iterative Linear Programming (iLP)

algorithm described in Chapter 2. It should also be considered that although these MD simulations allow to attain very high pressures, extrapolations to $1/p = 0$ might yield a *different* jammed state than the one obtained using iLP. This feature will also be analysed in Chapter 2.

On the other hand, the simple aspect of Fig. 1.15 can be misleading, given that a very rich phenomenology is “hidden” deep in the glass phase. However, to properly investigate this regime numerically, a special type of Monte Carlo (MC) algorithms^[97,100,101] must be used that, rather unfortunately, only work in polydisperse systems. The idea is that simulations are performed using the standard MC dynamics but, after a fixed number of steps a swap between two particles is performed with a given probability. Impressively, this type of SWAP algorithms are able to equilibrate liquids *well beyond* the glass transition of various systems. To explore the glass phase in HS configurations, once an equilibrated liquid has been generated at $\varphi > \varphi_g$ ⁹ an annealing compression is performed, *e.g.* using the same MD protocol as above. This results in the phase diagram of Fig. 1.16, taken from Ref. [99]. A lot of information is contained in this picture, so let us analyse it carefully. First of all, the (estimated) glass transition density is indicated with the star at φ_d , while the equilibrated liquid states are the green squares on top of the liquid EOS. Each of them is then compressed into the glass region until the jamming density is reached; notice that they follow the glass free-volume EOS, *i.e.* the analogous of Eq. (1.32) but for configurations with polydispersity. However, the most salient feature is that the glassy phase is divided in two by another transition (red line) that takes place *within* such phase. This corresponds to a Gardner transition^[102] –named after Elizabeth Gardner, who in [103] discovered the corresponding phenomenon in the p -spin– that divides a region of stable glasses from one where stability is only marginal.

Many interesting features come about inside the marginal, Gardner phase, with the phase diagram further complexified by considering different types of potentials and the effects of shear^[104–107]. Once again however, space-time limitations dictate that I should omit giving a detailed description and thus I only refer to the very illuminating accounts [102] and [108]. Suffice it to say that a peculiar characteristic of glasses in their Gardner phase is the formation of hierarchical basins in their free energy landscapes. Of course, detecting such structure is no trivial task^[109–111] and, as explained in Ref. [99], the best signature of the Gardner transition is the appearance of a peak in the *fluctuations of dynamical variables*. Specifically, the variance and skewness of overlap between the displacement of two clones of a configurations. The difference of such fluctuations is illustrated in the insets of Fig. 1.16, where at the centre of the i -th particle a sphere is drawn, with a radius proportional to $|\mathbf{r}_i^{(A)} - \mathbf{r}_i^{(B)}|$, *i.e.* the difference in displacement between two clones of the system. The snapshot that belongs to the stable glass phase does not show any peculiar structure. In contrast, beyond the Gardner line signatures of spatial heterogeneity are clearly present.

Importantly, jammed states are always located within the Gardner phase. This is not exclusive of HS system, but instead a general property of several models.

⁹Note that the value of φ_g depends on the amount and type of polydispersity employed. Nevertheless, φ_g and φ_J are generally larger compared to those of monodisperse configurations.

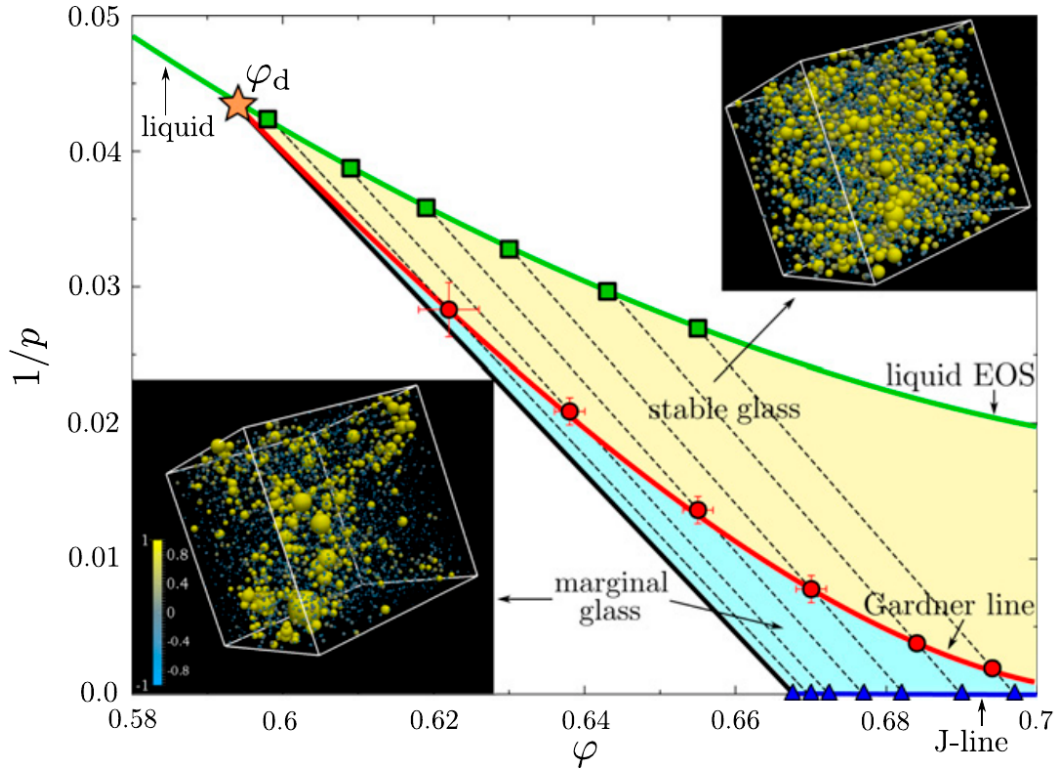


Figure 1.16. Phase diagram of $d = 3$ HS polydisperse mixture. The liquid and *equilibrated* glass EOS are shown in green, while the line where the Gardner Transition ensues is shown in red. Note that this line separates the glass phase in two regions: above the GT glasses are stable, but below they are only marginally so. This latter phase ends in the jamming line (in blue). Markers corresponds to values obtained through numerical simulations. See text for a more detailed discussion. Taken from Ref. [99].

Moreover, given that jammed configurations are identified with minima of the free energy landscape, understanding its structure is a fundamental step towards building a complete theory of jamming. In the $d \rightarrow \infty$ limit of mean-field, the situation has been carefully analysed^[16,112] and theory predicts a sharp transition in the Gardner line^[111]. Yet, in $d \leq 3$ systems the situation is less clear due to the possibility of a crossover or an avoided transition^[113,114] and is further muddled by the *absence* of the marginal phase in some exemplary glass-former models^[115]. Nevertheless, very near the jamming point numerical studies are promising^[3,111,116] and suggest that the fractal, ultrametric structure of the energy landscape, predicted by mean-field theory, remains valid.

1.3 Mean-field theory: The $d \rightarrow \infty$ limit

In this section, I will go through the main results of a recently developed mean-field (MF) theory for glasses and jammed systems. The theory itself is a *tour de force* –see *e.g.* Refs. [84, 111, 112, 117, 118] among others– that provides an *exact* description of the transition from a liquid to a glass, as well as the physics that ensues near and at the jamming point. In other words, it is an exact theory of the

glass and jamming transitions. Furthermore, it also provides an accurate account of what happens as a system goes across such transitions, through its glassy phase. Naturally, the Gardner transition (GT), introduced in the end of the last section, is the most important feature that a glass stumbles upon in such path. Actually, the GT was first predicted and verified in MF models^[110,112] and later discovered in other glass formers.

Yet, the cost paid for exactness is that systems are assumed to have very high dimensionality. It is undoubtedly intriguing that the $d \rightarrow \infty$ limit renders a given scheme more tractable and makes possible an analytical treatment. Hence, in this section I will try to give a very summarized account of why this is so. Happily for the curious reader, this MF theory has been recently put together in a coherent, self-contained, and pedagogical form in Ref. [16]. A short and very nice review is also available in Ref. [104], which conveniently focuses on HS systems in large dimensions, and therefore I will follow it in the rest of this section (except in Sec. 1.3.1 which is based on Ref. [16]). Besides, adhering to its spirit (and further shortening the theory) I will not present any details of the calculations involved, but will mostly restrict myself to showing how many of the phenomena that have been described in this chapter are the “low dimensional reminiscence“ of the MF predictions.

1.3.1 Liquids in high dimensions

In my opinion, the first thing that needs to be explained is why the high d regime simplifies the analysis of liquids and glasses. However, this is not an exclusive property of this kind of systems. For instance, consider N Ising spins, $\vec{S} = \{S_i\}_{i=1}^N$ with $S_i = \pm 1$, whose Hamiltonian is given by

$$\mathcal{H}(\vec{S}) = -\frac{1}{2} \sum_{i,j}^{1,N} J_{ij} S_i S_j - \sum_{i=1}^N B_i S_i, \quad (1.33)$$

where B_i is a local field acting on the i -th spin. $J_{ij} = J_{ji}$ are the exchange couplings between spins i and j , and naturally $J_{ii} = 0$. Eq. (1.33) is very general, since it can describe many spins systems with pairwise interactions by the appropriate choice of couplings. For instance, fully connected models are such that $J_{ij} \neq 0$ for all $i \neq j$, while the a graph or lattice structure can be easily included by setting $J_{ij} \neq 0$ if an edge connects spins i and j , and $J_{ij} = 0$ otherwise. An important case is the d -dimensional cubic lattice, where each spin interacts with its $2d$ nearest neighbours. It is a well known result from statistical mechanics that the free energy of the system, F , can be expressed, in principle, as a function of the configuration's local magnetizations, $\vec{m} = \{m_i\}_{i=1}^N$, with $m_i = \langle S_i \rangle$. This means that, thermodynamically, the configurations observed in such spin system are the ones that produce a given \vec{m}^* such that the free energy is minimized. That is, \vec{m} plays the role of an order parameter and the configurations that dominate the statistics are the ones satisfying

$$\left. \frac{\partial F(\vec{m})}{\partial m_i} \right|_{\vec{m}^*} = 0, \quad \forall i = 1, \dots, N. \quad (1.34)$$

Of course, finding a closed expression for $F(\vec{m})$ is extremely difficult, if not impossible, even in simple cases, making the equation above unusable. Resorting to

approximations is thus unavoidable. Among the several possibilities, high temperature (small β) expansions are quite fruitful. For instance, considering terms up to order β^3 it can be shown^[119] that

$$\begin{aligned} -\beta F(\vec{m}) = & -\sum_i \left[\frac{1+m_i}{2} \log \frac{1+m_i}{2} + \frac{1-m_i}{2} \log \frac{1-m_i}{2} \right] \\ & + \beta \left[\frac{1}{2} \sum_{i,j} J_{ij} m_i m_j + \sum_i B_i m_i \right] + \frac{\beta^2}{4} \sum_{i,j} J_{ij}^2 (1-m_i^2)(1-m_j^2) \\ & + \frac{\beta^3}{6} \left[2 \sum_{ij} J_{ij}^3 m_i (1-m_i^2)(1-m_j^2) + \sum_{i,j,k} J_{i,j} J_{i,k} J_{j,k} (1-m_i^2)(1-m_j^2)(1-m_k^2) \right]. \end{aligned} \quad (1.35)$$

In the case a cubic lattice of dimension d , with an uniform external field, and setting¹⁰ $J_{ij} = \frac{1}{2d}$, the situation is considerably simplified: translational invariance indicates that $m_i = m \forall i$. The expression above then reads

$$\frac{F}{N} = -\frac{1}{\beta} s_0(m) - \frac{1}{2} m^2 - Bm - \frac{\beta}{8d} (1-m^2)^2 - \frac{\beta^2}{12d^2} m^2 (1-m^2)^2. \quad (1.36)$$

I have introduced $s_0(m) = -\left[\frac{1+m}{2} \log \left(\frac{1+m}{2} \right) + \frac{1-m}{2} \log \left(\frac{1-m}{2} \right) \right]$ to denote the entropy associated to a single spin. Comparing Eqs. (1.35) and (1.36) it is easy to conclude that the $d \rightarrow \infty$ limit is equivalent to the $\beta \rightarrow 0$ one. Moreover, the series in powers of $1/d$ maps precisely to the high temperature expansion. And given that for high T the phenomenology is usually much simpler, this analogy suggests that by considering large values of d theory might end up being more tractable.

In the case of liquids the advantages of a large d are even clearer. To see why, let first us consider the virial expansion of the free energy. Combining Eqs. (1.12) and (1.14) we have that

$$-\beta f(\rho, T) = -\beta f^{(id)}(\rho, T) - \sum_{n=2}^{\infty} \frac{B_n}{n-1} \rho^{n-1}. \quad (1.37)$$

Of course, some clarification is in order to ease the (reasonable) doubts as to why we can use a set of $3d$ equations to describe large dimensional systems. First of all, we should consider that just as in the spin systems the free energy can be written in terms of \vec{m} , in liquids theory F can be expressed as a *functional* of the systems density, $F[\varrho(\vec{r})]$. The formalism thus obtained, termed *density functional theory*, is very powerful and actually easily generalizable to an arbitrary number of dimensions^[16,81]. The analogies with spin systems are indeed deep, since the equilibrium thermodynamic state can also be found as the density profile $\varrho(\vec{r})^*$ that makes the (functional) derivative of F vanish, *i.e.*

$$\left. \frac{\delta F[\varrho(\vec{r})]}{\delta \varrho(\vec{r})} \right|_{\varrho(\vec{r})^*} = 0.$$

¹⁰This choice guarantees that the interaction energy remains extensive in the thermodynamic limit, and of the same magnitude as the external potential.

Moreover, if we impose that solutions to the equation above should be uniform and isotropic, crystalline states can be readily discarded. And indeed it can be shown that $F[\varrho(\vec{r})]$ is minimized by $\varrho(\vec{r}) = \rho$, that is, by a system with uniform density^[16]. Notice that this is just mimicking the homogenous solution $m_i = m \ \forall i$ discussed above by enforcing an uniform external magnetic field.

Hence, once that we know that F can be expressed, for any d , as a function of the system's density, it is natural to write it as the series expansion of Eq. (1.37). Clearly, the values of the virial coefficients B_n depend, in general, on the dimensionality of the system and the interaction potential. But nevertheless, for sufficiently low densities and high temperatures, such virial expansion should be reasonably accurate.

The charm of the high dimensional limit is that it allows to truncate the series just after the first virial term. More precisely, if the density of the system does not scales exponentially with d , it can be shown that B_2 dominates the whole series. Therefore, under these conditions it is possible to only consider the ideal gas contribution and the first virial coefficient to obtain the free energy, and concomitantly the EOS. Besides, the value of B_2 is known exactly for HS systems in any dimensions:

$$B_2 = \frac{1}{2}2^d V_d; \quad V_d = \frac{\pi^{d/2} \sigma^d}{2^d \Gamma(1 + d/2)}. \quad (1.38)$$

In this last expression, V_d is the volume of a d -dimensional hypersphere of diameter σ and Γ is the Gamma function. Introducing the rescaled packing fraction,

$$\tilde{\varphi} = 2^d \varphi = 2^d \rho V_d,$$

we have that, for any d , the free energy and (reduced) pressure are, to leading order,

$$-\beta f(\rho, T) = -\beta f^{(id)}(\rho, T) - \frac{\tilde{\varphi}}{2}, \quad (1.39a)$$

$$p(\rho) = 1 + \frac{\tilde{\varphi}}{2}. \quad (1.39b)$$

The proof^[16] of this result involves (i) identifying the different terms of the virial expansion as (irreducible) Mayer diagrams; (ii) showing that ring diagrams dominate at each order; and (iii) that their contribution to the n order of the free energy is such that

$$|f_n| \sim \tilde{\varphi}^{n-1} \left(\frac{n^{n-2}}{(n-1)^{n-1}} \right)^{d/2}.$$

Now, because for $n \geq 3$ the term in parenthesis is smaller than 1, $|f_n| \rightarrow 0$ as $d \rightarrow \infty$. This means that Eqs. (1.39) are actually *exact* if conditions (ii) and (iii) can be met. Or, in other words, if the series (1.37) is convergent. Currently, a lower bound of its convergence radius is known: $\tilde{\varphi}^{(conv)} \geq 0.145$. Within the interval $\tilde{\varphi}^{(conv)} \leq \tilde{\varphi} < 1$ the series is conjectured to be convergent, while if $1 < \tilde{\varphi} < e^{\gamma d}$ with $\gamma < (1 - \log 2)/2$ the series is divergent, but this would imply that the packing fractions grows exponentially with d . Nonetheless, even in this case resumming the virial expansion results in a series that to leading order also agrees with Eqs. (1.39).

Importantly, very similar results hold for several other potentials, like LJ and SS whose interaction is a power-law of the overlap between them. See [16, Sec. 2.3] for details.

1.3.2 Dynamics and the glass transition

We have seen that the thermodynamics of liquids in high d can be solved exactly. But another important feature of the MF theory is that also the *dynamics* is amenable to an exact treatment^[120]. As a major outcome, the glass transition comes about within this perspective as a *sharp dynamical transition*. As we will see, in finite d systems such transition is blurred due to the finite lifetime of metastable states. Nevertheless, many of the phenomena described in Sec. 1.1 can be understood as the manifestation of the (avoided) large d transition.

The solution begins by considering that particles follow a Langevin dynamics, with a friction term ζ and white noise ξ with zero mean and variance $\langle \xi_i^\mu(t) \xi_j^{\mu'}(t') \rangle = 2T\zeta \delta_{ij} \delta_{\mu\mu'} \delta(t-t')$. Whenever equilibrium can be reached an important result follows: the long time limit of the MSD (Δ) can be used as the order parameter relevant to the dynamic transition. Let me briefly describe the argument. First, the equilibrium assumption is important to guarantee that both time translational invariance and the fluctuation-dissipation theorem hold. If this is the case, it can be shown^[104,120] that the dynamics of the MSD is given by

$$\hat{m}\ddot{\Delta}(t) + \hat{\zeta}\dot{\Delta}(t) = T - \beta \int_0^t du M(t-u) \dot{\Delta}(u). \quad (1.40)$$

Here $\hat{m} = \sigma^2 m / (2d^2)$ and $\hat{\zeta} = \sigma^2 \zeta / (2d^2)$ are the scaled mass and drag coefficients. Such rescaling with d is important in order to keep the variables finite as $d \rightarrow \infty$. For later use, let me now introduce the scaled packing fraction

$$\hat{\varphi} = 2^d \varphi / d = \frac{\tilde{\varphi}}{d}. \quad (1.41)$$

On the other hand, $M(t)$ is a memory kernel that is determined self-consistently by solving the following equations:

$$\hat{m}\ddot{y}(t) + \hat{\zeta}\dot{y}(t) = T - \tilde{v}'[y(t)] - \beta \int_0^t du M(t-u) \dot{y}(u) + \Xi(t) \quad (1.42a)$$

$$M(t-t') = \frac{\hat{\varphi}}{2} \int dy_0 \exp(y_0 - \beta \tilde{v}[y_0]) \langle \tilde{v}'[y(t)] \tilde{v}'[y(t')] \rangle_{\Xi}; \quad (1.42b)$$

where $\tilde{v}(d(x-\sigma)) = v(x)$ is the rescaled potential and \tilde{v}' its derivative. Additionally, the effective noise is correlated as $\langle \Xi(t) \Xi(t') \rangle = 2\hat{\zeta}T\delta(t-t') + M(t-t')$. Note that in this way, instead of dealing with dN degrees of freedom, a single one is needed: $y(t)$. Its dynamics is driven by an effective potential $\tilde{v}[y(t)] - Ty(t)$ and is influenced by the presence of coloured noise, whose memory kernel M is determined by the force-force correlation^[104].

From these equations a series of far-reaching predictions arise. First, there is a density $\hat{\varphi}_d$ at which a *dynamical glass transition* occurs. This is characterized

by the fact that the $t \rightarrow \infty$ limit value of the MSD remains finite. Thus, let¹¹ $\Delta_{\text{EA}} \equiv \lim_{t \rightarrow \infty} \Delta(t) < \infty$. Importantly, MF theory is able to predict both $\hat{\varphi}_d$ and Δ_{EA} . In fact, it is found that $\hat{\varphi}_d = 4.807$. Because $\hat{\varphi} = \tilde{\varphi}/d$, for sufficiently large d this packing fraction is clearly within the interval where Eqs. (1.39) are exact. Besides, the fact that the long time limit of the MSD remains finite implies that a given configuration can never escape the a region whence it departed, and thus remains out of equilibrium. The size of that region is measured by Δ_{EA} .

Second, from the scaled diffusion constant $\hat{D} = \frac{2d^2}{\sigma^2} D$ and viscosity $\hat{\eta} = \frac{2^d V_d}{d^2} \eta$ an expression resembling the Stokes–Einstein relation can be obtained:

$$\left. \begin{aligned} \hat{D} &= \frac{T}{\hat{\zeta} + \beta \int_0^\infty du M(u)} \\ \hat{\eta} &= \beta \hat{\varphi} \int_0^\infty du M(u) \end{aligned} \right\} \implies \hat{D} = \frac{T}{\hat{\zeta} + \hat{\eta}/\hat{\varphi}}. \quad (1.43)$$

In the glass state, the configuration keeps memory of its initial state for arbitrarily long times. In other words, $M_{\text{EA}} \equiv \lim_{t \rightarrow \infty} M(t) > 0$. When this happens, the integral of the expressions above clearly diverges and therefore diffusion is suppressed and the viscosity becomes infinite. This is a sharpened analogy of the drastic increase in the viscosity discussed in Sec. 1.1. Interestingly, notice that even in this case, a sort of Stokes–Einstein relation remains valid, $\hat{D}\hat{\eta} \sim T\hat{\varphi}$.

Another important outcome is that, for $\varphi \lesssim \varphi_d$, the way the MSD approaches and leaves its plateau value is different, namely,

$$\Delta(t) \simeq \Delta_{\text{EA}} - At^{-a}, \quad (1.44a)$$

$$\Delta(t) \simeq \Delta_{\text{EA}} + Bt^b. \quad (1.44b)$$

Note the very close resemblance with Eq. (1.5) for the intermediate scattering function, $F(q, t)$, derived from MCT. Eq. (1.44) also captures the two steps relaxation discussed above, and illustrated in Fig. 1.5a. Moreover, the two *dynamical critical exponents* a and b are related through Eq. (1.6), the same equation as in MCT. Importantly, for HS systems the ratio of Gamma functions defining the dynamical critical exponents can be compute exactly^[112]. It equals 0.707, whence it follows that $a = 0.324 \dots$ and $b = 0.629 \dots$.

Another signature of the transition is obtained by analysing the behaviour of the four-point susceptibility

$$\chi_4(t) = N \left[\overline{\Delta^2(t)} - \overline{\Delta(t)}^2 \right],$$

near φ_d . Theory shows that χ_4 diverges as φ_d is approached from either the glass phase (as $(\varphi - \varphi_d)^{-1/2}$) or from the liquid. In this latter scenario, however, it happens that χ_4 is peaked at τ_α , which also diverges as $\varphi \rightarrow \varphi_d^-$. (Compare with Fig. 1.9.) The form of this divergence can also be deduced analytically, $\tau_\alpha \sim \hat{\eta} \sim (\hat{\varphi}_d - \hat{\varphi})^{-\gamma}$, with $\gamma = 1/(2a) + 1/(2b)$.

In finite d however, *no real transition comes about*. This is caused by the fact that activation processes help to overcome dynamical arrest and, therefore, a configuration eventually escapes the basin whence it departed. That is, metastable states

¹¹The notation, now somewhat spread, honours Sam Edwards and Philip W. Anderson who in 1975 discovered an analogous order parameter in the theory of spin glasses; see [121].

can only persist indefinitely in a large d description, where energy barriers are always high enough to prevent relaxation towards equilibrium. Fortunately, many of the MF predictions remain *qualitatively* valid in low dimensional systems, due to the fact that activation processes are exponentially suppressed when increasing d . Furthermore, several of them have already been put to test in finite d systems since they coincide, or at least mimic, the predictions of MCT, which precedes MF theory. Some of the most important findings of this latter approach are: (i) as d increases, the range of validity of the power-law divergence of τ_α is extended; (ii) estimations of λ constantly approach the predicted value as systems increase their dimensionality; (iii) the product $\hat{D}\hat{\eta}$ seemingly converges to the modified Stokes-Einstein relation for higher d . See [104, Fig. 1] for a summary in $d = 3 - 8$.

1.3.3 Thermodynamics of mean-field glasses

Let us now describe what happens beyond $\hat{\varphi}_d$. Because diffusion is absent, there is no need to explicitly solve the dynamics and a statistical mechanical approach can be used instead. Once again, the relevant order parameter will be the long time limit of the MSD. As we will shortly see, the thermodynamic analysis is based on the fact that *equilibrium* configurations can be used to define, and follow, glassy states. The emphasis should warn about the fact that this is a strong assumption because, by definition, configurations cannot be equilibrated in the glass phase. Indeed, we just discussed that beyond $\hat{\varphi}_d$ the relaxation time diverges. Nevertheless, as a purely theoretical resource we can assume that equilibrated glasses exist (because they were prepared in such state since the beginning of time) and that we are fortunate enough to have some of them at hand for us to study. More realistically, as long as $d < \infty$, τ_α will never be infinite and thus equilibration beyond $\hat{\varphi}_d$ can be realized. As discussed in Sec. 1.2.3, one way of doing so is through MC-swap type of algorithms, which exploit an extended phase space to equilibrate systems at densities that would be otherwise unreachable. An even more important step in the “realism ladder” is accomplished via the ultra-stable glasses generated experimentally through vapour deposition^[122]. Finally, see Sec. 2.2.2 for an example of a MF model in which equilibration can be reached by a technique called *planting*. Hence, we can say that using equilibrated glass configuration is a strong, but plausible assumption.

So, let \vec{s} be one of such equilibrated glasses, with a density $\hat{\varphi}_g > \hat{\varphi}_d$ and pressure p_g ¹². We want to argue that \vec{s} can be used to precisely identify metastable states and to study their thermodynamics. In a first scenario, it will be the initial condition for a trajectory, $\vec{r}(t)$, with $\vec{r}(0) = \vec{s}$ and the same packing fraction, that evolves according to the Langevin dynamics of the liquid, as mentioned above. Given that beyond $\hat{\varphi}_d$ the dynamics is arrested, the distance of \vec{r} from such reference state will be finite, even for long times; *i.e.* $\lim_{t \rightarrow \infty} |\vec{r}(t) - \vec{s}|^2 = \Delta_r < \infty$. In other words, this means that the glass is unable to escape the region of phase space around \vec{s} . Now as a second scenario, let once again $\vec{r}(0) = \vec{s}$ but of a configuration with different density $\varphi' \neq \varphi_g$ and pressure $p' \neq p_g$. Then the long time limit of the distance between such compressed (or decompressed) configurations can be shown to be finite, $\Delta' \equiv$

¹²Note that because it is an equilibrium state, once $\hat{\varphi}_g$ is given, the corresponding pressure is automatically determined through the EOS. Thus, I will avoid making explicit reference to p_g .

$\lim_{\tau \rightarrow \infty} \left(\lim_{t \rightarrow \infty} \Delta[\vec{r}(t + \tau), \vec{r}(t)] \right) < \infty$, where $\Delta[\vec{r}_1, \vec{r}_2] = \frac{d^2}{N\sigma^2} \sum_{i=1}^N |\mathbf{x}_{1,i} - \mathbf{x}_{2,i}|^2$ is the mean squared distance between configurations \vec{r}_1 and \vec{r}_2 . That is, Δ_r measures the change in the configurations with respect to a reference one, while Δ' measures the change between configurations compressed up to $\hat{\varphi}'$. Therefore, a glass state is determined by the values of (Δ', Δ_r) because they specify the set of (isolated¹³) configurations whose mutual distance is Δ' , and are Δ_r apart from the reference state.

Once we have seen how to identify glass states by specifying \vec{s} and $\hat{\varphi}_g$, as well as the characteristic distances Δ_r and Δ' , we can define a probability density constrained by the values of these variables. This density will thus determine how configurations \vec{r} with a density $\hat{\varphi}'$, and a distance Δ_r away from \vec{s} , are sampled. More precisely,

$$\begin{aligned} P(\vec{r}, \hat{\varphi}' | \vec{s}, \hat{\varphi}_g) &= \frac{e^{-\beta \mathcal{H}(\vec{r}; \hat{\varphi}')}}{Z[\Delta_r, \hat{\varphi}' | \vec{s}, \hat{\varphi}_g]} \delta(\Delta_r - \Delta[\vec{r}, \vec{s}]); \\ Z[\Delta_r, \hat{\varphi}' | \vec{s}, \hat{\varphi}_g] &= \int d\vec{r} e^{-\beta \mathcal{H}(\vec{r}; \hat{\varphi}')} \delta(\Delta_r - \Delta[\vec{r}, \vec{s}]) \end{aligned} \quad (1.45)$$

We can then compute the free energy associated to the metastable state selected by \vec{s} from

$$f[\Delta_r, \hat{\varphi}' | \vec{s}, \hat{\varphi}_g] = -\frac{1}{N\beta} \log Z[\Delta_r, \hat{\varphi}' | \vec{s}, \hat{\varphi}_g]. \quad (1.46)$$

Certainly, there is nothing special about \vec{s} since there are many possible equilibrium configurations with the same density $\hat{\varphi}_g$. Hence, the real thermodynamic free energy is given by averaging Eq. (1.46) over all such possible initial configurations, conditioned by the values of $\hat{\varphi}_g$ and Δ_r . Performing such average we obtain the so called Franz–Parisi (FP) potential^[16,35]:

$$V_{\text{FP}}(\Delta_r, \hat{\varphi}' | \hat{\varphi}_g) = \overline{f[\Delta_r, \hat{\varphi}' | \vec{s}, \hat{\varphi}_g]}^{\vec{s}} = -\frac{1}{N\beta} \int d\vec{s} \frac{e^{-\beta \mathcal{H}(\vec{s}, \hat{\varphi}_g)}}{Z[\hat{\varphi}_g]} \log Z[\Delta_r, \hat{\varphi}' | \vec{s}, \hat{\varphi}_g], \quad (1.47)$$

where $Z[\hat{\varphi}_g] = \int d\vec{s} e^{-\beta \mathcal{H}(\vec{s}, \hat{\varphi}_g)}$ is the equilibrium partition function. Despite its rather abstract construction, the FP potential can actually be computed explicitly through the replica method^[78] and thus employed to analyse the stability of the states selected by \vec{s} . In fact, it is this latter step what introduces the dependence on Δ' as distances between replicas are considered. Once such computation has been carried out, the thermodynamics follows in the usual manner. That is, if $V_{\text{FP}}(\Delta_r, \hat{\varphi}' | \hat{\varphi}_g)$ has a minimum at Δ_r , then it is very likely that some metastable states are found within a distance Δ_r from \vec{s} . In contrast, if V_{FP} has no minimum, then states are not trapped around \vec{s} for long times, *i.e.* the system is a liquid.

Another advantage of the FP potential is that the EOS in the glass phase can be obtained via the relation $p = \hat{\varphi}' \frac{\partial(\beta V_{\text{FP}})}{\partial \hat{\varphi}'}$. The first important instance is clearly $\hat{\varphi}' = \hat{\varphi}_g$, in which case the resulting EOS corresponds to the pressure of the *equilibrated* glass, *i.e.* a liquid in equilibrium *beyond* $\hat{\varphi}_d$. In other words, the pressure in this

¹³They are isolated because glassy states belonging to different reference states cannot approach each other. That is, $(\hat{\varphi}_g, p_g)$ define a basin –since it is an equilibrium state– isolated from basins with different densities and pressures.

situation corresponds to the continuation of the liquid's EOS beyond $\hat{\varphi}_d$. This means that, thermodynamically, the liquid is replaced by a set of glassy metastable states, all with density $\hat{\varphi}_g$. The surprising result is that each of such metastable states has a pressure that equals the one the liquid would have at that density. The second relevant scenario is the behaviour of p upon adiabatic compression, for a fixed reference state $\hat{\varphi}_g$. That is, we want to explore what happens as $\hat{\varphi}'$ changes infinitely slowly, from a fixed reference density $\hat{\varphi}_g$. In contrast to the free-volume EOS of a introduced in Eq. (1.32), the MF approach allows to obtain the true thermodynamic EOS. Conceptually, this is evinced by the fact that the FP potential depends explicitly on equilibrium quantities (through $\hat{\varphi}_g$), and on “how faraway from equilibrium” we have driven the system (through Δ_r and $\hat{\varphi}'$). As mentioned above, no such information is contained in Eq. (1.32). Another relevant difference is that, while for finite d systems the curves finish at the equilibrium EOS, the construction obtained through the FP potential shows that the metastable glassy states can be *decompressed* beyond such equilibrium line. In other words, the system presents hysteresis. Each of these (out of equilibrium) glassy EOS ends at an spinodal point. The behaviour of the EOS is depicted in Fig. 1.17. Note that the reference glass states (grey squares) fall on the continuation of the liquid EOS. On the other hand, the dashed lines, above (resp. below) the solid one, correspond to the glass EOS obtained by adiabatic compression (resp. decompression) or, analytically, through the FP potential. The glass phase is delimited on one side by the jamming line at $p = \infty$, but before this boundary is met, the Gardner transition ensues as we will discuss in the next part. Before concluding, note the very close resemblance of this figure with its $d = 3$ counterpart, Fig. 1.16.

To recap, what we have seen in this part is that a glass prepared in equilibrium at a given density $\hat{\varphi}_g > \hat{\varphi}_d$ and pressure p_g , can be used to follow the evolution of glass states up to a target $(\hat{\varphi}', p')$. These two states are connected through (i) the typical distance (Δ_r) between the reference configurations at $(\hat{\varphi}_g, p_g)$ and the target one; and (ii) the typical distance (Δ') between different states with the same target $(\hat{\varphi}', p')$. Such distance is necessarily finite, because diffusion is suppressed in the glass phase, and consequently states can never get too far from their departing point. In other words, metastable glasses remain stuck near the equilibrium state used as reference. This is important because the $t \rightarrow \infty$ limit of dynamical observables, as well as thermodynamical variables, can thus be computed by defining a probability density in phase space, conditioned on the values of $\hat{\varphi}_g, p_g$, and Δ_r . This procedure is what the Franz–Parisi construction is about^[16].

1.3.4 Gardner Transition and fractal free energy landscape

I have just mentioned that the FP potential is computed through the replica method. More precisely, a 1-step Replica Symmetry Breaking (RSB) solution is found near the equilibrium line of glassy states. This solution always exists but, as the systems is further compressed to $\hat{\varphi}_G > \hat{\varphi}_g$, it becomes unstable and is replaced by a fullRSB one. This scenario is akin to the Gardner transition (GT) found in the mean-field p -spin model^[103,124] and constraint satisfaction problems^[125]. Very loosely speaking, a k -step RSB implies that there are k order parameters involved in the computation of the free energy. In turn, in a fullRSB scheme there is a continuum set of order




Figure 1.17. Phase diagram of HS systems in the $d \rightarrow \infty$ limit. Liquid and equilibrated glassy states fall on top of the black solid line. Glasses in equilibrium can be compressed or decompressed *beyond* the equilibrium line through state following (dashed lines). Deep in the glass phase, a Gardner Transition takes place (blue line), rendering the glasses marginally stable. Such marginal phase ends in the jamming line (red). See the close resemblance with Fig. 1.16. Taken from [104]; see also [123].

parameters. For the cases of glasses in particular, this feature applies to the MSD, which still acts as the order parameter. Importantly, this also determines how the states are organized within a basin, as I will explain next.

Let me first consider the case of the FP potential when the 1RSB solution is stable. In such situation, an equilibrium configuration corresponds to a minimum of the energy landscape and thus can be used to identify a basin. Metastable glassy states associated to it are then recognized as the configurations within a typical distance of $\Delta_r = \Delta_1$ from the equilibrium state. Let me now assume that the transition is such that, beyond a certain density, the new stable solution corresponds to a 2-step RSB. In this scenario, the equilibrium configuration would still define a basin, although one with richer structure, since smaller sub-basins would appear within it. Glassy states would now be determined by two order parameters –*i.e.* values of the MSD–, namely $\Delta_{1,2}$ with $\Delta_2 < \Delta_1$. This means that if two metastable states are part of the same sub-basin, then their typical distance would be Δ_2 , while if they belong to different sub-basins they will be apart Δ_1 . This process can clearly be generalized to a k -step RSB solution, and is illustrated in Fig. 1.18 for the case

of $k = 4$. So, for a given level k of RSB, basins are divided into a hierarchy of k sub-basins and glassy states are specified according to the set of order parameters $\{\Delta_k, \Delta_{k-1}, \dots, \Delta_1\}$ as possible values for the MSD. Because these values are such that $\Delta_m < \Delta_{m-1}$, glassy states are said to be ordered hierarchically. These means that if two configurations, \vec{r}_1 and \vec{r}_2 , belong to a common n -level sub-basin, then $\Delta[\vec{r}_1, \vec{r}_2] = \Delta_n < \Delta_{n-1} < \dots < \Delta_1$. Hence, if another configuration \vec{r}_3 is such that $\Delta[\vec{r}_1, \vec{r}_3] = \Delta_m$, *i.e.* shares with \vec{r}_1 a m -level sub-basin, with $m < n$, then it must also happen that $\Delta_m > \Delta_n$. Importantly, note that because of the hierarchical structure, it follows that $\Delta[\vec{r}_2, \vec{r}_3] = \Delta_m$. In general, we can therefore conclude that for any triplet of configurations, $\vec{r}_1, \vec{r}_2, \vec{r}_3$, the following relation holds,

$$\Delta[\vec{r}_1, \vec{r}_3] \leq \max(\Delta[\vec{r}_1, \vec{r}_2], \Delta[\vec{r}_2, \vec{r}_3]). \quad (1.48)$$

In other words, *the metastable glassy states have an ultrametric structure.*




Figure 1.18. MF picture of the fragmentation of the phase space at different densities. In the stable glass phase, *i.e.* for $\hat{\varphi}_d < \hat{\varphi} < \hat{\varphi}_G$, the phase space is clustered into disconnected basins, each of which defines a (metastable) glass. For $\hat{\varphi} > \hat{\varphi}_G$, *i.e.* in the Gardner, marginal phase, each basin is decomposed into a meta-basin. This means that a hierarchy of sub-basins appears (only 4 levels shown), furnishing an ultrametric structure to the phase space. Taken from [104].

A fullRSB scheme is the $k \rightarrow \infty$ generalization of the processes just described. Hence, basins that appeared featureless under the 1RSB solution are transformed, as the GT takes places, into meta-basins; *i.e.* a hierarchy of sub-basins within sub-basins. Such meta-basins are not only hierarchical, but also ultrametric according to Eq. (1.48). Even more, the free energy landscape acquires a *fractal structure*^[111] as a consequence of the infinite sequence of decomposition of a single basin into ever smaller sub-basins.

Naturally, the precise value of $\hat{\varphi}_G$ depends on the density $\hat{\varphi}_g$ at which the equilibrated glass was driven away from equilibrium. This is clear from the (conditional) dependence of V_{FP} on $\hat{\varphi}_g$ in Eq. (1.47). However, any equilibrium glass that is sufficiently compressed will undergo a GT sooner or later, independently of its initial value of $\hat{\varphi}_g$ ¹⁴. This is seen in Fig. 1.17 as the blue line running across the full glass phase. Now, as if all the intriguing features described in the previous paragraphs were not enough, there are a couple of extra properties that come about in the Gardner phase: dynamical heterogeneity and marginal stability. The first one is related to the fact that, beyond $\hat{\varphi}_G$, particles vibrate around non-fixed positions, although

¹⁴At least in the case of HS systems, although the situation can be much more complicated for other potentials or in presence of shear; see Refs. [123, 126, 127].

their displacement is very slow. The amplitudes of such vibrations show correlation over large regions and these glasses are considerably more heterogeneous than when the density is such that $\hat{\varphi}_g < \hat{\varphi} < \hat{\varphi}_G$. This is an important finding because the standard picture of a featureless basin is unable to explain the short timescales of inter-basin relaxation observed deep in the glass phase^[111]. Consider also the insets of Fig. 1.16 where the contrast in the mobility fluctuations below and above the Gardner line is shown.

On the other hand, the marginality of glasses in the Gardner phase is a direct consequence of the fact that at least one eigenvalue of the stability matrix, obtained by computing the replicated entropy, is zero^[118]. Physically, this is revealed by the fact that $\chi_4 = \infty$ throughout this phase: a result that, by the way, is also inconsistent with the simple basin picture^[111]. A diverging four-point susceptibility thus signals the marginal stability of glasses because it implies that they respond abruptly even to infinitesimal perturbations. As we will see next, this has deep consequences for the jamming transition. The criticality of this last transition is characterized by three non-trivial exponents (θ, γ, κ) and is caused by the fact that jammed states are always situated in the marginal phase. It has been a major success of MF theory to be able to predict all three exponents. The role of the first two will be thoroughly discussed in the next section. For the third one, suffice it to say that it equals $\kappa = 1.4157$ and that it controls the rate at which the innermost basin shrinks, as the system is compressed towards its jamming point^[104,111], $\Delta \sim p^{-\kappa}$ ¹⁵. Its value also determines the fractal dimension of the free energy landscape, it is^[111] $2/\kappa \approx 1.413$.

1.4 The jamming transition

We thus finally arrive at the jamming transition. I would like to begin by stressing that the glass and jamming transitions are different. I hope that, by comparing the material from the previous sections to the current one, it becomes clear that the physics of both phenomena markedly differ. When discussing the glass transition above, much emphasis was put on dynamical variables, relaxation times, etc. At jamming, such properties make no sense because there is no dynamics whatsoever¹⁶. Moreover, not even in the vicinity of the jamming transition should the same features arise, if only, because jammed states are contained in the marginal glass phase. That is, the Gardner transition (see Secs. 1.2.3 and 1.3.4) separates glasses and jammed configurations. Naturally, also the variables of interest will be different. For instance, I will discuss about the critical scalings of the energy, pressure, bulk and shear moduli, etc. in Sec. 1.4.2, while in Secs. 1.4.4 and 1.4.5 microscopic structural variables such as contact forces and interparticle gaps will be considered. As I will show in Sec. 1.4.3, these latter quantities are perfectly well defined at jamming and their value is obviously fixed. But how could the analogous structural variables, with the same precise definition, be defined in a glass? In any

¹⁵Note that this scaling contrasts with the ones expected for a normal glass, $\Delta \sim p^{-1}$, or a crystal, $\Delta \sim p^{-2}$.

¹⁶This is not strictly true if a finite temperature is present. However, I will only consider jammed packings at $T = 0$ or as $T \rightarrow 0$.

case, before beginning, I want to explicitly state that many of the critical properties I discuss next apply to what has been termed *isostatic* jamming point.

1.4.1 All the roads lead to jamming

Throughout this chapter, the path I have followed is clearly the one of HS glass formers as they are compressed deep in their glass phase. Yet, the jamming transition would also be a “stop” if the journey had proceeded instead by analysing SS. But not only. As I will discuss next, the jamming transition is shared by a wide variety of systems and under a broad range of circumstances; recall Fig. 1.1. What is more, several of its peculiar critical properties are common to all these systems as well. In other words, jamming is a critical point where physics from diverse systems and models coalesces. In this section, I will consider configurations of both SS and HS, favouring the first type because calculations are more transparent with such model. But the next chapter will be entirely devoted to HS configurations. This should not be taken as an example that jamming is only relevant for spheres. Quite the opposite. But systems made out of spherical particles are a very convenient *minimal* model to explore the non-trivial phenomenology that comes about near and at the jamming transition. As expected, when more complicated shapes or scenarios are considered, the phenomenology is even richer.

As a first standpoint, jamming can be considered as a purely mathematical satisfiability problem^[6,128,129], namely, that of finding possible packings of non-overlapping identical spheres. As simple as it may sound, it encompasses very hard problems, such as, finding the densest possible packing (also called close packing) in an arbitrary dimension d . The archetypical example is the $d = 3$ case, for which Kepler conjectured four centuries ago that the maximum possible density was, precisely, $\varphi_{max,3d} = \varphi_{FCC} = \pi/\sqrt{18} \approx 0.74$. Yet, a complete proof was made available less than 30 year ago by Hales^[83]. Even the “much simpler” analogous problem for disks was completely solved not earlier than 1940^[6]. In this case, the largest packing fraction is attained by placing the disks in a triangular lattice; it yields $\varphi_{max,2d} = \pi/\sqrt{12} \approx 0.907$. Besides these two results, and the trivial one $\varphi_{max,1d} = 1$, the value of φ_{max} is only known for $d = 8$ and $d = 24$ ^[16]. Importantly, in all these cases the optimal value of φ_{max} is realized by using Bravais lattices. In other dimensions the situation is identical: all the maximal density packings known correspond to periodic arrays; see [16, Chp. 8] and references therein for a detailed account. With this in mind, the question “what is the maximum density that *random* packings can attain?” seems more than challenging. Of course, such a question is not well posed until a proper definition of *random* is given. And once again, this is no trivial task. See, *e.g.*, [6] for an excellent review of the many intricacies of defining “disorder metrics”, as well as many other properties that are relevant for analysing jammed configurations, both of spherical and more complex shapes.

On the other hand, the mathematical approach pays little attention to how often a given possible packing is actually found empirically¹⁷. And while it is true

¹⁷I avoided making any reference to what is found *in nature*, because it is clear that packings in $d \geq 4$ would be thus irrelevant. However, considering dimensions other than 2 and 3 is very relevant (and doable) in numerical experiments and in theoretical models. Thus, I will refer to

that regular structures and crystals are common in our daily lives (*e.g.* salt), it is also true that the most common structures are in fact disordered. Therefore, from a physical point of view, analysing packings without any visible periodicity or long-range order is very relevant. Besides, random packings thus produced can also be thought of as instances of a satisfiability problem, in the mathematical sense. It is nonetheless astonishing that many of the features of such *random* close packings (RCP) are very robust and easily reproducible. Here, an intuitive notion of randomness is actually enough. For instance, consider spheres placed via a Poisson process in a fixed volume and at a density large enough that many of them initially overlap. The position of each sphere is independent of the others, so obviously no correlation is expected in their locations. If an interaction energy is introduced between pairs of overlapping spheres, it is clear that by finding the ground state¹⁸ a valid jammed packing is obtained. And because the initial condition was random, so will be the final configuration with very high probability. Experimentally, the situation is, *in principle*, straightforwardly feasible. Consider placing the spherical beads in a large container and tapping or vibrating it until no rearrangements take place. Clearly, no order is expected in the particles positions after such process. And yet, experiments done with hundreds to tens of thousands particles^[130,131], as well as numerical simulations in several dimensions and with different protocols, show that, on average, the properties of the packings thus produced are surprisingly similar. The most important examples are the packing fraction and the average coordination number \bar{z} . In $d = 3$ for example, all the instances have approximately the same value, $\varphi_{RCP} \approx 0.64$ –recall Fig. 1.14 or the experimental results reported in [131]– and $\bar{z} \approx 6$. These findings, as well as several others^[6,84], motivate an *ensemble approach* when considering jamming of RCP configurations. That is, the study of jammed packings of configurations that are most likely to be observed empirically. I will adopt such approach in this work since it is particularly useful for the type of systems I will consider; see however Refs. [132, 133] for a more detailed discussion. Furthermore, given that only disordered configurations will be considered, I will implicitly assume that any jammed packing is an RCP instance of the ensemble and, henceforth, use φ_J instead of φ_{RCP} .

Additionally, the ensemble approach is well suited for studying athermal systems, such as grains, beads, sand, etc. The reason is that it allows to apply, in theory, the full methodology of statistical mechanics. This approach was pioneered by Sam F. Edwards¹⁹, and is nicely reviewed in Ref. [134]. The basic idea is that, within such framework, all the relevant observables can be computed as ensemble averages of an appropriately defined probability measure. As a result, it is possible to define metastability of jammed states (by considering the changes in packing fraction resulting from particle displacements); compute the configurational entropy of packings; as well as to identify a hierarchy of basins connected through their degree of constraint. Interestingly, this method shows a very close correspondence with classifications relying mainly on mathematical aspects, such as the so called *local*,

both results coming from experiments and simulations as *empirical*.

¹⁸In practice, the situation is more complicated because besides an energy minimization, the particles' size should be simultaneously decreased. Nevertheless, several algorithms exist that can accomplish such task in configurations of tens or even hundreds of thousands of spheres.

¹⁹Yes, the same Sam F. Edwards of the Edwards–Anderson order parameter of Sec. 1.3.2.

collective, and *strictly* categories of jamming^[6]. Even more, it can be shown that the ground states of jammed granular materials are in an analogous fullRSB phase, just like glass formers (see Sec. 1.3.4), thus demonstrating that what connects jamming of these two type of systems is not just a mere analogy, but their underlying thermodynamical properties. Similarly, several of the scalings present at the jamming transition (see Sec. 1.4.2) are also found in soft materials such as foams, colloids, and emulsions^[4].




Figure 1.19. Liu–Nagel phase diagram. It sketches that the jamming transition can be identified with a critical point –common for several systems and reachable by tuning different physical variables– the so called “J-point”. Note however that the distinction between glasses and jammed states is not clear from this particular plot. Taken from [4], after the proposal in [135, 136].

Jamming thus seems to be rather pervasive, on the one hand, while it consistently exhibits the same features, on the other. This latter point will be discussed in the following sections, but before doing so, it is worth asking why the jamming phenomenology is so persistent in so many diverse systems. In their seminal works^[135–137], Liu, Nagel and co-workers provided an answer of far-reaching consequences: jamming is ubiquitous because it is a common critical point (also termed *J-point*) of all this kind of systems. This behaviour can be illustrated using the so called *Liu–Nagel plots*, as the one shown in Fig. 1.19. From this perspective it becomes clear that, although different systems reach jamming by tuning different physical variables, several properties near and at the onset of jamming are shared by

all of them. It is also worth mentioning that equivalent results can also be derived from MF theory, as in, *e.g.*, [126, 127].

To finish this part, let me mention that jamming also occurs outside the physics realm. Two prominent examples are the perceptron^[138–140] and the SAT-UNSAT transition of continuous constraint satisfaction problems^[141,142]. These problems exhibit all the critical properties expected from a MF perspective and are thus important theoretical models to analyse. Likewise, the Mari–Kurchan (MK) model^[143] is another MF-type of model that is very useful as a reference for comparison. It consists in particles interacting through a randomly shifted distance, $D(\mathbf{r}_i, \mathbf{r}_j) = |\mathbf{r}_i - \mathbf{r}_j + \zeta \mathbf{A}_{ij}|$, where \mathbf{A}_{ij} is a fixed random vector, different for each pair of particles. ζ is a parameter that allows to interpolate between fully MF ($\zeta \rightarrow \infty$) and finite d ($\zeta = 0$) behaviours. The MK model in $d = 2$ and $d = 3$ has been used to validate several of the predictions of MG theory^[109], in particular the GT^[110]. Some other aspects about the liquid and glass phases of this model will be given in Sec. 2.2, while in Chapter 4 I analyse in detail its microstructure at jamming, so for the moment I just mention that it has also been used to test some of the jamming criticality properties^[144]. Finally, jamming is also present in other MF models^[145], neural networks^[146–148], and statistical inference^[149]. This shows the relevance of jamming as a current and relevant research topic.

1.4.2 Criticality approaching the jamming transition

To analyse the critical behaviour of (isostatic) jammed configurations, I will first consider the case of soft spheres that interact only when there is an overlap between them. It is convenient to define the dimensionless gaps,

$$h_{ij}(\vec{\mathbf{r}}_i, \vec{\mathbf{r}}_j) = \frac{|\mathbf{r}_i - \mathbf{r}_j|}{\sigma_{ij}} - 1, \quad (1.49)$$

where $\sigma_{ij} = \frac{1}{2}(\sigma_i + \sigma_j)$ is the sum of the two particles' radii. That is, for the time being, I will consider the more general case of polydisperse configurations for reasons that will become clear shortly. In this way, an overlap corresponds to a negative gap. The contact potential modelling the interaction between particles can then simply be expressed as a function of the gaps. Following the literature^[16,137,150] I will consider a contact potential of the form

$$v(h) = \frac{\epsilon}{\alpha} |h|^\alpha \Theta(-h), \quad (1.50)$$

where Θ is the Heaviside step function, ϵ is a constant that fixes the energy scale, while α is a parameter that can be used to vary the “softness” of the spheres.

Once the potential has been defined, a valid jammed packing can be obtained with a suitable energy minimization procedure. Recently, the FIRE algorithm^[151] has been extensively used for this purpose, although several options are possible. Note however that, independently of the algorithm employed, jammed configurations correspond to an energy minimum with $T = 0$. That is, in scenarios where thermal noise is included to enhance the exploration of the configuration space and thus find lower minima, it is important to have in mind that a proper critical jammed configuration will only be achieved in the $T \rightarrow 0$ limit²⁰.

²⁰ At least in the sense that I will discuss here. But see Ref. [152] for a study with $T > 0$.

At this point, I want to bring attention to a peculiar feature of the jamming transition. So far, finding jammed packings seems only a minimization problem, similar to others found in statistical physics. Consider for the sake of simplicity the Ising model with no external field and equal couplings, $J_{ij} = J$. Finding its ground state in the $T \rightarrow 0$ limit, will always result in well defined physical variables, such as $\langle |m| \rangle = 1$, and analogously for its energy, heat capacity, etc. In contrast, for jammed packings the situation is very different, although this has been only implicit in the discussion up to this point. To make the matters clearer, consider the phase diagrams of Figs. 1.16 and 1.17. Note that in those figures, jammed states are not concentrated around a point, but actually lie along a *line*. This means that the jamming density can take any value within an interval²¹. In fact, the value of φ_J depends on both the protocol used and the configuration itself. In the case of monodisperse systems, this effect is considerably suppressed, given that particles are identical and consequently the *J-line* is shrunk to a narrow interval. Nonetheless, the dependence of the protocol remains. This is evinced by the results of Fig. 1.14, in which all the MD simulations departed from the same initial condition and only the compression rate (κ) was varied. It is evident that although the final packing fractions are similar, using a slower compression consistently produces denser packings. Similarly, in the case of the polydisperse configurations, letting the system remain longer in equilibrium allows to reach higher values of φ_J . Fortunately, these observations also provide the answer to understand the strange dispersion in the values of φ_J : that jamming is an *out of equilibrium* transition^[153]. In other words, the jammed state reached by whatever methods depends fundamentally on both (i) the state at which the system abandons equilibrium and (ii) how it evolved afterwards. Consequently, also some of the observables at jamming will show such dependence. Now, in the thermodynamic limit and with an infinitesimally slow compression or decompression protocol, φ_J has a well defined value, that nonetheless depends on the equilibrated glass state whence it departed (see previous section), *i.e.* $\varphi_J = \varphi_J(\varphi_g)$. When analysing jamming criticality, these considerations are important because we are interested in how thermodynamical variables scale as a function of $(\varphi - \varphi_J)$. So, in principle, we should use the thermodynamic limit value of the jamming density, $\varphi_J(\varphi_g)$, to properly observed the power-law scalings. In practice, however, one should keep in mind that the value of φ_J will be unique of each configuration.

Now, jamming criticality can be intuitively justified as follows. If SS interact through a potential like Eq. (1.50), then it is clear that for $\varphi < \varphi_J$ the energy is identically zero, while for densities larger than φ_J it is strictly positive. Thus, the energy must be singular at φ_J . The scaling of the energy and other thermodynamical

²¹Clearly, there is a lower bound for such interval, since at very low densities not even a glass can be formed. In contrast, deriving an upper bound is not an easy task. However, it can be shown that, the maximal density an amorphous packing can have is of order $\sim d \log(d2^{-d})$; see [16, Chp. 8].

variables are summarized in the following expressions^[4,136,137]:

$$e(\varphi) \sim (\varphi - \varphi_J)^\alpha, \quad (1.51a)$$

$$P(\varphi) = \rho^2 \frac{de}{d\rho} \propto \varphi^2 \frac{de}{d\varphi} \sim (\varphi - \varphi_J)^{\alpha-1}, \quad (1.51b)$$

$$G(\varphi) \sim (\varphi - \varphi_J)^{\alpha-3/2}, \quad (1.51c)$$

$$B(\varphi) - B_0(\alpha) \sim (\varphi - \varphi_J)^{\alpha-2}, \quad (1.51d)$$

$$\bar{z}(\varphi) - z_c \sim (\varphi - \varphi_J)^{1/2}. \quad (1.51e)$$

In the last equation, z_c is the average number of contacts observed for a particular configuration with φ_J . Finite size effects causes this variable to exhibit small sample to sample fluctuations. But in the thermodynamic limit its value is perfectly well defined and independently of the jamming protocol; it is $z_c = 2d$. In Sec. 1.4.3 I will show that this value is a consequence of the mechanical stability of the packings. On the other hand, the term $B_0(\alpha)$ in Eq. (1.51d) is such that $B_0(\alpha) > 0$ if $\alpha \leq 2$ and zero otherwise. Let me consider briefly the harmonic potential, $\alpha = 2$. Such potential is, quite reasonably, an important case and it is useful in modelling emulsions, colloids, etc.^[4] The scalings of Eqs. (1.51d) and (1.51c) imply that a jammed packing of harmonic SS is unable to withstand shear forces ($G \rightarrow 0$), and yet remains rigid ($B \rightarrow B_0 > 0$, *i.e.* it supports compression). In fact, the difference of the critical exponents of the shear and bulk moduli is a consequence of G being more susceptible to non-affine effects. In contrast, B is determined through the effective stiffness (k) of the spheres, which is usually determined by affine deformations; see [4] for a detailed discussion. Moreover, if the rigidity of a configuration is measured by normalizing with respect to k , then the behaviour at φ_J resembles a first order phase transition, given that B/k remains finite for any $\varphi > \varphi_J$ but vanishes for smaller densities. The same jump is observed in the coordination number since $\langle z \rangle = 0$, for $\varphi < \varphi_J$.

At any rate, the power-law scalings summarized in Eqs. (1.51) have been extensively verified in several works^[4,136,137,154,155] and, surprisingly, they are unaffected by dimensionality or polydispersity of the configurations. Additionally, the finite size effects expected from these relations have been confirmed^[154] and a Widom-like scaling ansatz has been derived for all of them^[156]. Furthermore, various studies have identified correlation lengths associated to the characteristic length scales of vibrational response to perturbations^[157,158], the fluctuations in the number of contacts^[144,159], and the fluctuations of particle mobility^[51], all of which diverge at the jamming point.

In the case of HS similar relations hold. However, because the interaction potential is always zero, the energy identically vanish for any $\varphi < \varphi_J$. Likewise, before the jamming point is reached, the configurations are not rigid^[160] nor do they sustain shears, hence both bulk and shear moduli vanish. Therefore, the relevant scaling relations only concern the reduced pressure and the intensive entropy^[84]:

$$p(\varphi) \sim (\varphi_J - \varphi)^{-1}, \quad (1.52a)$$

$$s(\varphi) \sim \log(\varphi_J - \varphi). \quad (1.52b)$$

Indeed, the second of these expressions can easily be obtained from the first one given that for HS (i) $p = \beta \rho \frac{\partial f}{\partial \rho}$ and (ii) $f = -Ts$. It is then straightforward to

obtain $p = -\varphi \frac{\partial s}{\partial \varphi}$ whence (1.52b) follows by integration. It is also worth noting that Eq. (1.52a) is the same as the free-volume of a HS glass, cf. Eq. (1.32) and Fig. (1.15). This implies that the algebraic dependence of p on $\varphi_J - \varphi$ is valid throughout the glass phase and not only near jamming.

Another important result is that effects of finite temperature can also be included near jamming^[51,152,161] and thus join the behaviours on both sides of φ_J . For instance, Eqs. (1.51b) and (1.52a) can be combined in a single scaling relation, that depends on both φ and T :

$$p(\varphi, T) = T^{-1/\alpha} \mathcal{P}\left(T^{-1/\alpha}(\varphi - \varphi_J)\right), \quad \mathcal{P}(x) \sim \begin{cases} |x|^{-1} & \text{for } x \rightarrow -\infty, \\ x^{\alpha-1} & \text{for } x \rightarrow \infty. \end{cases} \quad (1.53)$$

Scalings for other variables are listed in Ref. [152]. All of these results thus provide some of the strongest evidence in support of the critical nature of the jamming transition.

The picture that unfolds by putting together the exact MF description (Sec. 1.3), the scalings of Eqs. (1.51)-(1.53), and the robustness of numerical experiments for several dimensions and different protocols^[84,94,136,137,153] suggests that the jamming transition of spherical particles defines a universality class. Even more, it is a rather broad class that encompasses all of the different systems mentioned at the end of Sec. 1.4.1. Nevertheless, we also know that if inherent order is enforced – *e.g.* by using minimally polydisperse crystals^[162–164] – or particles are not spherical^[165,166] universality is broken. Properly demarcating such universality class is therefore an open problem that falls outside the scope of this work. Nevertheless, I will briefly come back to the subject in Chp. 4, which deals with a systematic assessment of the jamming criticality of the *microscopic variables* introduced in Sec. 1.4.4.

1.4.3 Network of contacts and isostaticity

In this section, I will analyse the properties of the network of contacts (NC) formed *exactly at a jammed state*. This is a fundamental section for several of the results that will be presented in the following chapters. From the analysis of the NC the most important properties that I will show are:

1. A jammed state reached with a given interaction potential is an equally valid jammed configuration with any other potential.
2. In the case of spherical particles, the minimal requirement for a jammed packing to be stable is that it has a *single state of self stress*. This means that the number of contacts (N_c) is equal to the number of degrees of freedom *plus one*; *i.e.* $N_c = N_{dof} + 1$.
3. Under such conditions, the forces exerted by pairs of particles in contact are determined solely by the particles positions.
4. The structure of the NC determines the Hessian of the system at jamming, and vice-versa.

Several of these results have been developed in a series of important works, *e.g.* [56, 57, 167–169]. While in [150] they have been put together and generalized. Thus, I will mostly follow the discussion presented there, as well as in [16, Sec. 9.2].

To show the properties mentioned above, let me begin by analysing the Hessian, H of a configuration *above* jamming, *i.e.* I will study a configuration of SS with $\varphi > \varphi_J$. To do so, I will use the following notation: $[ij]$ with $i < j$ will denote the contact between particles i and j and it will also be used as a contact index, *i.e.* $[ij] \in \{1, \dots, N_c\}$; $\mathbf{n}_{ij} = \frac{\mathbf{r}_i - \mathbf{r}_j}{|\mathbf{r}_i - \mathbf{r}_j|}$ is a unity vector pointing from particle j towards particle i (clearly $\mathbf{n}_{ij} = -\mathbf{n}_{ji}$); $f_{ij} = f_{ji}$ is the magnitude of the force of contact $[ij]$. Additionally, let $\mathcal{C} = \{[ij]_c\}_{c=1}^{N_c}$ be the set of contacts. With this notation, the total potential energy can be expressed as

$$V(\vec{\mathbf{r}}) = \sum_{c \in \mathcal{C}} v(h_c(\vec{\mathbf{r}})),$$

where, for the time being, I am assuming that the gap function can depend on the full configurational vector $\vec{\mathbf{r}} = \{\mathbf{r}_i\}_{i=1}^N$ and not just on the pair of particles involved in the contact c . Assuming a general power-law contact potential like the one of Eq. (1.50) we have that

$$\frac{\partial V(\vec{\mathbf{r}})}{\partial r_i^\mu} = \sum_{c \in \mathcal{C}} v'(h_c) \frac{\partial h_c}{\partial r_i^\mu} = -\epsilon \sum_{c=1}^{N_c} |h_c|^{\alpha-1} \frac{\partial h_c}{\partial r_i^\mu}; \quad \mu = 1, \dots, d. \quad (1.54)$$

The minus sign in the rightmost term comes from the assumption $h_c < 0$, given that $\varphi > \varphi_J$, so the Θ function is consequently omitted. Physically, it represents the fact that a force should increase in magnitude as the corresponding gap becomes more negative. Similarly, the Hessian is a $dN \times dN$ symmetric matrix $\frac{\partial^2 V(\vec{\mathbf{r}})}{\partial \delta \vec{\mathbf{r}}^2}$ whose entries are given by

$$H_{i,\mu}^{j,\nu} \equiv \epsilon \sum_{c=1}^{N_c} \left[(\alpha - 1) |h_c|^{\alpha-2} \frac{\partial h_c}{\partial r_i^\mu} \frac{\partial h_c}{\partial r_j^\nu} - |h_c|^{\alpha-1} \frac{\partial^2 h_c}{\partial r_i^\mu \partial r_j^\nu} \right]. \quad (1.55)$$

Now, let $\vec{\mathbf{r}}_0$ be an energy minimum, although not necessarily a jammed configuration. In that case, Eq. (1.54) vanishes and the energy cost of a small perturbation $\delta \vec{\mathbf{r}}$ around $\vec{\mathbf{r}}_0$ is well approximated by

$$\begin{aligned} \delta V(\vec{\mathbf{r}}) &= V(\vec{\mathbf{r}}_0 + \delta \vec{\mathbf{r}}) - V(\vec{\mathbf{r}}_0) \approx \frac{1}{2} \delta \vec{\mathbf{r}} \cdot H \cdot \delta \vec{\mathbf{r}} \\ &= \frac{\epsilon}{2} \sum_{c=1}^{N_c} \left[(\alpha - 1) |h_c|^{\alpha-2} \underbrace{\left(\delta \vec{\mathbf{r}} \cdot \frac{\partial h_c}{\partial \delta \vec{\mathbf{r}}} \right)^2}_{\text{harmonic}} - \underbrace{\delta \vec{\mathbf{r}} \cdot \frac{\partial^2 h_c}{\partial \delta \vec{\mathbf{r}}^2} \cdot \delta \vec{\mathbf{r}}}_{\text{prestress}} |h_c|^{\alpha-1} \right]. \end{aligned} \quad (1.56)$$

Note that the harmonic term is a stabilizing contribution to the perturbation, since each contact suppresses perturbations along the direction of its gradient. Thus, of the possible dN directions in which a perturbation could move away from \mathbf{r}_0 , N_c are blocked by the contacts. To proceed with this argument, I will restrict to the case where the system is subject to periodic boundary conditions, which reduces

in d the number of degrees of freedom, N_{dof} (see below a detailed argument), *i.e.* $N_{dof} = dN - d$. This means that $N_{dof} - N_c$ directions remain along which the system could be perturbed at zero energy cost. In other words, there are $\max\{N_{dof} - N_c, 0\}$ non-trivial soft or zero modes, and d trivial ones, related to uniform displacements of the configuration.

Because I assumed that $\vec{\mathbf{r}}_0$ is a minimum H must be a positive semi-definite matrix and $\delta V > 0$. To ensure stability, some constraints must hold for the prestress term. There are two possibilities, either $\frac{\partial^2 h_c}{\partial \delta \vec{\mathbf{r}}^2}$ is a positive-definite matrix, or its not. In the first case, it will have a *destabilising* effect in H . In particular, the prestress term is negative^[16] along the $N_{dof} - N_c$ (possible) zero modes of the harmonic contribution. Hence, it must be the case that $N_c \geq N_{dof}$ in order to guarantee stability. Otherwise $\vec{\mathbf{r}}_0$ would not be a minimum, in contradiction with the initial assumption. The second possibility in which $\frac{\partial^2 h_c}{\partial \delta \vec{\mathbf{r}}^2}$ is not positive-definite is more complicated because it can happen that the prestress term can actually stabilize some of the zero modes from the harmonic part. In this case, the situation $N_c < N_{dof}$ is in principle allowed, although it is system dependent. When this happens, the system is said to be *hypostatic*, the condition $N_c = N_{dof}$ corresponds to an *isostatic* configuration, while if $N_c > N_{dof}$ the system is called *hyperstatic*.

I will show now that for particles whose gaps are given by Eq. (1.49) the energy minima can never be hypostatic. In particular, this is true for SS, because symmetry constrains pairwise gaps to be only a function of the centres' distance. Hence, letting $h_{[kl]}$ be of form of Eq. (1.49), it is straightforward to derive that

$$\begin{aligned} \frac{\partial h_{[kl]}}{\partial r_i^\mu} &= \frac{1}{\sigma} \frac{r_k^\mu - r_l^\mu}{|\mathbf{r}_k - \mathbf{r}_l|} (\delta_{ik} - \delta_{il}), \\ \frac{\partial^2 h_{[kl]}}{\partial r_j^\nu \partial r_i^\mu} &= \frac{(\delta_{ik} - \delta_{il})(\delta_{jk} - \delta_{jl})}{\sigma} \left[\frac{\delta_{\mu\nu}}{|\mathbf{r}_k - \mathbf{r}_l|} - \frac{(r_k^\mu - r_l^\mu)(r_k^\nu - r_l^\nu)}{|\mathbf{r}_k - \mathbf{r}_l|^3} \right]. \end{aligned}$$

And thus, plugging these expressions into Eq. (1.56) and performing the dot products gives,

$$\begin{aligned} \delta V(\vec{\mathbf{r}}) &\approx \frac{\epsilon}{2\sigma} \sum_{[ij] \in \mathcal{C}} \left[\frac{(\alpha - 1)}{\sigma} |h_{[ij]}|^{\alpha-2} (\delta \mathbf{r}_{ij} \cdot \mathbf{n}_{ij})^2 - |h_{[ij]}|^{\alpha-1} \frac{|\delta \mathbf{r}_{ij}|^2 - (\delta \mathbf{r}_{ij} \cdot \mathbf{n}_{ij})^2}{r_{ij}} \right] \\ &= \frac{\epsilon}{2\sigma} \sum_{[ij] \in \mathcal{C}} \left[\frac{(\alpha - 1)}{\sigma} |h_{[ij]}|^{\alpha-2} (\delta \mathbf{r}_{ij} \cdot \mathbf{n}_{ij})^2 - |h_{[ij]}|^{\alpha-1} \frac{|\delta \mathbf{r}_{ij}^\perp|^2}{r_{ij}} \right], \end{aligned} \tag{1.57}$$

where $\delta \mathbf{r}_{ij} = \delta \mathbf{r}_i - \delta \mathbf{r}_j$, $r_{ij} = |\mathbf{r}_{i,0} - \mathbf{r}_{j,0}|$, and $\delta \mathbf{r}_{ij}^\perp$ is the component of $\delta \mathbf{r}_{ij}$ orthogonal to \mathbf{n}_{ij} . Because the second term inside the square brackets is the sum of non-positive terms the prestress matrix is thus negative semi-definite. And according to the discussion above, this implies that $N_c \geq N_{dof}$. Therefore, energy minima of spherical particles are at least isostatic.

Note that so far these results are valid for any type of interaction potential of the form (1.50) and for any state that defines an energy minimum. If, besides being an equilibrium point, $\vec{\mathbf{r}}_0 = \vec{\mathbf{r}}_J$ is also a jammed state, further important results can be derived. First of all, note that at jamming the overlaps go to zero, and as long

as $\alpha > 1$ ²², the prestress term vanishes, proportionally to the pressure, and quicker than the harmonic one. We can thus define the contact forces at jamming as,

$$f_{ij} = f_{ji} = \lim_{h_{[ij]} \rightarrow 0^-} - \left| \frac{\partial v'(h_{[ij]})}{\partial \mathbf{r}_{ij}} \right| = \lim_{h_{[ij]} \rightarrow 0^-} \frac{\epsilon}{\sigma} |h_{ij}|^{\alpha-1} > 0. \quad (1.58)$$

The last inequality follows from the fact that at the jamming transitions configurations are *rigid*, which means that there is a non-zero force between particles in contact. This a subtle issue, so let me analyse it more carefully. Recall that in the previous section I mentioned that the value of the bulk modulus at jamming is, in general, finite. This was a consequence of the spheres' stiffness, k , remaining finite at jamming. Now, k is proportional to the curvature of the energy minimum at $\vec{\mathbf{r}}_J$, which must also be positive due to stability. Explicitly, $k = \frac{\epsilon}{\sigma} |h|^{\alpha-2} > 0$, or equivalently, $f_{ij} = k_{ij}^{\frac{\alpha-1}{\alpha-2}} > 0$. Moreover, the finiteness of the contact forces can be deduced from mechanical stability considerations as I will now show. A very important outcome of the proof is that contact forces can indeed be obtained only from the particles' position $\vec{\mathbf{r}}$, without making any reference to the form of the interaction potential.

By definition, the jammed state corresponds to a state of mechanical equilibrium, and therefore the total force acting on particle i is given by

$$\mathbf{F}_i + \sum_{j \in \partial i} \mathbf{n}_{ij} f_{ij} = 0$$

where ∂i is the set of particles in contact with i and \mathbf{F} is any possible external force. If we denote by $\underline{\mathbf{f}} = \{f_{[ij]}\}_{[ij] \in \mathcal{C}}$ the N_c -dimensional vector of contact forces, the equilibrium condition can be expressed for the complete configuration as $\vec{\mathbf{F}} + \mathcal{S}^T \underline{\mathbf{f}} = 0$, where \mathcal{S} is the $N_c \times dN$ *contact matrix*, whose entries read

$$\mathcal{S}_{[jk]}^{i,\mu} = (\delta_{ij} - \delta_{ik}) n_{jk}^\mu, \quad \text{or} \quad \mathcal{S}_{[jk]}^i = (\delta_{ij} - \delta_{ik}) \mathbf{n}_{jk}. \quad (1.59)$$

Note that the same expression, except for a proportionality factor of $\frac{\epsilon}{\sigma}$, would follow by plugging in the expression given for $\frac{\partial h_{jk}}{\partial \mathbf{r}_i^\mu}$ into Eq. (1.54), with the definition of f_{ij} according to Eq. (1.58).

From now on, I will focus on systems with periodic boundaries and with no external forces acting on them, but see Ref. [150, 167] for the general analysis. The second of these conditions makes the mechanical equilibrium to be succinctly expressed as

$$\mathcal{S}^T \underline{\mathbf{f}} = 0 \quad \Longleftrightarrow \quad \sum_{[jk]} \mathcal{S}_{[jk]}^{i,\mu} f_{jk} = 0, \quad \forall i = 1, \dots, N \text{ and } \forall \mu = 1, \dots, d. \quad (1.60)$$

The last expression corresponds to a system of dN homogeneous linear equations in N_c unknowns. But not all of them are independent because the total force in

²²The case $\alpha = 1$ is very interesting, since it corresponds to a singular potential. It can be shown however, that all the results derived so far also apply. See Refs. [139, 170].

the system must be zero, due to absence of external forces. Because this condition holds independently for each spatial dimension we have that,

$$\sum_{i=1}^N \sum_{[jk]} S_{[jk]}^{i,\mu} f_{jk} = 0, \quad \forall \mu = 1 \dots d.$$

Hence, there are only $d(N - 1)$ independent equations. But this is precisely the number of (relevant) degrees of freedom! To see why, we just need to consider that the periodic boundary conditions make the system invariant under translations, so d degrees of freedom should be subtracted from the configurational dimension, thus leading to

$$N_{dof} = d(N - 1). \quad (1.61)$$

Equivalently, the same conclusion is obtained by noting that a uniform translation $\delta \mathbf{r} = \mathbf{a}$ implies that $\delta \mathbf{r}_{ij} = 0$, in which case Eq. (1.57) vanishes identically. This means that uniform translations are zero modes of the Hessian, thus effectively reducing the degrees of freedom as in Eq. (1.61). On the other hand, we already know that a condition valid for all energy minima is that $N_c \geq N_{dof}$. In particular, for a jammed state the mechanical equilibrium condition, Eq. (1.60), defines a set of homogeneous equations that therefore has $\max\{N_c - d(N - 1), 0\}$ non-zero solutions. But experience teaches us that jammed packings do exist, so there is at least one non-trivial solution to Eq. (1.60). Therefore the number of contacts must be, at least,

$$N_{1SS} \equiv N_{dof} + 1 = d(N - 1) + 1. \quad (1.62)$$

This condition corresponds to the so called *single state of self-stress* (1SS), although in literature it is more commonly termed *isostaticity*. Nevertheless, I will use the former name and reserve the latter for the case $N_c = N_{dof}$. This shows that, whenever $N_c \geq N_{1SS}$ it necessarily happens that $f_{ij} > 0 \forall [ij] \in \mathcal{C}$.

Before moving forward, I would like to mention that a precise enumeration of the degrees of freedom is of utmost importance to correctly analyse jamming critical properties, because several of them manifest when $N_c = N_{1SS}$. In fact, at the beginning of this section I mentioned that the results discussed here were valid for isostatic configurations, but there I implicitly assumed that $N = \infty$. However, the correct statement is that the number of constraints should be 1 above isostaticity. (Of course, in the thermodynamic limit such distinction is irrelevant.) At any rate, to properly count the degrees of freedom it is mandatory to consider the symmetries as extra constraints of the system, which begs the question: should central potentials not further reduce N_{dof} ? The answer is: it depends on the boundary conditions. In fact, for closed systems, N_{dof} is further reduced by $\frac{d(d-1)}{2}$. However, in systems with periodic boundary conditions angular momentum is *not* conserved^[171], and such extra symmetries need not to be included. I am of the mind that these considerations are commonly overlooked, or simply stated without further thought, sometimes leading to apparently different criteria for N_{1ss} , *e.g.* [56, 154, 172–174]. (However, once the inherent symmetries of each case are accounted for, the criteria are consistent.) So I wanted to spend some time addressing this issue, at least for the case of periodic boundaries, for which Eq. (1.62) provides the correct value. A careful discussion in other scenarios can be found in Refs. [154, 172].

Now, despite its simplicity, Eq. (1.62) is an important result. So let me discuss it further. First, because each contact involves two particles, the average number of contacts is

$$\bar{z}_J = \frac{2N_{1SS}}{N} = 2d + \frac{1-d}{N} \xrightarrow{N \rightarrow \infty} 2d = z_c. \quad (1.63)$$

The critical value of the average coordination number in Eq. (1.51e) is thus recovered. As discussed above, this is the smallest number of contacts required for a configuration to be stable. This type of counting argument was first derived by Maxwell, and applies to the stability analysis of a large class of problems, like networks of harmonic springs^[174]. Second, mechanical stability also dictates that each particle must have $d + 1$ contacts, otherwise at least one degree of freedom remains unconstrained. Particles that do not fulfil this bound are called *rattlers* and their contribution to the NC should be removed since they do not contribute to the rigidity of the packing. Nevertheless, they are scarce in the vast majority of the configurations, and their quantity decreases with dimensionality. Along these lines, it is an interesting fact that the coordination number is a self-averaging quantity, which means that at jamming, each particle very likely has $2d$ contacts. Moreover, the fluctuations around this value are suppressed rapidly as d increases^[150]. As a final remark, I should mention that these last expressions make clear that in the thermodynamic limit there is no distinction between isostaticity and 1SS. Nevertheless, such difference will be relevant for the results of Chp. 4, in which the effects of finite size in jammed packings are studied. At any rate, the role of the additional contact of 1SS configurations, with respect to isostaticity, cannot be overstated because it implies that $\underline{\mathbf{f}}$ is the unique zero mode of \mathcal{S} , which only depends on $\vec{\mathbf{r}}_J$. Consequently, *the forces magnitudes are entirely determined by the particles' position.*

Notice that in deriving these results I have not made any assumption about the potential between spheres, thus showing that once a jammed configuration is reached, it is a valid jammed state for any potential. Nevertheless, a direct connection with the energetic approach is possible by noticing that the Hessian can be expressed in terms of a rescaled contact matrix. For instance, in the most usual case of an harmonic interaction, from Eq. (1.55) and the expression for $\frac{\partial^2 h_{[kl]}}{\partial r_j^\nu \partial r_i^\mu}$, it readily follows that

$$H_{i,\mu}^{j,\nu} = \delta_{ij} \sum_{k \in \partial i} n_{ik}^\mu n_{ik}^\nu - n_{ij}^\mu n_{ij}^\nu \delta([ij]) = (\mathcal{S}^T \mathcal{S})_{i,\mu}^{j,\nu}$$

In this last expression, I have used $\delta([ij])$ to denote a function that is 1 if particles i and j are in contact, and 0 otherwise. Even more, for a jammed configuration $\vec{\mathbf{r}}_J$ obtained with a general potential, we can construct \mathcal{S} and then solve Eq. (1.60) for $\underline{\mathbf{f}}$. Concomitantly, we can compute the effective stiffness $k_{ij} = \frac{(\alpha-1)\epsilon}{\sigma} f_{ij}^{\frac{\alpha-2}{\alpha-1}}$ and thus define

$$\tilde{\mathcal{S}}_{[jk]}^{i,\mu} = k_{jk}^{1/2} (\delta_{ij} - \delta_{ik}) n_{jk}^\mu = k_{jk}^{1/2} \mathcal{S}_{[jk]}^{i,\mu}. \quad (1.64)$$

Then, the relation

$$H = \tilde{\mathcal{S}}^T \tilde{\mathcal{S}} \quad (1.65)$$

for expressing the Hessian in terms of the (rescaled) contact matrix is still fulfilled. Notice that once the Hessian is obtained, we can obtain all the normal modes and the spectrum of the system; see Sec. 1.4.5.1. This result will be used in Chp. 3.

As a final comment, it should be mentioned that the contact matrix, the Hessian, and the closely related symmetric matrix

$$\mathcal{N} = \mathcal{S}\mathcal{S}^T, \quad (1.66)$$

are very useful to analyse the mechanical properties of jammed configurations. It is easy to show that H and \mathcal{N} have the same spectra. Besides, because \mathcal{N} and \mathcal{S} have the same zero modes, another way to calculate \underline{f} is to extract the eigenvector associated with the zero eigenvalue of \mathcal{N} . Numerically, this is often more convenient^[150]. Moreover, all these matrices provide information about the floppy modes, response to a dipolar force opening a contact, etc. A very complete discussion can be found in [150, 167–169].

1.4.4 Criticality of the microscopic structure *at* the jamming point

Jamming exhibits another, rather peculiar, type of criticality. It involves the contact forces of the NC and the gaps between near contacts. Both of these variables also exhibit a power law scaling once the configuration *has reached* its jamming point. Importantly, MF theory predicts the exponents of such scalings to be universal. Note that this is different from the criticality discussed in Sec. 1.4.2, because in that case, the scaling relations (1.51) and (1.52) were derived as the jammed state is *approached*. In other words, the usual type of scalings studied in critical phenomena^[175,176]. In contrast, the results of this part concerns the situation where $\varphi = \varphi_J$ and a jammed configuration \vec{r}_J has already been realized. In this situation, it has been found that the distributions of interparticle gaps and contact forces are power-laws with non-trivial critical exponents.

I begin with the gaps, defined according to Eq. (1.49). At variance with the previous section, only positive gaps will be considered because the aforementioned scaling pertains particles that are almost touching. Now, gaps values are randomly distributed because packings are disordered, but theoretical predictions^[111] state that the distribution of small enough gaps should scale as

$$g(h) \sim h^{-\gamma}, \quad \text{with } \gamma = 0.41269\dots \quad (1.67)$$

(The distribution g just introduced should not be confused with the radial distribution function, $g(r)$, defined in Eq. (1.17), although they are closely related; see Eq. (2.21) in the next chapter.) Current numerical evidence supports the universality hypothesis of γ , since the probability density function (pdf) of Eq. (1.67) has been verified in different dimensions and using different protocols^[94,111,153].

Similarly, the distribution of small contact forces is predicted to scale as, $p(f) \sim f^\theta$, but a strong dependence of θ on dimensionality and jamming protocol was initially reported, in apparent contradiction with the theoretical expectation. This paradox was resolved by recognizing that two different types of forces contribute in this regime^[150,169]. This is a consequence of the two distinct responses that can occur when opening the contact between a pair of particles: (i) a localized

rearrangement of neighbouring particles; or (ii) a displacement field that extends over the whole configuration, without decaying with distance; see Fig. 1.20. The former is associated with a buckling motion, and hence remains localized. The latter is associated with a correlation length of the same order as the system size, and hence is a clear example of the criticality of jammed packings. Considering these two types of forces separately yields two power laws with different exponents,

$$p_\ell(f) \sim f^{\theta_\ell}, \quad \text{with } \theta_\ell \simeq 0.17, \quad (1.68a)$$

$$p_e(f) \sim f^{\theta_e}, \quad \text{with } \theta_e = 0.42311 \dots; \quad (1.68b)$$

for localized and extended excitations, respectively.

The ability of MF theory^[16,104,111] to predict the non-trivial values of γ and θ_e is a major success. No MF prediction, however, exists for θ_ℓ , because bucklers are an intrinsically low-dimensional feature^[150], and are therefore subdominant in the $d \rightarrow \infty$ description. In any case, the MF prediction for θ_e and γ is based on the solutions of the fullRSB differential equations, in the limit of $\Delta \rightarrow 0$. The calculations are far from trivial, but when an scaling ansatz is proposed for the form of the solutions, it can be shown that conditions leading to g and p_e are indeed universal. The details of the calculations are presented in [118, 123, 141] and in abbreviated manner in [16, Secs. 9.2-9.3]. On the other hand, besides the theoretical predictions several works provide numerical results supporting the relations of Eqs. (1.67) and (1.68) in several dimensions and in different scenarios; see [94, 118, 138–141, 150, 153, 167, 169, 170].

1.4.5 Marginal stability

The last aspect I would like to discuss concerns the marginal stability of jammed packings. The mathematical derivation follows mostly Ref. [168], which offers a more general result than the one presented in a previous, very nice work^[177]. Additionally, the role of marginal stability in jammed packings and in other common disordered systems has been reviewed recently in [178].

The argument goes as follows. Let us focus on a given contact $\lambda \equiv [kl]$ and open it an amount s . We are interested in analysing the response of the system to such contact opening, while leaving the rest of them unchanged. The rearrangement of particles caused by opening λ is called a *floppy mode*. Notice that if the floppy mode associated to λ leads to a denser packing, then the original configuration would be unstable. Denoting the displacement experienced by particle i after opening λ as $\delta \mathbf{r}_i^{(\lambda)}$, and $\delta \mathbf{r}_{ij}^{(\lambda)} = \delta \mathbf{r}_i^{(\lambda)} - \delta \mathbf{r}_j^{(\lambda)}$ being the relative displacements, it can be shown that the equation such displacements must obey, up to the leading nonlinear contribution, reads:

$$\delta \mathbf{r}_{ij}^{(\lambda)} \cdot \mathbf{n}_{ij} + \frac{(\delta \mathbf{r}_{ij}^{(\lambda)} \cdot \mathbf{n}_{ij}^\perp)^2}{2r_{ij}} + \mathcal{O}(s^3) = s\delta_{[ij],\lambda}, \quad (1.69)$$

where \mathbf{n}_{ij}^\perp is a vector orthogonal to \mathbf{n}_{ij} . Multiplying this last expression by f_{ij} and summing over all the contacts we obtain the virtual work theorem:

$$p \delta V^{(\lambda)} \approx s f_\lambda - N \frac{c_\lambda \langle f \rangle}{\sigma} s^2. \quad (1.70)$$

In this equation $\delta V^{(\lambda)}$ is the volume change caused by the force unbalance when the contact λ was opened, while c_λ is a dimensionless number, containing the total contribution of the $\left(\delta \mathbf{r}_{ij}^{(\lambda)} \cdot \mathbf{n}_{ij}^\perp\right)^2$ terms, in the limit $s \rightarrow 0$. Physically, it is related to the magnitude of the displacement field averaged over the whole system and, as argued in Ref. [177], its value is determined by average force in the system, and is of order s^2 .

Notice that the nonlinear term is important, otherwise Eq. (1.70) would imply that any floppy mode would increase the volume and therefore packings would be inherently stable. In contrast, when such term is taken into account it is easy to see that there exists s^* such that for

$$s < s^* \equiv \frac{\sigma f_\lambda}{N c_\lambda \langle f \rangle}, \quad (1.71)$$

no packing with smaller volume can be formed. To prevent such situation, a new contact must be formed elsewhere in the configuration before λ is opened as much as s^* . Denoting by s^\dagger the distance at which such new contact is formed, a configuration will be stable as long as $s^\dagger < s^*$. Otherwise, a denser packing could be formed before another contact is closed, signalling that the configuration was initially unstable. The conditions for stability are thus derived by (i) estimating s^\dagger using the gaps distribution, and (ii) computing s^* from the distributions of the contact forces. It is convenient to consider a *stability index*

$$\Sigma_\lambda \equiv s^*/s^\dagger = \frac{\sigma f_\lambda}{N s^\dagger c_\lambda \langle f \rangle}, \quad (1.72)$$

so that Σ_λ greater (smaller) than 1 corresponds to a stable (unstable) packing. It is then reasonable that we centre our analysis on the smallest forces, since they are likeliest to produce destabilising floppy modes.

As argued in Ref. [167], there are two (approximately independent) possibilities for having a small force; whence let us write $f_\lambda \sim b_\lambda W_\lambda$, where the meaning of these two factors will be explain next. The first possibility is that a contact force between a pair of particles is weakly coupled to the rest of the forces exerted on them by the rest of their neighbours. For instance, if all but one contacts of a particle are nearly coplanar, in which case the remaining force should be orthogonal to that plane and very small. This would lead to a bucking motion, whose displacement field thus remains mostly localized, as mentioned in the previous part; see also the analysis of Ref. [150]. If we let b_λ measure the median of the displacements, then it is clear that for these forces $b_\lambda \ll 1$. More precisely, it can be shown^[167] that a localized floppy mode is such that $b_\lambda \ll 1/\sqrt{N}$. I should emphasize that localized modes produce displacements of considerable magnitude, but only near the originating contact. It is farther away that the displacement field is negligible. This justifies using the median and not the mean for properly identifying them, given that it would be dominated by the large displacements near the opened contact. In contrast, it could also happen that a contact force is small because it is weakly coupled to the *bulk* –not just forces of other neighbour contacts– and thus produces a very small, but *uniform and non-decaying*, floppy mode. W_λ is thus a measure of the bulk coupling to contact λ . Calculating it is not straightforward, and its definition

even depends on the boundary conditions of the system. But in any case, it can be estimated through the response of the full configuration to an extended floppy mode; the details are given in [167]. The two type of floppy modes are illustrated in Fig. 1.20, for a configuration of 1024 disks with periodic boundary conditions. Note that both values of f_λ (lower panels) are similar, but the displacement fields they produce are very different.

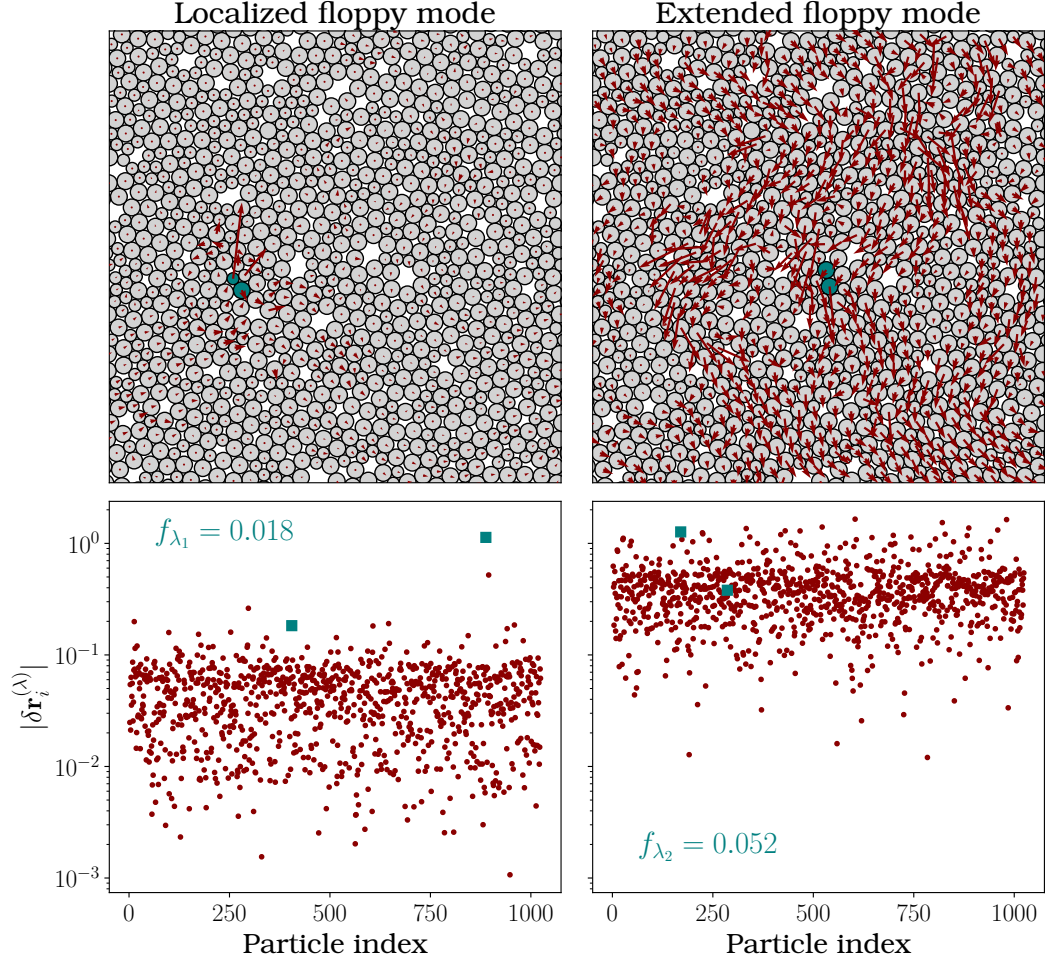


Figure 1.20. Upper panels: Localized (left) and extended (right) floppy modes obtained by opening the contact involving highlighted particles. The displacement field is obtained following the methods developed in Ref. [150]. Lower panels: Norm of each particle's displacement in the floppy mode. The data of the particles involved in the contact openings are similarly highlighted. Notice that both forces have a similar magnitude, but their corresponding floppy modes are clearly different.

Another important assumption for $\{b_\lambda\}_{\lambda \in \mathcal{C}}$ and $\{W_\lambda\}_{\lambda \in \mathcal{C}}$ is that, besides being statistically independent, they are power-law distributed as

$$\begin{aligned} p(b) &\sim b^{\theta_b}, \\ p(W) &\sim W^{\theta_e}. \end{aligned} \tag{1.73}$$

So the same distributions as the bucklers and extended forces in Eq. (1.68). In fact,

if MF theory is not considered, the distributions of contact forces can only be found empirically and the factorization $f_\lambda \sim b_\lambda W_\lambda$ is thus justified *a posteriori*. Besides, given that MF calculations^[118] make no prediction for localized modes, in practice the effects of $\{b_\lambda\}_{\lambda \in \mathcal{C}}$ should be checked after the full NC is known. In any case, the factorization assumed for the forces also allows to split c into two contributions as

$$c_\lambda \sim b_\lambda^2 + \frac{1}{N},$$

in such a way that for localized modes $b \ll 1/\sqrt{N}$ we have $c \sim 1/N$, while for extended modes $c_\lambda \sim b_\lambda^2$.

Let us now estimate the maximal displacement s^\dagger from the distribution of gaps. If h_{min} is the smallest gap in the configuration and gaps are distributed according to Eq. (1.67), then

$$\frac{1}{\sigma^{1-\gamma}} \int_0^{h_{min}} dh h^{-\gamma} = \left(\frac{h_{min}}{\sigma} \right)^{1-\gamma} \sim \frac{1}{N} \quad \implies \quad h_{min} \sim \sigma N^{-\frac{1}{1-\gamma}}.$$

On the other hand, from the arguments above we have that $s^\dagger \sim h_{min}/b_\lambda \sim \frac{\sigma}{b_\lambda} N^{-\frac{1}{1-\gamma}}$. The last step needed to estimate the stability index, defined in Eq. (1.72), is to compute $f_\lambda/\langle f \rangle$. To do so, let me first consider that opening f_λ causes an extended response in the system, so (1) the smallness of f_λ is caused by a small value of W_λ and (2) $c_\lambda \sim b_\lambda^2$. Assuming the extreme case of which f_λ is of the order of the smallest force in the system we have that

$$\int_0^{f_{min}} \frac{df f^{\theta_e}}{\langle f \rangle^{1+\theta_e}} \sim \int_0^{W_{min}} \frac{dW W^{\theta_e}}{\langle W \rangle^{1+\theta_e}} \sim 1/N \quad \implies \quad \frac{W_{min}}{\langle W \rangle} \sim \frac{f_\lambda}{\langle f \rangle} \sim N^{-\frac{1}{1+\theta_e}},$$

where the $\langle \bullet \rangle$ terms are introduced for normalization purposes. Plugging these results into Eq. (1.72) we obtain

$$\Sigma_\lambda \sim \frac{f_{min}}{N h_{min}} \sim \frac{N^{\frac{\gamma}{1-\gamma}}}{N^{\frac{1}{1+\theta_e}}}.$$

Analogously for a localized mode, if f_λ is small due to a very small value of b_λ , a worst case scenario analysis leads in this case to

$$\Sigma_\lambda \sim \frac{f_{min}^2}{\frac{1}{N} h_{min}} \sim \frac{N^{\frac{1}{1-\gamma}}}{N^{\frac{2}{1+\theta_e}}}.$$

From these two last expressions, it is easy to verify that if we require $\Sigma_\lambda > 1$ to guarantee stability, then the following pair of inequalities must hold:

$$\gamma \geq \frac{1}{2 + \theta_e}, \tag{1.74a}$$

$$\gamma \geq \frac{1 - \theta_e}{2}. \tag{1.74b}$$

Values derived from MF theory, reported in Eqs. (1.67) and (1.68b), as well as estimations from simulations suggest that Eq. (1.74a) is actually an equality.

The same is valid for (1.74b), although in this case the value of $\theta_\ell \approx 0.17$ is obtained directly from numerical results. Therefore, given that both inequalities are saturated, we conclude that jammed packings are *marginally* stable; see [178] for a detailed discussion. Importantly, this result is consistent with the MF picture mentioned above, which always locate the jamming line within the Gardner phase. Recall that, as discussed in Sec. 1.3.4, the meta-basin structure acquired by the free energy landscape is a consequence of the (fullRSB) decomposition of a basin into an infinite hierarchy of smaller sub-basins, where each minimum identifies a marginal state. Marginality is observed, for example, in the divergence of χ_4 throughout the Gardner phase. Jammed states thus inherits the marginality property, but it is not longer visible through a correlation of dynamical variables, since no dynamics takes place whatsoever. In turn, it is a feature exhibited by the microscopic structure of the network of contacts, as I have just argued. More precisely, a mathematical analysis of the fullRSB equations that fix the values of the exponents $(\theta_e, \gamma, \kappa)$ show that they are equivalent to the marginal stability condition; see [111, 118].

1.4.5.1 Density of states at jamming

As a final topic, I will briefly consider the vibrational density of states (DOS), $D(\omega)$. In amorphous solids, the DOS has been extensively studied (see Refs. [53, 55–59, 179–184] for a *partial* selection) because the normal modes and their spectrum control several physical quantities of interest, such as specific heat, transport properties, response to shear on compression. For a well defined Hamiltonian, the normal modes are obtained easily by diagonalizing the corresponding Hessian. The spectrum (*i.e.* the set of frequencies $\{\omega\}_{i=1}^{dN}$) is simply related to the corresponding eigenvalues ($\{\lambda_i\}_{i=1}^{dN}$) of the Hessian by $\omega_i = \sqrt{\lambda_i}$. Whenever $\omega_i = 0$, the corresponding mode is termed *soft mode* to emphasize that particles can be perturbed at no energy cost.

It has been known for some time now that in glasses and amorphous solids $D(\omega)$ differs from the expected behaviour of a crystal, namely, the Debye scaling at low frequencies, $D_{\text{crystal}} \sim \omega^{d-1}$, with d the dimensionality of the system. However, the actual form of $D(\omega)$ has been amply debated. For instance, MF analytical calculations^[53] (in the perceptron) and effective medium theory^[183] predict that the DOS present different behaviours parametrized by three characteristic frequencies, $(\omega_0, \omega_{\text{max}}, \omega^*)$, and has approximately the following form^[16]:

$$D(\omega) = \begin{cases} 0, & \omega \notin [\omega_0, \omega_{\text{max}}], \\ \left(\frac{\omega}{\omega^*}\right)^2, & \omega_0 \ll \omega \ll \omega^*, \\ \text{constant}, & \omega^* \ll \omega \ll \omega_{\text{max}}. \end{cases} \quad (1.75)$$

The role of ω_0 and ω_{max} is thus to define the domain of the DOS, while ω^* determines a characteristic frequency at which the distribution detaches from the constant value plateau. Now, a fullRSB calculation shows that $\omega_0 = 0$, implying a *gapless* spectrum. Moreover, numerical studies^[56, 58, 179] in both the over- and under-jammed^[59] phases have shown that $\omega^* \sim (\varphi - \varphi_J)^{1/2}$. Therefore, *at* jamming, $D(\omega)$ does *not* decay to zero, even as $\omega \rightarrow 0$.

Such behaviour is actually another manifestation of the marginal stability of jammed packings, because a non-vanishing DOS as $\omega \rightarrow 0$ implies that excitations of arbitrary low frequencies are always present –apart from the trivial d zero modes from uniform translations. Hence, jammed configurations have an extensive response to external perturbations. See Ref. [53] for a detailed calculation of $D(\omega)$ and how marginal *mechanical* stability –derived from the 1SS criterion– and marginal *landscape* stability –characterized by the soft modes of $D(\omega)$ – are related.

For completeness, I should also mention that another regime, proportional to $(\frac{\omega}{\omega^*})^4$ has been identified far away from jamming, and related to localized modes. However, the extent of this additional regime has not been fully established due to its dependence on d ^[59,181]. On the other hand, no theoretical argument has been found to account for such behaviour, given that it mainly is a low dimensional feature, and thus out of the scope of MF theory.

Chapter 2

Linear Programming algorithm to generate jammed configurations

In this chapter I will describe in detail a method we developed to reach the jamming point of a system through an iteration of Linear Programming (LP) optimizations. LP-based algorithms have been used before to generate jammed packings (see for instance [185, 186], or [187] for an iterative approach similar to ours). Nonetheless, our implementation provides a simple and robust mechanism to reach the jamming point, as well as accessing the critical structural variables, *i.e.* the network of contact forces and of interparticle gaps. It has been described briefly in Refs. [1, 3], but here I will present a more careful analysis, putting special attention in demonstrating that it produces stable packings. The code with the algorithmic implementation will be made publicly available shortly.

I should anticipate that the applicability of LP algorithms is limited by the system's size, given that each particle introduces several constraints to be considered for the optimization. Nonetheless, we were able to generate configurations of more than 10^4 particles with a generic optimization software. Despite such disadvantage, LP jamming algorithms are a reliable way of producing packings of hard spheres (HS) because the non-overlapping constraints are explicitly introduced in the optimization problem. Carefully exploring HS systems as jamming ensues is important because, as discussed in the previous chapter, most of the numerical studies have focused on soft sphere systems. Therefore, developing algorithms capable of precisely probing the jamming criticality as approached from the unjammed phase is still relevant.

This chapter is structured as follows. First, in Sec. 2.1 I discuss the connections of jamming with (linear) optimization theory and I show that jammed packings can be produced by iteratively maximising the density of a configuration, which acts as the objective function. To continue, I enumerate few results from Convex optimization theory (Sec. 2.1.1) to show afterwards that, upon convergence, our iterative LP algorithm (iLP) generates well defined and mechanically stable jammed packings (2.1.2). With the [explicit algorithmic description](#) of our method, this section contains the main theoretical results of this chapter. As exemplified by Figs. 2.3

and 3.1, our algorithm is able to inflate a configuration from a low density phase all the way to its jamming point. However, using a configuration with a very high pressure –and thus closer to jamming– as initial condition drastically enhances its performance. Sec. 2.2 is devoted to describing an efficient molecular dynamics (MD) compression protocol that allows to produce such high p configurations. In this section I also present some results of the glassy phase of the Mari–Kurchan model, for which I developed the corresponding MD code to implement the compression. By complementing our iLP algorithm with the MD compression, we are able to obtain, rather quickly, *typical* jammed packings; that is, without any signatures of partial ordering. Later, in Sec. 2.3 I present a recent characterization of our iLP algorithm. I first consider the influence on the jammed state of the final pressure and compression rate of the MD part (Sec. 2.3.1) and show that the results agree with the fractal energy landscape framework^[111] mentioned in Sec. 1.3.4; then I describe the phenomenology of the MD+iLP route to jamming (Sec. 2.3.2) ; and finish by analysing the computational performance of our algorithm as a function of the system’s size (Sec. 2.3.3). Sec. 2.4 closes the chapter presenting some conclusions and topics to study in future works.

2.1 Jamming as a Linear Programming problem

The methodology I will present in this section has been previously described in [1, 3, 188]. Nevertheless, here I will present a complete account of how the iLP algorithm works and show that, upon convergence, it reproduces the conditions for mechanical stability discussed in Sec. 1.4.3; specifically, see Eq. (1.60). As mentioned in that section, jammed packings must have a single state of self-stress (1SS), or, in other words, have an extra contact with respect to isostaticity. Naturally, this property is also present in the packings produced with our iLP algorithm. The idea behind it is fairly simple. Suppose that we begin with a given configuration of N infinitely hard (hyper) spheres (HS) in d dimensions, whose positions¹ $\vec{\mathbf{r}}$ and diameters $\vec{\sigma}$ are known. Because HS can never overlap, the configuration’s packing fraction necessarily satisfies $\varphi \leq \varphi_J$, where φ_J is the jamming density that will be eventually reached. Without loss of generality I will assume the strict inequality, otherwise the jammed state has already been realized.

One way to approach the jamming point is to inflate all the particles by a factor $\sqrt{\Gamma} > 1$. (The reason for including the square root will be evident soon.) Because particles cannot overlap, Γ is limited by the distance between the closest pair of particles, $\sqrt{\Gamma} = \min_{ij} \frac{1}{\sigma_{ij}} |\mathbf{r}_i - \mathbf{r}_j|$. But such process would leave the vast majority of particles without neighbours in contact. A more clever strategy would be to let particles move in such a way that a larger value of $\sqrt{\Gamma}$ was possible. Even more, we could look for the optimal displacement $\vec{\mathbf{s}}$ such that, when particles are rearranged as $\vec{\mathbf{r}} \rightarrow \vec{\mathbf{r}} + \vec{\mathbf{s}}$, Γ is maximal. (Because the square root is a monotonic increasing function, if Γ reaches its maximum, so does $\sqrt{\Gamma}$.) Mathematically, this is equivalent

¹I will follow the notation introduced in Sec. 1.2.1.

to the following optimization problem (OP):

$$\max \Gamma \quad (2.1a)$$

$$|\mathbf{r}_i + \mathbf{s}_i - (\mathbf{r}_j + \mathbf{s}_j)|^2 \geq \Gamma \sigma_{ij}^2 \quad \forall i \neq j = 1, \dots, N. \quad (2.1b)$$

In this OP the *design variables*, i.e. the variables on which the objective function and constraints depend, are $\vec{\mathbf{s}}$ and Γ . Additionally, any vector $\vec{\mathbf{s}}$ that satisfies all the constraints is called *feasible*. Note that by including the square root in the inflating factor, this OP depends on Γ only linearly. At this point, this is no big advantage, but it will prove useful later. Expectedly, in such general setup this problem is incredibly hard to solve. Moreover, even if we consider a simplified version where all particles have the same size, we know from Sec. 1.4.1 that the optimal solution is, very likely, a crystal. But this is a configuration we want to avoid in the jamming configurations we are looking for². To overcome both difficulties and be able to obtain useful jammed packings from this OP some manipulations are needed.

The first thing to mention is that (2.1) effectively poses the generation of a jammed packing as an instance of an OP. Its objective function, Γ , is trivial. So what makes the problem hard to solve are the non-overlapping constraints, Eq. (2.1b), because they are *non-convex* inequalities. To see why, let $\mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j$ be the vector pointing from \mathbf{r}_j towards \mathbf{r}_i , with an analogous definition for \mathbf{s}_{ij} and $\sigma_{ij} = \frac{\sigma_i + \sigma_j}{2}$ the sum of the radii. For a given pair of particles, inequality (2.1b) thus reads

$$|\mathbf{s}_{ij}|^2 + 2\mathbf{r}_{ij} \cdot \mathbf{s}_{ij} - \Gamma \sigma_{ij}^2 = \vec{\mathbf{s}} \cdot \mathbf{Q} \cdot \vec{\mathbf{s}} + 2\vec{\mathbf{b}} \cdot \vec{\mathbf{s}} - \Gamma \sigma_{ij}^2 \geq -|\mathbf{r}_{ij}|^2. \quad (2.2)$$


Where $\vec{\mathbf{b}} = (\mathbf{0}, \dots, \mathbf{r}_{ij}, \mathbf{0}, \dots, -\mathbf{r}_{ij}, \mathbf{0}, \dots, \mathbf{0})$ is a dN dimensional vector whose only non-zero entries are the d components associated to particles i and j . Similarly, \mathbf{Q} is a $dN \times dN$ symmetric matrix with entries

$$Q_{k,\mu}^{l,\nu} = \begin{cases} \delta_{\mu,\nu} & \text{if } i = k \text{ or } j = l, \\ -\delta_{\mu,\nu} & \text{if } j = k \text{ or } i = l, \\ 0 & \text{otherwise.} \end{cases}$$

In other words, it is a matrix of $N \times N$ blocks, each of dimension $d \times d$. The i th and j th ones along the diagonal are \mathbb{I}_d (the identity matrix in d dimensions), while the ones in the i th and j th rows and column are $-\mathbb{I}_d$. All the remaining entries are zero. Now, linear terms are always convex, so what determines the convexity of the constraints is the quadratic term, $\vec{\mathbf{s}} \cdot \mathbf{Q} \cdot \vec{\mathbf{s}}$. But \mathbf{Q} is a positive semi-definite matrix, since for any vector $\vec{\mathbf{x}} \in \mathbb{R}^{dN}$, $\vec{\mathbf{x}} \cdot \mathbf{Q} \cdot \vec{\mathbf{x}} = |\mathbf{x}_i - \mathbf{x}_j|^2 \geq 0$. Therefore the right hand of Eq. (2.2) is a convex function of the design variables. It is thus anticlimactic that the *non-convexity* is rooted in the inequality sign. More precisely, the *feasible domain* of the OP (2.1) is a non-convex set, as a consequence of the inequality sign of its constraints. Fig. 2.1a illustrates this feature in the case of 4 fixed (coloured) disks constraining the available location of another movable (black) particle, all of different sizes. In this diagram, the set of constraints (2.2) corresponds to the

²This might not be the case in general. For instance, analysing jammed packings with some inherent ordering is interesting; see [187, 189]. Nevertheless, for this thesis only amorphous systems will be relevant.

white space inside the four circles, and the feasible region –*i.e.* the space where the centre of the black disk is allowed to be– is shaded in grey. It is clearly non-convex, since there are many line segments that pass through the space excluded by the constraints.



(a) Exact constraints of the OP (2.1).
The feasible region is *non-convex*.

(b) Linearised (solid) and exact (dashed) constraints.

Figure 2.1. Exact (left) and linearised (right) non-overlapping constraints according to Eqs. (2.1) and (2.3), respectively. Each particle induces on the black disk’s centre a constraint identified by its colour, with the resulting feasible region shaded in grey. Only in the second case it is convex. Coloured particles are assumed to be fixed, so only the constraints affecting the black one are shown.

Fortunately, we can sidestep the difficulty of solving a non-convex OP by assuming that the given configuration $\vec{\mathbf{r}}$ and its diameters $\vec{\sigma}$ are sufficiently close to the jamming point (a method to achieve this will be discussed in Sec. 2.2 below). In this situation, the quadratic term can be neglected and thus the original OP can be transformed into a Linear Programming (LP) problem,

$$\max \Gamma \quad (2.3a)$$

$$2\mathbf{r}_{ij} \cdot \mathbf{s}_{ij} - \Gamma \sigma_{ij}^2 \geq -|\mathbf{r}_{ij}|^2, \quad \forall i \neq j = 1, \dots, N. \quad (2.3b)$$

As mentioned few lines above, OP’s involving only linear functions of the design variables are always convex, for which many efficient solving algorithms exist; see, *e.g.* [190–193]. Moreover, because we have neglected a non-negative term, any feasible solution of (2.3) also satisfies the constraints of the original OP (2.1). Geometrically, each of the linear non-overlapping constraints defines the half-space of an hyper-plane perpendicular to the segment joining the centres of the corresponding pair of particles. Hence, the complete set of constraints forms a polytope as sketched in Fig. 2.1b. Given that each bisection of space by a hyper-plane is itself convex, and that the intersection of convex sets is also convex^[190], the resulting polytope is guaranteed to be convex.

On the other hand, neglecting the quadratic term causes any solution of the LP problem to be sub-optimal with respect to the original OP. In other words, an optimal solution (\vec{s}^*, Γ^*) of (2.3) is (very likely) not a jammed configuration. Yet, the new configuration thus produced $\vec{r}' = \vec{r} + \vec{s}$, with diameters $\vec{\sigma}' = \sqrt{\Gamma^*} \vec{\sigma}$, is definitely closer to the jamming point. Then, if this new configuration is used to define another instance of the LP (2.3), the second optimal solution will be even closer to jamming. If this process is repeated, and it converges, then the resulting configuration will correspond to a proper jammed state. This is, essentially, the method we developed to generate jammed packings. I mention in passing that if the initial condition is random, then this algorithm is also able to overcome, in most of the cases, the issue of possibly ending up with an ordered packing. This is so because the linear constraints impose much tighter bounds to the possible displacements of each particle, even at relatively low densities; cf. the feasible regions of both figures in Fig. 2.1. This means that at each LP step, only local changes are performed, which are consequently unable to rearrange a disordered configuration into a regular one. For this to happen, major collective rearrangements are needed, which are significantly disfavoured in our algorithm. Put it simply, it means that by linearising the non-overlapping constraints, our algorithm suppresses jumps across (entropic) barriers. The only exception we have found so far are configurations of monodisperse disks, because they are particularly prone to partially crystallize. But it is widely known that crystallization in monodisperse disks is very hard to avoid, so it is not an issue specific to our algorithm. In contrast, in $3 \leq d \leq 5$ we have verified that with this method crystallization does not occur if a random initial condition is used.

In summary, our algorithm works as follows (see however Alg. 1 for more detailed description):

1. Choose an initial random configuration \vec{r} , and its set of diameters $\vec{\sigma}$.
2. Use $(\vec{r}, \vec{\sigma})$ to construct an instance of the LP problem (2.3).
3. Obtain the optimal solution of the LP problem (2.3), corresponding to displacements \vec{s}^* and growing factor Γ^* .
4. **If** $\vec{s}^* = \vec{0}$ and $\Gamma^* = 1$ **break** (convergence was reached).
5. **else**
 - (a) Update the particles position according to $\vec{r} \rightarrow \vec{r} + \vec{s}^*$ and $\vec{\sigma} \rightarrow \sqrt{\Gamma^*} \vec{\sigma}$.
 - (b) Repeat steps 2-5.

However, the considerations of the previous paragraph hint that the solution thus obtained is not globally optimal, since only a “local search” has been performed. Nevertheless, from the discussion of the previous chapter, we know that jammed states correspond to local minima of the free energy landscape (FEL), because they are obtained from the metastable glass states. Even more, recall that because of their abundance, local minima actually dominate the probability measure. So the local optimality obtained through our method is an expected property from any other technique and, moreover, it is a desired feature for sampling correctly the FEL of HS. Note also that in our algorithm convergence is reached when the only

solution to (2.3) is $\bar{s}^* = 0$ and $\Gamma^* = 1$, and this implies that particles cannot be further inflated because all the degrees of freedom are blocked. Heuristically, this is precisely what defines a jammed state according to the discussion of Sec. 1.4. Furthermore, upon convergence, the final configuration obtained from a sequence of LP optimizations of (2.3) coincides with a (possibly local) optimal solution of (2.1).

All of these arguments hint that our algorithm is capable of producing jammed packings, although so far I have not shown that such packings are stable, nor that they fulfil the mechanical equilibrium condition. To do so, I will introduce some mathematical results of convex optimization theory. Then in Sec. 2.1.2 I will show that from such results the full network of contacts can be obtained and the stability of packings can be established (in the sense of the results of Sec. 1.4.3). Once again, linear constraints will play a crucial role because their associated dual variables correspond to the contact forces in the jammed configuration.

2.1.1 Some results of Convex Optimization

Here I will mention (without proving) some mathematical results needed to analyse the solution obtain from the iLP algorithm described above. I will follow Ref. [190], a very good reference on convex optimization problems, both from a theoretical and practical point of view.

Let me first introduce some basic terminology. Consider the following general OP:

$$\min F_0(\vec{x}) \quad (2.4a)$$

$$\text{subject to } F_a(\vec{x}) \leq 0, \quad a = 1, \dots, m; \quad (2.4b)$$

$$H_a(\vec{x}) = 0, \quad a = 1, \dots, q; \quad (2.4c)$$

where $\vec{x} \in \mathbb{R}^n$. The function F_0 is called the *objective function*, while F_a and H_a are the constraint equations. The domain of the OP (2.4) is $\mathcal{D} = \cap_{a=0}^m \text{dom } F_a \cap \cap_{a=1}^q \text{dom } H_a \subseteq \mathbb{R}^n$. A point $\vec{x} \in \mathcal{D}$ is said to be *feasible* if it satisfies the equality and inequality constraints, while the OP is called feasible if there is at least one feasible point. The *feasible set* or *feasible domain* is the collection of all possible feasible points. Additionally, the optimal value of the OP (2.4), denoted p^* , is defined as

$$p^* = \inf \{ F_0(\vec{x}) \mid F_a(\vec{x}) \leq 0, \ a = 1, \dots, m; \ H_a(\vec{x}) = 0, \ a = 1, \dots, q \}. \quad (2.5)$$

It is customary to assume that p^* is within the *closed* interval $p^* \in [-\infty, \infty]$, because $p^* = -\infty$ means that F_0 is an unbounded objective function, while $p^* = \infty$ for an infeasible OP. Analogously, a point \vec{x}^* is termed *optimal point* if (i) \vec{x}^* is a feasible point; and (ii) $F_0(\vec{x}^*) = p^*$. Finally, given a feasible point \vec{x} , each of the inequality constraints $\{F_a\}_{a=1}^m$ can be classified as: *active* at \vec{x} if $F_a(\vec{x}) = 0$, or *inactive* at \vec{x} if $F_a(\vec{x}) < 0$. A constraint is *redundant* if the feasible set is unchanged when it is deleted.

All of this applies to any OP, but I will be interested only in *Convex* optimization problems (COP). But first, some extra terminology. A set C is convex if for any $\vec{x}_1, \vec{x}_2 \in C$ and $0 \leq \theta \leq 1$, then $\theta \vec{x}_1 + (1 - \theta) \vec{x}_2 \in C$. Similarly, a function

$F : \mathbb{R}^n \rightarrow \mathbb{R}$ is said to be *convex* if its domain, $\text{dom } F$ is convex and, if for all $\vec{x}, \vec{y} \in \text{dom } F$ and $0 \leq \theta \leq 1$ we have that

$$F(\theta\vec{x} + (1 - \theta)\vec{y}) \leq \theta F(\vec{x}) + (1 - \theta)F(\vec{y}). \quad (2.6)$$

If only the inequality holds in this last expression F is *strictly convex*. On the other hand, F is called (strictly) *concave* if $-F$ is (strictly) convex. Now, COP's have the same form as Eq. (2.4) but the functions $\{F_a\}_{a=0}^m$ are required to be convex, while the function of the equality constraints are required to be *affine*. That is, $H_a(\vec{x}) = \vec{k}_a \cdot \vec{x} - \vec{c}_a$, for constant vectors \vec{k}_a and \vec{c}_a . So, a COP in *standard form* is stated as

$$\min F_0(\vec{x}) \quad (2.7a)$$

$$\text{subject to } F_a(\vec{x}) \leq 0, \quad a = 1, \dots, m; \quad (2.7b)$$

$$\vec{k}_a \cdot \vec{x} = \vec{c}_a, \quad a = 1, \dots, q. \quad (2.7c)$$

Because the intersection of convex sets is convex, the domain of this COP is also convex. This means that we are minimising a convex function over a convex set. Note that this type of problem also includes trivially the case of maximising a *concave* function \tilde{F}_0 , because it suffices to consider the objective function $-\tilde{F}_0$ and minimise it. One of the most important properties of COP's is that *any locally optimal point is also globally optimal*. Moreover, if $F_0(\vec{x})$ is strictly convex, then the optimal set of (2.7) contains at most one point.

LP problems are a particular kind of COP, for which some special results hold. Yet, I will not deal here with such particular cases because the most important result I will use is valid for COP's in general. (But Refs. [191, 192] contain a detailed account of LP problems and special methods for their solution.) So, to continue, I will mention some results of (Lagrangian) *Duality theory*, where for the sake of generality I first assume a general OP as (2.4). The idea of Lagrangian duality is to incorporate the inequality and equality constraints into a new objective function called the *Lagrangian*, $\mathcal{L} : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^q \rightarrow \mathbb{R}$ and defined as

$$\mathcal{L}(\vec{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) = F_0(\vec{x}) + \sum_{a=1}^m \lambda_a F_a(\vec{x}) + \sum_{a=1}^q \nu_a H_a(\vec{x}). \quad (2.8)$$

$\boldsymbol{\lambda}$ and $\boldsymbol{\nu}$ are called *dual variables*. A closely related function is the (Lagrange) *dual function* $G : \mathbb{R}^m \times \mathbb{R}^q \rightarrow \mathbb{R}$ defined as the infimum of \mathcal{L} over \vec{x} , *i.e.*

$$G(\boldsymbol{\lambda}, \boldsymbol{\nu}) = \inf_{\vec{x}} \left\{ F_0(\vec{x}) + \sum_{a=1}^m \lambda_a F_a(\vec{x}) + \sum_{a=1}^q \nu_a H_a(\vec{x}) \right\}. \quad (2.9)$$

It is very important to note that G is defined as a pointwise infimum of affine functions of the dual variables. This implies that G is concave even if the OP (2.4) is *not* convex. Additionally, the dual, G , is closely related to the original OP because it satisfies that

$$G(\boldsymbol{\lambda}, \boldsymbol{\nu}) \leq p^*; \quad \forall \boldsymbol{\lambda} \succeq \mathbf{0}, \quad \forall \boldsymbol{\nu}.$$

(I will use \succeq and \preceq to denote the element wise inequalities \geq and \leq , respectively.) Hence, $G(\boldsymbol{\lambda}, \boldsymbol{\nu}) = -\infty$ if \mathcal{L} is unbounded from below.

From the concavity of G and the upper bound just given, it follows that by maximising G we can “get close to” or obtain –in the best case– the optimal value p^* of the original OP. So we can define the *dual* OP of (2.4) as

$$\max G(\boldsymbol{\lambda}, \boldsymbol{\nu}) \quad (2.10a)$$

$$\text{subject to } \boldsymbol{\lambda} \succeq \mathbf{0}. \quad (2.10b)$$

In analogy with the *primal* OP – i.e. (2.4) –, $(\boldsymbol{\lambda}, \boldsymbol{\nu})$ are dual feasible if (i) $\boldsymbol{\lambda} \succeq \mathbf{0}$; (ii) $(\boldsymbol{\lambda}, \boldsymbol{\nu}) \in \text{dom } G$; and (iii) $G(\boldsymbol{\lambda}, \boldsymbol{\nu}) > -\infty$. Additionally, let d^* and $(\boldsymbol{\lambda}^*, \boldsymbol{\nu}^*)$ denote the optimal value and optimal solution, respectively, of (2.10). Whenever $d^* = p^*$ we say that we have *strong duality*.

Strong duality may occur even for non-convex problems, but in the case of COP and when the inequality constraints $\{F_a\}_{a=1}^m$ are affine, it is possible to show that strong duality holds. This is a corollary of Slater’s Theorem; see [190, Secs. 5.2-5.3]. But, in general strong duality yields a very useful relation between constraints and dual variables:

$$\begin{aligned} p^* &= F_0(\bar{\mathbf{x}}^*) = G(\boldsymbol{\lambda}^*, \boldsymbol{\nu}^*) = d^* \\ &= \inf_{\bar{\mathbf{x}}} \left\{ F_0(\bar{\mathbf{x}}) + \sum_{a=1}^m \lambda_a^* F_a(\bar{\mathbf{x}}) + \sum_{a=1}^q \nu_a^* H_a(\bar{\mathbf{x}}) \right\} \\ &\leq F_0(\bar{\mathbf{x}}^*) + \sum_{a=1}^m \lambda_a^* F_a(\bar{\mathbf{x}}^*) + \sum_{a=1}^q \nu_a^* H_a(\bar{\mathbf{x}}^*) \\ &\leq F_0(\bar{\mathbf{x}}^*) = p^*, \end{aligned}$$

where the last line follows from the fact that $H_a(\bar{\mathbf{x}}) = 0$, while $\boldsymbol{\lambda}^* \succeq \mathbf{0}$ and $F_a(\bar{\mathbf{x}}) \leq 0$, making $\sum_{a=1}^m \lambda_a^* F_a(\bar{\mathbf{x}}^*)$ a sum of non-positive terms. But from the first and last lines, it is obvious that this sum equals zero. That is,

$$\sum_{a=1}^m \lambda_a^* F_a(\bar{\mathbf{x}}^*) = 0 \quad \implies \quad \lambda_a^* F_a(\bar{\mathbf{x}}^*) = 0 \quad \implies \quad \begin{cases} \lambda_a^* > 0 & \implies F_a(\bar{\mathbf{x}}^*) = 0 \\ F_a(\bar{\mathbf{x}}^*) < 0 & \implies \lambda_a^* = 0 \end{cases}. \quad (2.11)$$

This property is called *complementary slackness* and it is essential for relating the dual variables with the contact forces in the LP problem (2.3).

A final results that I will make use of are the so called Karush–Kuhn–Tucker (KKT) optimality conditions. From them, the stability analysis of the packings generated through our iLP algorithm follows easily. They are stated in the following

Theorem 2.1.1 (KKT optimality conditions) *Assume $\{F_a\}_{a=0}^m$ and $\{H_a\}_{a=1}^q$ are differentiable, and let $\bar{\mathbf{x}}^*$ and $(\boldsymbol{\lambda}^*, \boldsymbol{\nu}^*)$ be any primal optimal and dual optimal points, respectively, for which strong duality holds. We know that since $\bar{\mathbf{x}}^*$ minimizes $\mathcal{L}(\bar{\mathbf{x}}, \boldsymbol{\lambda}^*, \boldsymbol{\nu}^*)$ its gradient vanishes*

$$\nabla \mathcal{L}(\bar{\mathbf{x}}, \boldsymbol{\lambda}^*, \boldsymbol{\nu}^*) \Big|_{\bar{\mathbf{x}}^*} = \nabla F_0(\bar{\mathbf{x}}^*) + \sum_{a=1}^m \lambda_a^* \nabla F_a(\bar{\mathbf{x}}^*) + \sum_a \nu_a^* \nabla H_a(\bar{\mathbf{x}}^*) = 0. \quad (2.12)$$

Whence the KKT conditions follow

1. $F_a(\bar{\mathbf{x}}^*) \lambda_a^* = 0$, so either $F_a(\bar{\mathbf{x}}^*) \leq 0$ or $\lambda_a^* \geq 0$; for all $a = 1, \dots, m$.

2. $H_a(\vec{\mathbf{x}}^*) = 0$ for all $a = 1, \dots, q$.
3. $\nabla F_0(\vec{\mathbf{x}}^*) + \sum_{a=1}^m \lambda_a^* \nabla F_a(\vec{\mathbf{x}}^*) + \sum_a \nu_a^* \nabla H_a(\vec{\mathbf{x}}^*) = 0$.

The KKT conditions are necessary conditions to be fulfilled by $\vec{\mathbf{x}}^*$ and (λ^*, ν^*) in any OP. However, in the case of a COP where the objective and inequality constraint functions are differentiable, such conditions are also *sufficient*. These relations are important, because in several cases, it is easier to solve the KKT conditions for $\vec{\mathbf{x}}$ and the dual variables than the original COP. Notice also that condition (2.1.1-1) just restates the complementary slackness condition derived above. That property as well as (2.1.1-3) lie at the basis of the proof of the stability of jammed packings obtained through LP, as I will show next. However, their relevance for physical applications is not limited to the problem discussed in this thesis, but actually is of great generality. Indeed, if F_0 is a potential function, while $\{F_a\}_{a=0}^m$ and $\{H_a\}_{a=1}^q$ are physical constraints, condition (2.1.1-3) is nothing but the force balance equation for each degree of freedom and the dual variables play precisely the role of contact forces. See [190, Sec. 5.5] for a more detailed discussion of the KKT conditions as well as other applications to physical and other type of problems.

2.1.2 Stability and network of contacts from iLP

Using the KKT theorem (2.1.1), it is now easy to analyse the stability of packings obtained upon convergence of the LP algorithm described above. In principle, this means that we need to apply the KKT conditions (2.1.1-1) and (2.1.1-3) in the limit $\vec{\mathbf{s}} \rightarrow 0^*$, $\Gamma^* \rightarrow 1$. (Because the LP problem (2.3) does not contain equality constraints, there is no need to consider the second condition of Theorem (2.1.1).) Let me begin with the latter of these conditions, involving the gradient with respect of the design variables, *i.e.*

$$\nabla = \left(\frac{\partial}{\partial \mathbf{s}_1}, \frac{\partial}{\partial \mathbf{s}_2}, \dots, \frac{\partial}{\partial \mathbf{s}_N}, \frac{\partial}{\partial \Gamma} \right).$$

Notice that in order to convert the jamming LP problem (2.3) into a COP in standard form, like (2.7), we need to minimise the objective function $F_0 = -\Gamma$. Hence,

$$\nabla F_0 = (0, \dots, 0, -1). \quad (2.13)$$

Next, letting

$$F_{(i,j)}(\vec{\mathbf{s}}, \Gamma) = F_{(j,i)}(\vec{\mathbf{s}}, \Gamma) = -2\mathbf{r}_{ij} \cdot \mathbf{s}_{ij} + \Gamma \sigma_{ij}^2 - |\mathbf{r}_{ij}|^2 \leq 0, \quad \text{with } i \neq j$$

be the linear constraints associated to (2.3b) in standard form, we have that

$$\nabla F_{(ij)} = \left(0, \dots, 0, \underbrace{-2\mathbf{r}_{ij}}_{i\text{th entry}}, 0, \dots, \underbrace{2\mathbf{r}_{ij}}_{j\text{th entry}}, 0, \dots, \sigma_{ij}^2 \right). \quad (2.14)$$

Plugging in Eqs. (2.13) and (2.14) into (2.1.1-3) and considering the i th component of the gradient we have that³

$$0 = \sum_{j \neq i}^{1,N} -2\lambda_{(i,j)}^* \mathbf{r}_{ij} = -2 \sum_{\substack{j \in \partial i, \\ [ij] \in \mathcal{C}}} \lambda_{(i,j)}^* \mathbf{r}_{ij} ,$$

where the last equality follows from the complementary slackness condition, (2.11). In this case, such condition means that the sum over all active constraints –*i.e.* the constraints whose dual variables are positive– is equivalent to a sum over all *linear* contacts. Linear, because they saturate the linear non-overlapping constraints, $F_{(i,j)}$ defined above. Nevertheless, at convergence $\bar{\mathbf{s}}^* = 0$ and $\Gamma^* = 1$, so $F_{(i,j)}(\bar{\mathbf{s}}^*, \Gamma^*) = 0$ clearly implies $|\mathbf{r}_{ij}| = \sigma_{ij}$. In other words, in the final configuration of the iLP algorithm, linear contacts are equivalent to real physical contacts. But the correspondence is not only geometrical, but also mechanical. Rescaling the dual variables as $\lambda_{(i,j)}^* = \frac{f_{ij}}{|\mathbf{r}_{ij}|}$, where f_{ij} are yet unknown (although they clearly correspond to the contact forces), the last expression reads

$$0 = \sum_{j \in \partial i} \mathbf{n}_{ij} f_{ij} . \quad (2.15)$$

And this is nothing but the mechanical equilibrium condition in the absence of external forces. Moreover, when all the spatial components of $\sum_{[ij]}^{1,N_c} \lambda_{(i,j)}^* \nabla F_{(i,j)}(\vec{0}, 1) = 0$ are considered, the resulting expressions correspond precisely to the force balance equations of the whole jammed configuration, *i.e.* Eq. (1.60)! This equation was derived in Sec. 1.4.3, where I also showed that jammed packings are mechanically stable provided that it has a non-zero solution. To verify that this is the case in the output of our algorithm notice that the component of (2.1.1-3) corresponding to $\frac{\partial}{\partial \Gamma}$ fixes the scale of the contact forces:

$$\sum_{[ij] \in \mathcal{C}} \frac{\sigma_{ij}^2 f_{ij}}{|\mathbf{r}_{ij}|} - 1 = 0 \quad \implies \quad \sum_{[ij] \in \mathcal{C}} f_{ij} \sigma_{ij} = 1 , \quad (2.16)$$

with the last equation being valid at jamming. But given that all dual variables are non-negative, this last expression implies that at least one of them is positive. Hence, the homogeneous solution to the force balance equation is never a solution of the LP problem. From the arguments of Sec. 1.4.3 we can directly conclude that jammed packings thus produced are mechanically stable because their number of contacts satisfies $N_c \geq N_{dof} + 1 = d(N - 1) + 1$, where I assumed periodic boundary conditions for the last relation.

Here, two remarks are in order. First, although it is possible that more contacts are present, in practice we found that in the vast majority of cases this last inequality is saturated, just as with other algorithms. Hence, most of the configurations have a single state of self-stress (1SS). Second, as a consequence of this feature the packings thus obtained are, actually, predominantly isostatic, *if the whole set of*

³As in Sec. 1.4.3, ∂i denotes the set of particles in contact with i , while $[ij]$, with $i < j$ is the index of the contact between particles i and j . The full network of contacts is denoted as \mathcal{C} and $N_c = |\mathcal{C}|$ is the number of contacts in a jammed configuration.

design variables are considered, *i.e.* (\vec{s}, Γ) . The reason is that Γ acts just as another degree of freedom upon which the optimization is carried out. Physically, this is equivalent to using the density as an extra variable of the configuration, whence the need of the extra contact can be justified. Or, from another perspective, the extra constraint comes from the requirement of jammed states to be rigid.

Incidentally, for monodisperse systems, or whenever $\sigma_{ij} \sim \bar{\sigma}$ —*i.e.* the distribution of radii is peaked around its mean— Eq. (2.16) also determines the typical scale of contact forces. Because for a given configuration contained in a fixed volume its packing fraction scales as $\varphi \sim N\bar{\sigma}^3$, Eq. (2.16) implies that

$$N\bar{f} \sim \frac{1}{\bar{\sigma}} \implies \bar{f} \sim N^{-2/3}. \quad (2.17)$$

Therefore, the results of this part show that configurations obtained with our method are well defined jammed states (*i.e.* determined by the same mechanical equilibrium conditions introduced in Sec. 1.4.3 of the previous chapter), whose stability is guaranteed because they have, at least, 1SS. Moreover, the *active* dual variables corresponding to the non-overlapping constraints provide direct access to the contact forces between particles, whence the full network of contacts can be easily constructed; see Fig. 2.3.

2.1.3 Details of the algorithm and some examples

Some additional remarks are in order to better understand the properties of our algorithm. Adopting first a mathematical point of view, note that the LP problem (2.3) satisfies the Slater’s condition given that its constraints are linear and thus affine, so strong duality always holds. This is important because the solution of the primal problem can be obtained simply by maximising its dual, which amounts to solving the set of linear equations obtained from (2.1.1-3). Second, such equations hold at *each* step of LP optimization, and not only at convergence. This is due to the fact that neither ∇F_0 nor $\nabla F_{(ij)}$ depend on the design variables, and therefore Eq. (2.15) is fulfilled even if $\vec{s}^* \neq \vec{0}$ and $\Gamma^* > 1$. Thus, at each step, configurations are “mechanically stable” with respect to the linear constraints. Geometrically, this feature represents the fact that linear constraints are always saturated, even if no physical contact occurs between particles. In Fig. 2.2a this situation is illustrated after one optimisation step of the LP problem of Fig. 2.1b. The initial feasible set (defined by constraints with the initial positions and sizes) appears in grey; the straight lines are the same linear constraints but using the optimal value of the inflating factor, Γ^* ; and the dashed curves are initial the exact constraints, but also evaluated with Γ^* . The centre of the black disk has been translated according to the optimal displacement vector obtained from the solution and is indicated by the black dot. Notice that it saturates three of the linear constraints but *none* of the exact ones, *i.e.* there is not a single real contact, but three “linear” ones. The solution, nonetheless, satisfies “force balance” with respect to this latter type.

This last consideration evinces the importance of performing several steps of LP, until linear contacts are equivalent to the physical ones. For this to happen, it is not enough for the linear constraints to have “zero gap” with respect of the exact ones, as the blue constraint of Fig. (2.2a) shows. Here, both type of constraints coincide at a

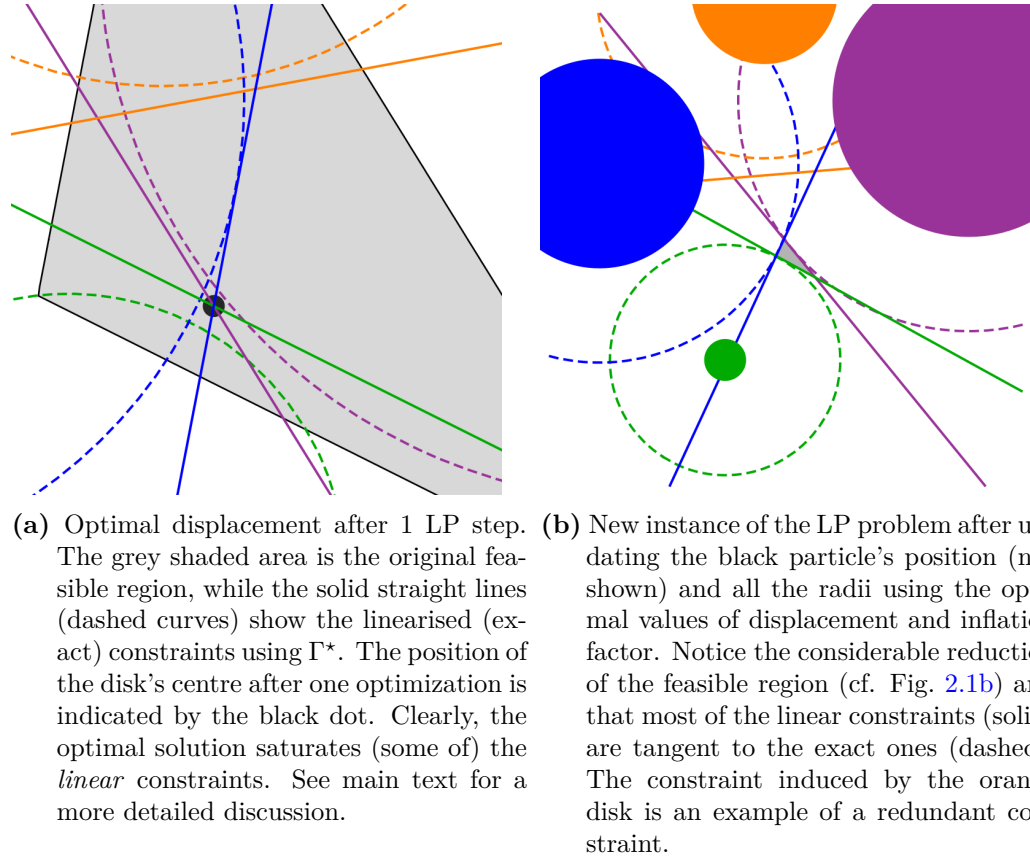


Figure 2.2. Optimal solution and new constraints after the first LP optimization of the linear problem in Fig. 2.1b.

point, but the optimal solution is not on it, so neither in this case a physical contact is realized. It is only when the optimal position coincides with the tangency points of both type of constraints that real contacts take place. And for this to happen it is necessary that $\vec{s}^* = 0$ and $\Gamma^* = 1$. On the other hand, Fig. 2.2b shows that, if particles initially have a lot of free volume, the initial inflation is so large that the feasible region reduces rather quickly. Consequently, linear and exact constraints are expected to match after few steps. However, one should also kept in mind that these figures only consider one movable particle. In the real scenario where all of them are allowed to move more steps are required to ensure real contacts. Fig. 2.2 also reveals that constraints that were initially relevant, might become redundant closer to the jamming point (compare orange constraints in Figs. 2.1b and 2.2b). Even more, the set of inequalities (2.3b) of the jamming LP problem establishes a non-overlapping constraint for *each pair* of particles. But whenever particles are far apart such constraints are obviously redundant and can be omitted so the performance of the algorithm is improved. It is therefore convenient to define a radius of influence for each particle, $\ell_i(\varphi)$, whose value (possibly) depends on the system's packing fraction. Such dependence is justified because for low densities, most of the particles will be displaced by a large amount, and therefore even if a given pair is not relatively close, it might overlap after such large displacement. Conversely, for densities close to φ_J , only the closest particles need to be considered

because very small rearrangements are performed after each LP solution. Note that, in principle, ℓ_i could be different for each particle, hence the subscript. This might be useful, for instance, if a given particle is small enough to escape the cage formed only by its closes neighbours. However, such situation never occurred in the configurations considered here, so I will assume $\ell_i = \ell \forall i$. With this in mind, in what follows the constraints (2.3b) are to be included only if $|\mathbf{r}_{ij}| < \ell(\varphi)$. This is the so called “neighbours list” approach, in which a given particle i only interacts with the set of neighbouring particles $\tilde{\partial}i \equiv \{j \text{ such that } |\mathbf{r}_{ij}| \leq \ell(\varphi)\}$. A final remark: for our periodic systems we used the *nearest image convention*^[194] to compute distances and thus identify the relevant constraints. After all these considerations, our iLP method can be described algorithmically as follows:

Algorithm 2.1: iterative LP (iLP) algorithm to generate jammed packings

Input: Particles’ position $\vec{\mathbf{r}}$ and diameters $\vec{\sigma} \succ 0$

Result: Jammed configuration $(\vec{\mathbf{r}}_J, \vec{\sigma}_J)$ and network of contacts forces

$$\mathcal{C} = \{(f_{ij}, \mathbf{n}_{ij})\}_{[ij]=1}^{N_c}.$$

Compute initial density, φ , and radius of influence, $\ell(\varphi)$;

Define convergence tolerance (tol_s, tol_Γ) ;

Initialize displacements $\vec{\mathbf{s}}$ and inflation factor Γ so they are outside tolerance range;

while $\max_i |\mathbf{s}_i| > tol_s$ **or** $\Gamma - 1 > tol_\Gamma$ **do**

Constrs.	<pre> /* Find relevant constraints according to $\ell(\varphi)$ */ for $i = 1, \dots, N$ do Construct $\tilde{\partial}i$, defined as the set particles j such that $\mathbf{r}_{ij} < \ell(\varphi)$ and $j > i$; /* So each potential contact is counted once. */ for j in $\tilde{\partial}i$ do Include the corresponding linearised non-overlapping constraint, Eq. (2.3b); end end </pre>
LP step	<pre> Solve LP problem (2.3) with the constraints defined in Constrs. and obtain optimal displacements $\vec{\mathbf{s}}^*$ and inflation factor Γ^*; Store the active dual variables $\boldsymbol{\lambda}^*$; $\vec{\mathbf{s}} \leftarrow \vec{\mathbf{s}}^*$ and $\Gamma \leftarrow \Gamma^*$; Update particles’ position $\vec{\mathbf{r}} \leftarrow \vec{\mathbf{r}} + \vec{\mathbf{s}}$ and diameters $\vec{\sigma} \leftarrow \sqrt{\Gamma} \vec{\sigma}$; Update system’s packing fraction using the new diameters; Recompute $\ell(\varphi)$; </pre>

end

/* Construct output variables from data at convergence */

$(\vec{\mathbf{r}}_J, \vec{\sigma}_J) \leftarrow (\vec{\mathbf{r}}, \vec{\sigma})$;

for $[ij]$ in $\{\lambda_{(i,j)}^* \text{ such that } \lambda_{(i,j)}^* > 0\}$ **do**

$\mathbf{n}_{ij} \leftarrow \frac{\mathbf{r}_{ij}}{r_{ij}};$	$f_{ij} \leftarrow \frac{\lambda_{ij}^*}{\sigma_{ij}}$
--	--

end

In this way, packings such as the one depicted in Figs. 2.3 and 3.1 can be generated. On the other hand, notice that Alg. 1 does not mention anything about the initial condition. I did it on purpose to stress that our algorithm is robust enough to produce jammed packings even if the initial condition is not very close to the jamming point. In fact, the configuration of Fig. 2.3 was produced with an initial density of $\varphi_0 \approx 0.4$ (cf. $\varphi_{J,2d} \approx 0.84$), while in the one of monodisperse spheres of Fig. 3.1 the initial packing fraction was $\varphi_0 \approx 0.2$ ($\varphi_{J,3d} \approx 0.64$). Naturally, the closer the initial condition is to the jamming point the faster the iLP algorithm converges. But besides being a matter of speed, we should also keep in mind the physics simulated during the implementation of our method. Considering that all particles' displacements come from the series of LP solutions and not from any thermal energy, our algorithm acts as an immediate quench to $T = 0$ and a very fast compression of the input configuration. In other words, it is a crunching algorithm that, accordingly, does not necessarily follows the glass EOS. And even though the final configuration is a valid jammed state, if only the iLP method is used it may happen that it is rather untypical. For instance, it could be one of the highest minima of the fractal free energy landscape^[3]. Therefore, if we want to explore the typical properties of jammed packings it is convenient to use a glass configuration as thermalised as possible. So I will next describe a method to obtain such configurations.

2.2 Approaching jamming using molecular dynamics

In order to generate a sufficiently well thermalised glass HS configuration we used the Molecular Dynamics (MD) algorithm described in Ref. [94]. Such algorithm works in an event-driven fashion, taking advantage of the fact that particles only interact when they collide with each other and the corresponding collision times can be easily computed. I will first describe why event-driven MD methods are very efficient when using the HS potential, Eq. (1.23), and then discuss how a compression protocol can be implemented. Besides the standard HS configurations, I will also present results of MD simulations in the Mari–Kurchan (MK) model^[143] of HS with infinite range random shifts. In this latter case, the simulations were performed from my own implementation of the event-driven MD algorithm. Producing numerous MK glasses near their jamming point was essential for the results of Chp. 4. Expanding on what I mentioned in Sec. 1.4.1, in the MK model particles interact according to a randomly shifted distance,

$$D_\zeta(\mathbf{r}_i, \mathbf{r}_j) = |\mathbf{r}_i - \mathbf{r}_j + \zeta \mathbf{A}_{ij}|, \quad (2.18)$$

where \mathbf{A}_{ij} is a quenched, uniformly distributed random vector and ζ is a parameter that tunes the effect of the shift. Setting $\zeta = 0$ the usual Euclidean distance is recovered, while $\zeta \rightarrow \infty$ corresponds to the Mean-field (MF) limit. In this latter case, an exact solution to its thermodynamics is available, whence a simple closed-form of its equation of state (EOS) in the liquid phase can be derived^[143]. This is possible because by introducing the random shifts three body interactions (leading to short range correlations) are avoided. Being explicit, if σ denotes the particles




Figure 2.3. Jammed configuration and network of contact forces of $N = 1024$ bidisperse disks and periodic boundary conditions. Rattlers, *i.e.* particles with less than $d + 1 = 3$ contacts, are coloured in orange. The edges of the network are the contact forces, whose thickness indicates the magnitude of the force. The configuration was produced using the Algorithm 1 with particles' initial position drawn uniformly from the box area, initial density $\varphi_0 = 0.4$, and keeping a big/small radii ratio of 1.4.

diameter⁴, in MK configurations even when $D_\infty(\mathbf{r}_i, \mathbf{r}_j) \sim \sigma$ and $D_\infty(\mathbf{r}_i, \mathbf{r}_k) \sim \sigma$, it is almost certain that $D_\infty(\mathbf{r}_j, \mathbf{r}_k) \gg \sigma$. So neighbours of a given particle are very unlikely neighbours between themselves. MK configurations are very useful because they allow to test many of the MF predictions in finite dimensional systems, while being amenable for numerical simulations. I will henceforth assume that $\zeta = \infty$ whenever referring to MK systems. As mentioned earlier, it has been shown that MK configurations of HS display all the rich phenomenology of real glass formers^[109] as well as having a Gardner^[110] and jamming transitions^[144]. Along this section, I

⁴For simplicity, I will henceforth deal with monodisperse configurations. But the results can be easily generalized to systems with polydispersity.

will discuss few other properties of the MK model, as they become relevant to the topics of the MD simulations.

2.2.1 Basics of event-driven MD

MD simulations are particularly well suited for studying the thermodynamics of configurations of HS because their potential renders the dynamics trivial. More precisely, given that particles only interact through elastic collisions, each particle moves ballistically until it bounces off another particle. Compared with most of other types of interaction potentials –for which integrating the equations of motion is mandatory– in HS systems only the computation of collision times is required. Even more, the collision time between a pair of particles is readily obtained by solving a quadratic equation. Such equation is easy to derive considering the diagram of Fig. 2.4, which depicts a (blue) particle colliding with another one (in red), that can be used as reference. That is, let \mathbf{r} be the vector from the reference particle to the incident one, and \mathbf{v} their relative velocity. I will also assume that both particles have the same mass and diameter, σ .




Figure 2.4. A (blue) particle collides with another one (red). To determine the collision time t_c , only their relative position \mathbf{r} and velocity \mathbf{v} are needed. The collision (lighter blue) occurs at the smaller of the solutions of Eq. (2.19), t_c^- . Although the second solution t_c^+ is valid mathematically, it lacks physical meaning (dashed).

The two particles will only collide (lighter blue) if the equation

$$|\mathbf{r} + \mathbf{v}t|^2 = \sigma^2 \quad \Longleftrightarrow \quad |\mathbf{r}|^2 + 2\mathbf{r} \cdot \mathbf{v}t + |\mathbf{v}|^2 t^2 - \sigma^2 = 0 \quad (2.19)$$

admits at least one real, positive solution for t . Using the general formula for quadratic equations it is straightforward to show that this will only happen if: (i) the discriminant is positive (otherwise they do not cross with each other); and (ii)

$\mathbf{r} \cdot \mathbf{v} < 0$, which means that they are approaching each other. If only (ii) does not hold, a collision “would have occurred in the past”, *i.e.* for $t < 0$. Because Eq. (2.19) is a quadratic equation, two solutions exist, t_c^\pm , of which only the smaller (positive) one, t_c^- , is physically relevant. As illustrated with the dashed particle, the other solution also provides a *geometrically* valid collision, but physically unfeasible.

In this way, the collision times between each pair of particles are easily obtained, and by selecting the smallest of such times, the next collision is identified. Hence, the state of the full configuration can be advanced up to that instant and the collision simulated by exchanging the component of the particles’ momenta parallel to \mathbf{r}_c at that time. With the updated velocities, the next collision time is obtained. The dynamics is thus evolved following successive collision events. Moreover, because momentum and energy are conserved at each collision, and the number of particles and system’s volume are fixed, the dynamics corresponds to trajectories sampling the microcanonical ensemble (also called *NVE* ensemble in the MD jargon).

This is the basic idea of *even-driven* MD simulations. It is a very efficient algorithm when dealing with HS systems and its performance can be improved using the following two techniques. First, only potential collisions between particles nearby need to be taken into account. This can be done using the “neighbours list” approach described above for the LP jamming algorithm. However, a faster program is obtained if the simulation volume is divided into cells of fixed size and hence collisions can only occur between particles inside the same cell or the ones next to it. The gain in efficiency is mainly due to the updates of the neighbours lists in the first approach, which requires, every so often, recomputing the distances between all the particle pairs. In turn, in the grid approach the cells remain fixed, so there is only need of bookkeeping the cell to which each particle belongs to. Transfers between cells can also be easily computed and handled as a second type of event. The MD algorithm of Ref. [94] that I used for HS systems belongs to the second class. However, the results of the MK configurations to be considered here and in Chp. 4 were obtained with an implementation of event-driven MD using the neighbours list method, following the algorithm described in [194]. Indeed, in this model the grid method leads to a poorer performance because each particle would need to have an independent grid, given that random shifts are different for each pair. Moreover, a cell transfer for a given grid does not necessarily correspond to a change of cell in the grid of another particle. Hence, every time a particle moves one must check which of the remaining $N - 1$ cell assignments are still valid and which ones need to be updated.

The second way to accelerate event-driven MD is to use asynchronous dynamics (Refs. [194, 195] provide detailed and very clear descriptions). Actually, this is by far the major performance improvement and the reason why event-driven approaches are so useful. Asynchronous dynamics works by computing the impending event for each particle independently and storing the respective times. Because a collision involves only a pair of particles, each event can be processed by updating the state of only those two particles, while leaving the rest of the configuration unaltered. Note that after a collision has taken place, only the new events of the particles involved need to be computed, so most of the remaining collision times previously found remain valid. After the new times have been calculated, the next event on the list is processed and the dynamics continues. Besides, storing and extraction of

collision times can be done very efficiently employing a *binary heap*⁵. In this work, all the MD simulations were carried out in this way.

In general, MD simulations can be used to compute many variables of interest by averaging over time and/or different realizations. As discussed in Sec. 1.2, for HS systems the relevant variables are the (reduced) pressure p , Eq. (1.26), and packing fraction φ , because the temperature is just a scale parameter. In fact, its role is restricted to assigning the initial velocities, which are sampled from a Maxwell–Boltzmann distribution before the MD takes places. Because the kinetic energy is⁶ $K = \frac{d}{2}NT = \frac{1}{2}Nm|\vec{v}|^2$, for a fixed volume, the temperature of a HS system simply determines the temporal scale according to $t \sim 1/\sqrt{T}$. Now, computing the packing fraction is straightforward, but the pressure requires some care when the system is subject to periodic boundary conditions as I will assume throughout this work. Under such conditions no walls are containing the system, so p cannot be measured as the force (per unit area) exerted on them. Nevertheless, a kinetic pressure, associated to the momentum exchange during collisions, can be computed. In fact, from the Virial Theorem^[81] we have that

$$p = 1 + \frac{\beta}{dN} \left\langle \sum_{i < j}^{1,N} \mathbf{r}_{ij} \cdot \mathbf{F}_{ij} \right\rangle = 1 + \frac{1}{2K} \left\langle \sum_c \frac{\mathbf{r}_c \cdot \Delta \mathbf{p}_c}{\Delta t_c} \right\rangle_t. \quad (2.20)$$

To transform the Virial relation into the last expression, the sum over interaction forces has been converted into a sum of the momentum exchanged during each collision, $\Delta \mathbf{p}_c$, and divided by the time passed between successive the collisions, Δt_c . The vector \mathbf{r}_c is the contact vector at the collision time, illustrated in Fig. 2.4. Note also that the average in the rightmost equation is taken over long times (or, equivalently considering a large amount of collisions), while the one in the virial equation is the usual ensemble average.

In Figs. 1.14 and 1.15 I showed that the MD algorithm of Ref. [94] correctly reproduces the liquid and glass EOS, estimating the pressure according to Eq. (2.20), and using as reference the Carnahan–Starling formula (1.31) for the liquid and Eq. (1.32) for the glass⁷. Similarly, with the analogous algorithm for the MK systems the liquid phase can be accurately simulated as illustrated in Fig. 2.5. In this case, the respective EOS of an MK liquid is given, exactly, by the MF equation (1.39b) in the $d \rightarrow \infty$ limit. The excellent agreement between the simulations and the closed expression confirms the MF nature of the model. In other words, random shifts effectively suppress correlations between three or more particles, so only the leading term of the virial expansion, associated to pairwise correlations, is non-zero. For the same reasons. although in a different situation, within the MF theory this is the only non-vanishing correction to the ideal gas pressure; see Sec. 1.3.1

⁵ A heap is a data structure that allows to store, extract, and update ordered data more efficiently than by using a sorted list. To compare, given a set of n data, implementations of both quick sorting and a binary heap scale, on average, as $\mathcal{O}(n \log n)$. However, once the objects are constructed the time complexity of inserting and deleting elements in a sorted array scales as $\mathcal{O}(n)$, while using a heap structure it is reduced to $\mathcal{O}(\log n)$.

⁶ Recall that through the text I am using $k_B = 1$.

⁷ Although these figures were obtained from MD simulations with a compression protocol (as described in the next subsection), the compression was very slow so a good estimate of p was available for each density.

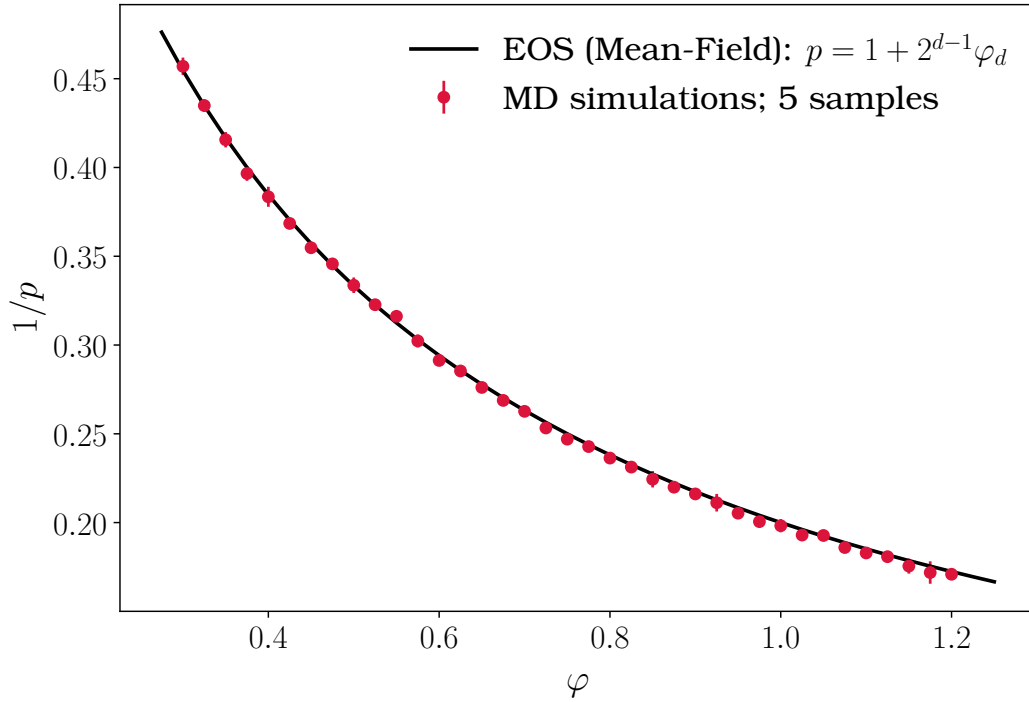


Figure 2.5. Comparison of event-driven MD simulations (markers) and the EOS, Eq. 1.39b (black curve). MD data corresponds to the average and standard deviation over 5 independent estimations of p for each φ . Note that because of the MF nature of the MK configurations, truncating the virial series after the second term leads to the *exact* EOS. Additionally, because of the random shifts, MK configurations can reach densities that would be unphysical in standard HS systems, *i.e.* $\varphi > \varphi_{FCC}$ in 3d or even $\varphi > 1$. All of this while remaining in the liquid phase.

Fig. 2.5 also shows that, due to random shifts, MK configurations can attain densities that are prohibitively high in usual models, *e.g.* $\varphi > 1$. Once again, this is caused by the fact that neighbours of any given particle are not mutual neighbours between themselves, so they can be overlapping from the point of view of such particle. Consequently, a particle has much more neighbours than if random shifts were absent. From an algorithmic point of view, this feature hinders the MD simulations because neighbours lists are considerably larger than for usual HS systems, so more potential collisions are to be considered when obtaining the set of potential t_c . I should also mention that for $\varphi > 0.7$, the initial configuration of the MD simulations was generated via *planting*. That is, for a given value of φ , the set of particles' initial positions \vec{r}_0 is drawn from a random distribution, uniform over all the volume. And only *afterwards* the shifts \mathbf{A}_{ij} are drawn and assigned, but with the condition that no overlap occurs for any pair of particles, so \mathbf{A}_{ij} is repeatedly drawn until such condition is met. With this trick, liquid MK configurations of arbitrarily high densities^[109] can be produced without altering the equation of state. This method will be important for simulating correctly the glass phase as I discuss next.

2.2.2 Compression protocol

For compressing the liquid into a glass we used a Lubachevsky–Stillinger (LS) compression protocol^[196]. It consists in increasing the particles’ diameter at a given rate $\dot{\sigma}(t)$ and it is therefore easy to implement in periodic systems as we consider here. As shown from the results of Ref. [94], in HS systems a glass can be readily obtained from a liquid using the simplest compression protocol, *i.e.* one with a constant growing rate $\dot{\sigma} = \kappa$. In that work, MD simulations with the LS protocol were tested in monodisperse systems in several dimensions, and latter, in [99], the same methodology was used for $3d$ polydisperse configurations. Analogously, such LS compression method has been used to study MK configurations in their glass phase; see [109, 110]. Similarly, the implementation I developed of the LS protocol is able to reproduce the expected EOS of MK liquids and glasses, as shown in Fig. 2.6. An important feature of the MK model is that it exhibits a dynamical glass transition, just as the one described by MF theory in Sec. 1.3.2, at $\varphi_d \approx 1.776$. Notice that a slow compression, say $\kappa = 10^{-4}$ (blue line in Fig. 2.6), yields an out of equilibrium liquid at densities well below φ_d , and decreasing the growth rate by an order of magnitude barely improves the situation^[109]. This is a manifestation of the true dynamical transition predicted by MF theory, and therefore no compression protocol is able to produce an equilibrated state for $\varphi > \varphi_d$.

On the other hand, these results also show the importance of using planted initial conditions in order to access the full glassy phase at high densities. When such method is used with no compression (*i.e.* $\kappa = 0$), MD simulations correctly estimates the pressure of the *equilibrated* glass (cyan diamond) for densities considerably above φ_d (vertical dash-dotted line). Yet, independently of whether planting is used or not, the LS protocol generates a dynamics that correctly reproduces the metastable glass behaviour, as indicated by the excellent agreement with the red dashed lines. These curves were obtained from a fit to Eq. (1.32), which describes the glass EOS in HS systems. However, when fitting the data I treated d in this equation as a free parameter, given the lack of inherent dimensionality in MK systems.

Before continuing, some clarification of the glass EOS is in order. In Figs. 1.15 and 2.6, I have used Eq. (1.32) as the expected theoretical behaviour of a glass, and showed that it is nicely verified by numerical simulations. However, it is important to stress that this equation is *not* derived from thermodynamic considerations, but only from free volume (*i.e.* geometric) arguments; see Sec. 1.2 and [98]. (Calling it “Equation of State” is actually an abuse of terminology.) As discussed in Sec. 1.3.2 and 1.3.4, the true EOS is obtained, at least for MF models, by the technique of state following, and requires using the Franz–Parisi potential in a 1RSB and fullRSB schemes. Restricting the discussion to the MK model, for which a full MF treatment is possible^[109,110,197], notice that the true thermodynamic EOS of a glass depends on the values (p, φ) at equilibrium –for instance, the cyan diamond in Fig. 2.6. In contrast, the glass pressure using the free volume EOS depends, quite naturally, on the density at jamming. Consequently, it does not provide information about any equilibrium properties, not even when crossing the liquid EOS line⁸. As

⁸Beyond φ_d , the “liquid” branch actually corresponds to an equilibrated glass. See Sec. 1.3 for a more detailed discussion.

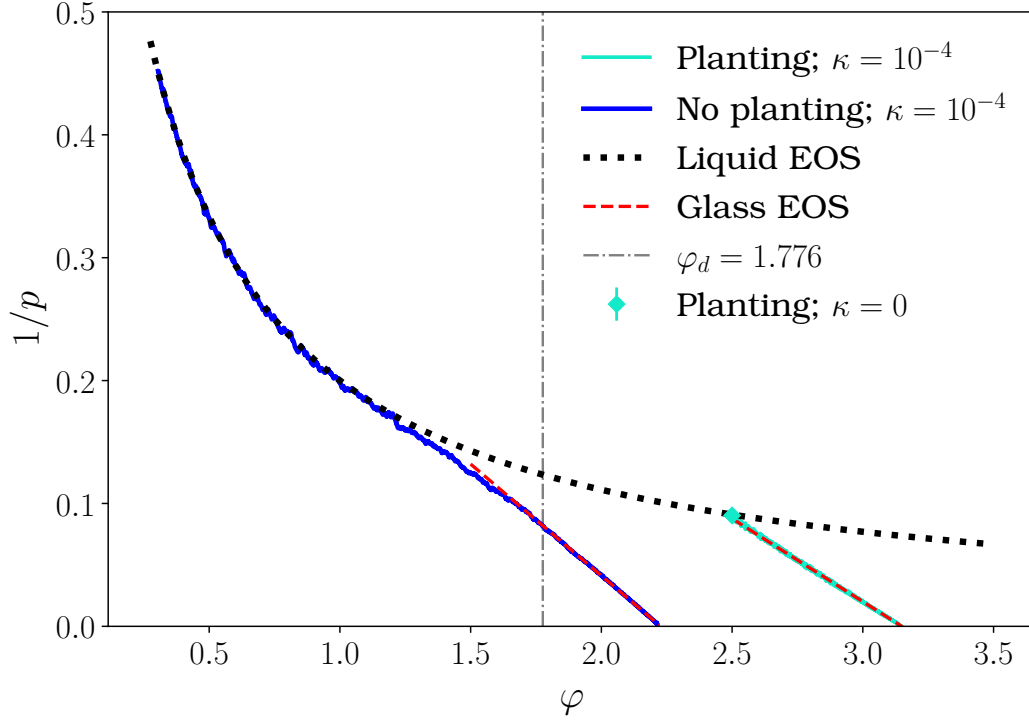


Figure 2.6. Liquid (dotted, black) and glass (dashed, red) branches in the MK model. The dynamical glass transition of this model, characterized by a divergent relaxation time, occurs at ϕ_d (dash-dotted). If no planting is used, the liquid is unable to equilibrate beyond ϕ_d using any compression protocol, including the LS one. However, planting allows to obtain an equilibrated liquid (cyan diamond) at an arbitrarily high density, which can then be used to probe the glass phase when compression is “switched on” (cyan curve). The glass EOS curves were obtained from a fit of the numerical data to Eq. (1.32).

expected, no trace of the Gardner transition is present in this latter EOS, while it is a salient feature of the thermodynamic construction based on state-following; see [110, 197, 198]. Even more, no other thermodynamic variables or free energies can be obtained from Eq. (1.32). As a side note, I take the opportunity to mention that planting is another type of method that allows to equilibrate *some* systems beyond the dynamical glass transition; see the discussion at the beginning of Sec. 1.3.3. Unfortunately, this is not true in general since the free energy of a planted system usually differs from the real thermodynamic one. But in the case of the MK model, planting is a very useful trick because the annealed (*i.e.* planted) and quenched free energies of the liquid coincide^[109].

Let me now briefly discuss the main implications of using a LS compression protocol in the event-driven MD algorithm discussed above. Given that the particles’ diameter is also a function of time, the most obvious consequence is that Eq. (2.19) cannot be solved analytically for a general growth function $\dot{\sigma}(t)$. However, when a constant compression is used, as in all the cases studied here, the diameter is simply $\sigma(t) = \sigma_{t_0} + \kappa t$, so Eq. (2.19) remains quadratic in time. In contrast, the effect of particles’ growth on the momentum transfer due to collisions is not so obvious. If

two particles i and j collide, the new velocity of particle i (resp. j) has an extra contribution^[94,196] equal to $\hat{\mathbf{r}}_{ij}\kappa$ (resp. $\hat{\mathbf{r}}_{ji}\kappa$), where $\hat{\mathbf{r}}_{ij} = \frac{\mathbf{r}_i - \mathbf{r}_j}{|\mathbf{r}_i - \mathbf{r}_j|}$ is a unit vector connecting the particles' during the collision. This extra contribution can be understood by considering the case of two static particles, where clearly no collision is expected. However, if their size is being increased, sooner or later they will come into contact. When this happens, the only source of “impulse” they have is due to the speed at which their radius is increasing, and the “collision” takes place along the line connecting their centres. The same argument applies when particles are not static, thus the extra contribution to their momenta is justified. Because $\hat{\mathbf{r}}_{ij} = -\hat{\mathbf{r}}_{ji}$, momentum is still conserved, but the energy is not⁹. In fact, the difference in kinetic energies before and after the collision scales as κ^2 . A simple solution is to rescale, after a fixed number of events, all the particles velocities so that their kinetic energy matches $K = \frac{d}{2}NT$. On the other hand, 2κ should be subtracted from $\Delta\mathbf{p}$ in each collision in order to avoid the extra velocity contribution to affect the estimation of p in Eq. (2.20).

As a final verification that the compression algorithm described here is able to produced well thermalised configurations for very high pressures, I will compare the gaps distributions obtained from MD simulations with the radial distribution function (RDF) at jamming^[94]:

$$g(h) = \frac{\bar{z}\sigma}{\rho s(\sigma)}\delta(h) + cg(h). \quad (2.21)$$

In this equation, $\bar{z} = 2d$ is the average coordination number, $\rho = N/V$ is the number density, $s(\sigma)$ is the surface of a hypersphere of radius σ , c is a geometric constant and $g(h)$ is the critical gaps distribution defined in Eq. (1.67). (Notice that this latter quantity is written in italic, while the RDF is in regular typeface.) Additionally, instead of considering the RDF as a function of r , as I did when it was introduced in Sec. 1.2, I am using $h = \frac{r}{\sigma} - 1$, which amounts to rescaling the distances by σ and translating the plot of $g(r)$ by the same value. Notice that this is just the gap definition given in Eq. (1.49) and used to study jammed packings in Sec. 1.4, but particularized for monodisperse systems. Now, Eq. (2.21) tells us that, at jamming, the RDF is made of a δ singularity due to particles in contact ($r = \sigma$) that decays following a non-trivial power-law: $g(h) \sim h^{-\gamma}$, with $\gamma \approx 0.41$ as discussed in Sec. 1.4.4. This is nothing but a consequence of $g(r)$ being proportional the probability of finding a pair of particles a distance r apart. Indeed, as I will show next, when p is high enough, these characteristic properties of g are reflected by the probability distribution function (pdf) of h .

To estimate the pdf of h , I obtained the histogram of interparticle gaps using the compression method just described in 20 systems of $N = 1024$ HS, at different final pressures and with different values of κ . The results are reported in Fig. 2.7, where a comparison with the power-law of $g(h)$ for $10^{-7} \leq h \leq 1$ is also included. The numerical results are in excellent agreement with the theoretical prediction over a range that, naturally, extends as p increases. On the other hand, the effect of changing κ , even by an order of magnitude is barely noticeable. This is an important

⁹Thus, strictly speaking the LS compression protocol modifies the sampling ensemble of the MD simulations.




Figure 2.7. Probability distribution function of interparticle gaps $h = \frac{r}{\sigma} - 1$ obtained from the MD based compression. The different values of the final pressure are indicated by different colours, while the distinct line styles identify the values of the particles' growth rate. A comparison (in cyan) with the critical gaps distribution expected at jamming is also included, while the predicted onset of the plateaus at $h = 1/p$ is indicated by the dash-dotted vertical lines. See main text for a discussion of why these two features are important.

point to which I will come back later. A second relevant attribute is the formation of plateaus whose height also increments with p . To understand this behaviour, recall that Eq. (1.25) relates the pressure of a HS system with the value of the RDF at $r = \sigma$, g^+ . This expression was obtained for a liquid, but the only underlying assumption is that the pressure can be obtained via the Virial theorem and, as argued few paragraphs above, this is the case in MD simulations, provided that κ is small enough so that averaging p over a large number of collisions matches its thermal average. Besides, note that when the final pressure of the MD simulations is increased by, say, two orders of magnitude the respective plateau is raised by the same amount. We can then conclude that, except for geometric factors, the height of the left part of the histograms equals $g^+ \sim p$. The plateaus are thus caused by a nascent singularity that gives rise to the δ function in Eq. (2.21). This means that as $p \rightarrow \infty$ the length of such plateaus should vanish, whence it is expected that their onset is reduced with p . In fact, when substituting \bar{z} and $s(\sigma)$ by their value in $d = 3$, some simple algebra leads to

$$\lim_{\epsilon \rightarrow 0} \int_0^\epsilon g(h) dh = \frac{1}{4\varphi}.$$

In turn, from the histograms obtained numerically, we have that the same integral is approximately ϵg^+ . But from Eq. (1.25) we know that for $p \gg 1$, $g^+ = \frac{p}{4\varphi}$, whence

we obtain that $\epsilon \sim 1/p$. This result is confirmed by the dashed-dotted, vertical lines of Fig. 2.7, all of which are drawn at $1/p$ and that are in excellent agreement with the rightmost part of the plateaus.

To make an explicit connection with the inners of the MD simulations, let me give an additional argument of why it is expected that the plateaus begin when $h \sim 1/p$. It goes as follows: according to Eq. (2.20), the value of p depends on how often collisions occur in the system. Given that velocities follow a Maxwell–Boltzmann distribution their variance is finite, so typical temporal intervals between collisions are of order $\Delta t \sim \frac{h}{v} \sim \frac{h}{\sqrt{T/m}}$. Because $p \sim 1/\Delta t$, it then follows that gaps which contribute the most to the kinetic estimation of p are, at most $\mathcal{O}(1/p)$. These considerations show that before jamming is reached, the δ singularity of Eq. (2.21) is smoothed out as a plateau of increasing height but decreasing width, caused by particles colliding with their nearest neighbours. These interactions define (some of) the contacts that will form at φ_J . Additionally, there are other set of very small gaps (although several orders of magnitude larger than the former), that begin to delineate the critical gaps distribution; see Eq. (1.67) in the previous chapter. Fig. 2.7 shows that if jamming has not been reached, there is a crossover between these two sets, where it is impossible to know if a given gap should be classified as a contact or as a true gap. Moreover, this intermediate interval is not affected by using different compression rates, so it is not caused by a lack of thermalisation of the configurations.

Interestingly, this same argument also shows that the MD-based compression identifies the $p = \infty$ point with the situation where $\Delta t = 0$. But notice that this situation does *not* necessarily correspond to a jammed state. For instance, there could be a percolating chain of contacts, leading to a set of particles blocked by the event-driven dynamics, and yet be an unjammed configuration. The case of the contact forces is harder to solve: even if contacts could be perfectly identified by the $h \sim 1/p$ criterion, the problem of defining the forces magnitude would remain. Therefore, despite the similarities of the pdf of h at $p \gg 1$ but finite, and the expected behaviour at jamming, the configurations obtained using MD are detectably different from real jammed packings. Because all the results of this thesis are based on a precise identification of gaps and forces, we needed to complement the MD method with our iLP algorithm to guarantee that systems are *at* jamming. On the other hand, as I have argued along this part, the compression protocol considered here is a good approximation of the state-following scheme. Therefore, using the configurations thus produced as initial conditions of the iLP algorithm is a way to ensure that typical jammed states are obtained. Combining both algorithms yields a robust and relatively fast method, as I will now describe.

2.3 Characterization of MD+iLP algorithm

As mentioned in the last part of the previous section, configurations compressed employing the MD algorithm will be used as initial conditions of the iLP algorithm described in Sec. 2.1. Figs. 1.15, 2.7 (for HS) and 2.6 (in the MK model) show that when a small value of κ is used, configurations can be compressed very close to their jamming point, identified by the $1/p = 0$ line. Hence, such initial conditions will

be parametrized by the target pressure p_{tar} at which the LS compression protocol stops, and by the diameter growth rate κ used in such protocol. We decided to use the pressure instead of φ because the packing fraction at jamming is system dependent, and might even be distributed within an unbounded interval (as in the MK configurations). In turn, jamming of rigid particles always ensues when $1/p = 0$. Additionally, the required fine tuning in φ is much harder to achieve, since p can increase by 4 or more orders of magnitude, while φ changes its value by less than 1%, see Fig. 1.15. Hence, fixing a threshold value of the packing fraction would yield a sample with very heterogeneous pressures.

The results that I will present in this section are a detailed characterization of the properties and performance of the MD and iLP algorithms combined. I will only consider HS configurations, but the results should be qualitatively analogous for the MK model. As I have mentioned before, in $d = 3$ monodisperse HS configurations tend to crystallize if a very slow compression is utilized. To avoid it, all the configurations I will consider were produced after an initial and fast ($\kappa = 5 \times 10^{-3}$) compression up to $p_{\text{tar}}^{(0)} = 500$. From the results of Figs. 1.14 and 1.15, it is clear that this value of κ is enough to suppress the formation of crystalline domains, while $p_{\text{tar}}^{(0)}$ is well into the glass phase. All the MD compressions were performed using the implementation of Ref. [94], while each LP optimization step of our method was carried out using Gurobi^[199].

2.3.1 Influence of the initial condition's parameters

Possibly the first thing to consider is how the jamming packing fraction φ_J depends on the parameters of the initial condition, (p_{tar}, κ) . The results obtained for $p_{\text{tar}} \in [10^3, 10^{11}]$ and for several values κ are reported in the upper panels of Fig. 2.8. Each point corresponds to the average value of φ_J (and the standard error) over 20 (resp. 10) samples for $N = 1024$ (resp. $N \geq 2048$); hence throughout this section $\langle \cdot \rangle$ denotes the average over samples. As anticipated, all the values are close to 0.64, and changing the target pressure of the compression protocol has little or negligible influence for $p_{\text{tar}} \geq 10^6$. In contrast, the effect of changing the growth rate (left panel) and system size are notorious (right plot), albeit similarly small. As expected, lowering κ yields a higher value of φ_J because it amounts to a better thermalised compression, so it can reach lower minima of the free energy landscape (FEL). Likewise, the jamming density attains higher values as N augments. This result is also unsurprising given that the larger the system, the less it is constrained by border or periodic effects; hence it can better rearrange its particles to achieve a higher φ_J . Of course, φ_J cannot grow indefinitely as $N \rightarrow \infty$, and indeed the data suggest a rather quick convergence to its thermodynamic limit value. I should emphasize that no crystallization occurs in these configurations (see Fig. 2.10), so for a fixed N , φ_J is narrowly distributed around a well defined value^[132,133,136]. If this was not the case, a much broader distribution of values could be obtained; see [187] for an example where also an LP-based jamming algorithm is used.

Let me discuss in greater detail the dependence of the final configurations on p_{tar} . Such analysis is relevant since it puts to test the hypothesis of the fractal structure of the FEL. As explained in Ref. [111] a careful investigation is required to explore such scenario, and while the results presented here are not intended for




Figure 2.8. Upper panels: Dependence of the jamming density φ_J on the target pressure from the compression protocol, p_{tar} . Lower panels: Statistics of the difference in packing fractions, $\Delta \varphi_J(p_{\text{tar}}) \equiv \varphi_J^{(p_{\text{max}})} - \varphi_J^{(p_{\text{tar}})}$; because for some samples such difference is not monotonic, the absolute value of the difference is considered. Several growth rates with $N = 1024$ are considered (left) and different system sizes with fixed $\kappa = 10^{-4}$ are analysed (right). Markers denote the sample average, while the error bars are the standard error.

those purposes, I will argue that they nevertheless support such hypothesis. To do so, I will consider as a reference for comparison the configurations obtained from the largest target pressure $p_{\text{max}} \equiv 10^{11}$. First I will focus on the difference $\Delta \varphi_J(p_{\text{tar}}) \equiv \varphi_J(p_{\text{max}}) - \varphi_J(p_{\text{tar}})$ as a measure of how much the final configuration changes if the iLP crunching begins at a different p_{tar} . The lower panels of Fig. 2.8 depict the dependence of $\langle |\Delta \varphi_J(p_{\text{tar}})| \rangle$ of the same samples considered before. The absolute value is necessary because, unexpectedly, the difference is not monotonic. That is, some samples attain a larger density when $p_{\text{tar}} < p_{\text{max}}$, so in order to avoid deviations to larger values cancelling the ones to smaller values the absolute value is used. In any case, when $N = 1024$ and the value of κ is varied, our results show (bottom left panel) that the difference of jamming packing fractions decreases almost steadily as p_{tar} grows. This behaviour is precisely the expected one from a fractal FEL, because it implies that the minima –attained from different p_{tar} and that define the jammed states– are similar between themselves (because they all belong to the same meta-basin), but their specific features are realized gradually (because the configuration navigates the different sub-basins as p_{tar} increases). Conversely, if the FEL possessed a simple basin structure, once p_{tar} was high enough all the configurations would end up in the same minimum, so the final packings would

be identical. Second, for $p_{\text{tar}} \leq 10^6$ results are rather insensitive to the compression rate. However, while $p_{\text{tar}} \geq 10^8$ is enough to suppress fluctuations if $\kappa = 3 \times 10^{-4}$, detectable differences are present even if $p_{\text{tar}} = 10^{10}$ when the slowest compression is used. Our data thus suggest that the hierarchical structure of the FEL can only be fully resolved with an infinitesimally slow compression. Similarly, when κ is fixed and different system sizes are considered (bottom right panel), $\Delta\varphi_J(p_{\text{tar}})$ is monotonically decreasing up to $p_{\text{tar}} \sim 10^7$. For higher pressures, the difference with respect to p_{max} seem to increase with N . This is, very likely, once again a manifestation of the finite value of κ , which determines the *radius* growth rate. Because the volume of all the systems is fixed, configurations made of more particles have smaller spheres and, consequently, using the same value of κ for all of them produces different compression rates. Additionally, the discrepancy found for different N is also possibly caused by finite size effects, given that the marginal fractal phase is strictly valid only in the thermodynamic limit.




Figure 2.9. Overlap (defined in the main text) of the networks of contacts using as reference the one corresponding to p_{max} . The overlaps increase *gradually* as the configurations go deeper in the FEL, suggesting that it has a fractal structure. The data reported here correspond to the same configurations of Fig. 2.8.

As a second measure of the similarity between jammed configurations I will consider the overlap of the network of contacts. Thus, let $n_{ij}(p_{\text{tar}}) = 1$ if particles i and j are in contact in the jammed configuration obtained from an initial condition with p_{tar} , and 0 if they are not in contact. Taking once again p_{max} as reference, the overlap between two networks of contacts is then $\overline{n_{ij}(p_{\text{tar}})n_{ij}(p_{\text{max}})}$, where the overline denotes an average over all the contacts of the reference network. Notice that this variable is more sensitive to variations in the jammed state, given that the full microstructure of the packings is considered. In fact, this is the same quantity analysed in Ref. [111] to investigate the fractal structure of the FEL. However, the packings produced there were not precisely at jamming –but at a very high p – and the authors were concerned with the overlap of *different* configurations at the *same* pressure. Here I am instead considering the jamming point of the same system, reached from different pressures. Fig. 2.9 shows the results of the similarity of the

contact networks obtained from the configurations considered above. Also in this case the data support the fractal FEL hypothesis, because as p_{tar} increases the network of contacts are more alike. Once again, this can be explained assuming that the broad structure of the contacts network is shared by packings within a given meta-basin, but the precise particle contacts are determined *progressively* as a configuration goes down in the hierarchical structure of minima. It is worth emphasizing that the discrepancies found when varying κ and N are in tune with the explanation given above, namely: the need to use an adiabatic compression and be in the thermodynamic limit in order to probe the fractal structure of the FEL at all scales and for every value of p . In fact, when this is not the case, our results show that for some values of κ and N , and at sufficiently high pressure, the networks of contacts are *identical in all the samples*. At any rate, I should warn that even if these results are very promising and complimentary to the ones of Ref. [111], they are still preliminary given the few samples considered. Additionally, a more systematic analysis of the influence of the compression protocol and system size is needed to explain satisfactorily the dispersion of values found for $p_{\text{tar}} \geq 10^8$.

As a final feature, I will show that no crystalline domains are present in the jammed packings produced by our MD+iLP method. To do so, I will resort to the RDF computed using Eq. (1.17) and averaging over the 20 jammed packings of $N = 1024$ particles. Fig. 2.10 contains several plots of $g(r)$ obtained using different values of p_{tar} to generate the jammed packings, and fixing $\kappa = 3 \times 10^{-4}$. However, changing the compression rate leaves the RDF unaltered, as evidenced by the dashed, dash-dotted, and dotted black lines (corresponding to $\kappa = 10^{-4}$, 3×10^{-5} , 10^{-5} , respectively) superposed to the curves of the respective pressure. If crystallization did occur, a peak would appear exactly at $\sqrt{2}$, due to the distance between the pair of particles in the diagonal of a square^[200,201]. Given that $g(r)$ has a minimum very close to $r = \sqrt{2}\sigma_J$ (red dashed line), such peak would be easily noticeable. For instance, the inset shows the RDF of 10 configurations with a small degree of crystallization, $\varphi \in (0.65, 0.675)$, all of which show an evident peak at the predicted position. Another signature of crystallization is the peak at $\sqrt{7}$ (pink dashed, line), although its presence is less pronounced and it is only smoothed out in amorphous packings. In any case, these configurations were not included in any of the analyses I present here. On the other hand, the peaks at $\sqrt{3}$ and 2 are expected in disordered configurations^[94,172,200,201]. At $r = 3\sigma_J$ a discontinuity is expected, although it is hardly visible at this scale. A detailed analysis of the behaviour of the RDF in jammed packings can be found in [172, 200, 201], while a nice example of how the peaks disappear as a configuration moves away from the crystal is presented in Ref. [163], albeit using polydispersity as the source of disorder. At any rate, the behaviour of $g(r)$ obtained from our method agrees with previous results far from contact – that is, for $r/\sigma_J - 1 \gg 1$ – confirming that the MD+iLP algorithm fully suppresses crystallization.

2.3.2 The MD+iLP route to jamming

In this subsection I continue analysing some features of the jammed packings obtained by combining the MD and iLP algorithms. But now, instead of considering the properties of the final jammed configuration, I will focus on how such configu-




Figure 2.10. Radial distribution function (with the distance scaled by the diameter at jamming, σ_J) and its dependence on the target pressure of the initial configuration. Each curve is obtained by averaging 20 samples of $N = 1024$ particles. The curves have been displaced vertically for clarity, because they overlap almost perfectly. All of them were obtained using the largest value of κ in the MD compression, but changing the compression rate leaves $g(r)$ unaltered as shown by the dashed ($\kappa = 10^{-4}$), dash-dotted ($\kappa = 3 \times 10^{-5}$), and dotted ($\kappa = 10^{-5}$) lines. Vertical lines at $\sqrt{2}$ and $\sqrt{7}$ identify the position of expected peaks if crystallization occurs, while the grey dotted lines are present also in amorphous systems; see text. Inset: Sample of 10 configurations (in grey) and their average (cyan) where partial crystallization occurs. The presence of the peak at $\sqrt{2}$ is a clear signature of partial ordering, while the one at $\sqrt{7}$ is only smoothed out.

rations are reached. In other words, I will describe some of the physics underlying the MD+iLP protocol. For this purpose, I begin by arguing that, *based solely on a (p, φ) parametrization* of a jammed state, the packings obtained through the two algorithms combined cannot be distinguished from the ones corresponding to the $p \rightarrow \infty$ limit using only the MD compression. To perform such comparison, note first that the density value extrapolated to this limit, $\varphi_J^{(MD)}$, can be readily estimated by fitting the values of $\frac{1}{p}$ and φ obtained from MD simulations with a given κ to Eq. (1.32). The results that I present next were obtained from the 20 configurations of $N = 1024$ particles and the data of $p \in [10^5, 10^{11}]$ for each $\kappa \in \{3 \times 10^{-4}, 10^{-4}, 3 \times 10^{-5}, 10^{-5}\}$. Instead of fixing $d = 3$, I considered d as a free parameter for the fit in order to account for the deviation with respect to the EOS and thus better estimate of $\varphi_J^{(MD)}$. Nonetheless, in all the fits I found $1 \leq d/3 \leq 1.07$ and $0.999 \leq R^2$, indicating that they are reliable enough. $\varphi_J^{(MD)}$

is then equal to d divided by the estimated slope. The average densities resulting from such fits are reported in Table 2.1, as well as the analogous packing fractions with target pressures $p_{\text{tar}} = 10^3$ and $p_{\text{tar}} = 10^5$. Results with larger target pressures are the same, considering the standard error.

κ	$\langle \varphi_J^{(MD)} \rangle$	$\langle \varphi_J(p_{\text{tar}} = 10^3) \rangle$	$\langle \varphi_J(p_{\text{tar}} = 10^5) \rangle$
3×10^{-4}	0.64223(19)	0.64212(17)	0.64223(19)
10^{-4}	0.64255(21)	0.64252(21)	0.64255(21)
3×10^{-5}	0.64291(24)	0.64288(24)	0.64291(24)
10^{-5}	0.64301(23)	0.64299(23)	0.64301(23)

Table 2.1. Average values of the jamming density of the MD compression protocol alone, and complimented by iLP. Each row corresponds to a different growth rate. Packing fractions in the $p \rightarrow \infty$ limit are estimated through a fit, as described in the text. The uncertainties reported are the standard errors, since these are larger than the ones obtained from the error propagation of the fits estimates.

It is clear that when the standard error associated to the sample is considered, jamming densities obtained from MD cannot be distinguished from the ones of MD+iLP, even for the smallest p_{tar} we considered here. Yet, because of sample to sample fluctuations of the values of φ_J and $\varphi_J^{(MD)}$ (although always within a small range), it might be useless to compare their ensemble average. An alternative is to consider the difference $\varphi_J(p_{\text{tar}}) - \varphi_J^{(MD)}$ from individual samples. However, the errors associated to the fitted parameters, whence $\varphi_J^{(MD)}$ is calculated, lead to an uncertainty several times bigger than such difference. These findings are depicted in Fig. 2.11, where $|\varphi_J(p_{\text{tar}}) - \varphi_J^{(MD)}|$ is plotted for each sample, all values of κ used so far, and $p_{\text{tar}} \geq 10^5$. For comparison, the range where uncertainty values of $\varphi_J^{(MD)}$ lie is indicated by the grey shaded region, showing that it is at least two orders of magnitude larger than¹⁰ $|\varphi_J(p_{\text{tar}}) - \varphi_J^{(MD)}|$. This implies that the two jammed states cannot be told apart *macroscopically*.

On the other hand, note that the average values of $\Delta\varphi_J$ are several orders of magnitude smaller than $|\varphi_J(p_{\text{tar}}) - \varphi_J^{(MD)}|$; cf. with the black curves included in Fig. 2.11. This means that even when there is a measurable difference in the iLP packing fraction, it is impossible to discern it only through a comparison with the values corresponding to MD. This would call for a *microscopic* comparison between the two types of jammed states, if only the MD jammed packings were realizable. Unfortunately, this is clearly not the case. But even if a $p = \infty$ configuration could be produced through MD compression, there is no guarantee that it will be strictly jammed. As I argued in Sec. 2.2.2, whenever $\kappa > 0$, a $p = \infty$ *could* correspond to a configuration in which only few particles are in contact, leaving several degrees of freedom unconstrained, and thus far from the 1SS requirement for jamming criticality (see Sec. 1.4.2). In other words, although our simulations do not show any inconsistency with the hypothesis that φ_J and $\varphi_J^{(MD)}$ identify the same jammed state, proving that this is the case is far from easy.

¹⁰I should mention that I have introduced the absolute value only because for $p_{\text{tar}} \leq 10^5$ some samples have $\varphi_J < \varphi_J^{(MD)}$, while for greater pressures the opposite is true in the vast majority of cases. Hence, removing the absolute value makes no difference in our conclusions.

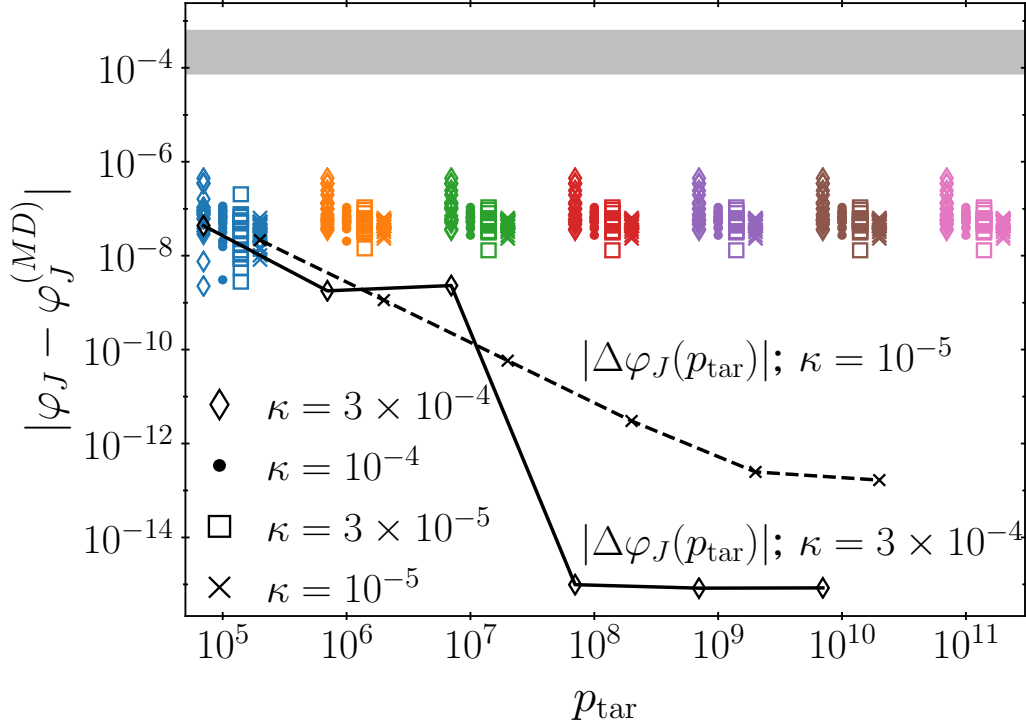


Figure 2.11. Difference in jamming densities obtained from iLP, φ_J , and fitting data of MD simulations, $\varphi_J^{(MD)}$. Markers are results obtained using 20 samples of $N = 1024$. Each marker style identifies a value of the growth rate, with the corresponding data horizontally displaced for clarity. To compare with the intrinsic variance of φ_J data of $\Delta\varphi_J(p_{\text{tar}})$ is also included. See text for more details.

Let me then analyse in detail the jammed states produced with iLP, for which micro- and macroscopic data are readily accessible. Recall that results of the previous subsection reveal that, when using the two smallest values of κ in the $N = 1024$ systems, an increase in p_{tar} is accompanied by a measurable change in the jammed state reached through iLP. Such change is detectable both by the values of φ_J and of the similarity of contact networks (see Figs. 2.8 and 2.9). Nevertheless, when larger values of κ were used, these differences persist up to a certain threshold pressure, $p_{\text{tar}}^{(\text{th})}$, above which all jammed states are virtually identical. That is, for $p \geq p_{\text{tar}}^{(\text{th})}$ it happens that $\phi_J(p) = \phi_J(p_{\text{max}})$ and $n_{ij}(p) = n_{ij}(p_{\text{max}})$ for all configurations, at least within numerical precision. It is natural to assume that $p_{\text{tar}}^{(\text{th})} = p_{\text{tar}}^{(\text{th})}(\kappa, N)$, with a larger threshold pressure associated to a smaller growth rate and larger N because the hierarchical structure of the FEL should be realized in the thermodynamic limit with an infinitely slow compression. For instance, our simulations indicate that for $N = 1024$ systems $p_{\text{tar}}^{(\text{th})}(\kappa = 3 \times 10^{-4}) \simeq 10^8$, $p_{\text{tar}}^{(\text{th})}(\kappa = 10^{-4}) \simeq 10^{10}$, and $p_{\text{tar}}^{(\text{th})} > p_{\text{max}} = 10^{11}$ for slower compressions. We can then conjecture that for any finite κ , there should be a threshold pressure above which iLP jammed states become identical, even using a *microscopic* characterization. This latter claim is further supported by the results presented soon below in Figs. 2.13 and 2.14.

The scenario that emerges from the discussion so far is that jammed states potentially reached by pure MD compression as well as those generated using iLP are different from the “thermodynamic” jammed state (φ_J^*) obtained via, say, adiabatic compression. The full picture is outlined in Fig. 2.12 in the $(1/p, \varphi)$ plane close to jamming, *e.g.* $p \gg 10^3$. The thermodynamic EOS is shown in black, while the path followed by MD compression is blue coloured. κ controls the relative slope between these two curves, which end at φ_J^* and $\varphi_J^{(MD)}$, respectively, with $\varphi_J^* > \varphi_J^{(MD)}$. Now, we have already seen that if p_{tar} is deep in the glass phase but not extremely high (say $p_{\text{tar}} \leq 10^5$), different initial conditions in the iLP algorithm lead to different packings; this is illustrated by the green paths, ending in $\varphi_{J,1}$ and $\varphi_{J,2}$. However, when the initial pressure is very large ($p_{\text{tar}} \geq p_{\text{tar}}^{(\text{th})}$), different values of p_{tar} lead to the same final jammed state $\varphi_{J,3}$ (red), even in its microscopic structure. Of course, the value of such $p_{\text{tar}}^{(\text{th})}$ depends on the value of κ ; cf. the cyan and green curves in Figs. 2.8 and 2.9. In most of the cases, $\varphi_{J,1} < \varphi_{J,2} < \varphi_{J,3}$, while the uncertainty in the estimation of $\varphi_J^{(MD)}$ does not allow to derive a similar relation with respect to $\varphi_{J,3}$. Nevertheless, it is likely that $\varphi_{J,3} = \varphi_J^{(MD)}$. Besides, given that the iLP algorithm works by crunching a given configuration, and in doing so it necessarily takes it out of equilibrium, it is also expected that $\varphi_{J,3} < \varphi_J^*$. That is, because iLP functions by applying an immediate quench to $T = 0$ and then bringing the configuration to a local minimum, essentially without going over barriers in the FEL (see Sec. 2.1.3), it is reasonable that a smaller density is achieved, but given that no precise calculation of φ_J^* is available for monodisperse HS, I could not verify this explicitly. In any case, note that this picture mimics the very recent findings of *algorithmic* memory formation in jammed packings reported in [202].

To proceed with the microscopic characterization of the iLP algorithm, it is worth to analyse the particles rearrangements performed by the iLP algorithm as the jamming point is reached. If the scenario of Fig. 2.12 is true, for sufficiently high p_{tar} we should observe that the particles displacements $\{\mathbf{s}_i\}_{i=1}^N$ become (roughly) independent of the target pressure of the initial condition. To test such hypothesis, I will consider the distribution of LP displacements, $\{|\mathbf{s}_i|\}_{i=1}^{N'}$ (*i.e.* without including rattlers) mainly during the first and last iterations. For later use, I will denote as n the number of LP optimizations carried out to reach the jamming point of a configuration. (The dependence of n on p_{tar} and κ is postponed to Fig. 2.16 in Sec. 2.3.3.) Given that $\langle n \rangle \approx 3$ for $p_{\text{tar}} \geq 10^7$ and $\kappa \leq 10^{-4}$, in practice only the first and last LP steps can be consistently compared in all the samples. Now, we need to distinguish between these two steps because, as argued in Sec. 2.1, our iLP algorithm quickly reduces the free volume (*i.e.* the feasible region) of the configuration. It is hence expected that the largest displacements take place during the initial steps. In contrast, just before convergence iLP performs displacements that are, very likely, employed to “fine tune” the contacts (for instance, by matching linear and real contacts, identified by tangency of constraints; see Sec. 2.1.3), but not on relatively large rearrangements. Furthermore, from the examples given in Sec. 2.1.3 about how the feasible region is monotonically reduced, we can reasonably assume that during the iLP process, configurations do not really jump over free energy (or actually, entropic) barriers. With all this in mind, if initial displacements are similar for very large p_{tar} , while the last ones are minor “adjustments” (for all values of p_{tar}), we can




Figure 2.12. Sketch of iLP and MD paths to jammed states, identified by their density. The equilibrium line obtained by adiabatic compression (solid black line) is followed closely by the MD compression protocol (blue), but due to the finite growth rate, it eventually detaches from the equilibrium line. Its $p = \infty$ extrapolated value is identified by $\varphi_J^{(MD)}$ (although it is possibly not a strictly jammed state; see text), using a big circle to highlight the uncertainty in its estimation. When p_{tar} is high, but not too much, using the configurations from MD as seeds of the iLP algorithm leads to different packings (green) with different packing fractions $\varphi_J^{(1)} < \varphi_J^{(2)}$. However when p_{tar} is very large, different initial conditions lead to the same jammed state (red), $\varphi_J^{(3)}$. Yet, an accurate comparison of $\varphi_J^{(MD)}$ and $\varphi_J^{(3)}$ with the available data is impossible given the uncertainty of the former. In any case, they are likely smaller than the jammed density estimated of an adiabatic protocol, φ_J^* . The values of p_{tar} included are representative of the samples I consider here with $\kappa > 10^{-4}$, and therefore are not universal.

conjecture that iLP works by crunching the configuration to the nearest minimum. Happily, the histograms of $|\mathbf{s}_i|$ in the central and uppermost panels of Fig. 2.13 agree precisely with this picture. Specifically, the effect of κ (central panel) is only visible for the first LP step, while if $p_{\text{tar}} \geq 10^7$ a small to negligible effect is observed even in such initial part (top panel). On the other hand, for the last LP iteration, there is essentially no difference when κ or p_{tar} are changed. The only visible effect is that the amount of null displacements, which for convenience are drawn around 10^{-18} , increases monotonically with κ . This means that when the initial condition of iLP is better equilibrated –*i.e.* when a smaller κ is used– more contacts can be fine tuned simultaneously. In contrast, if a faster compression is employed a larger portion of

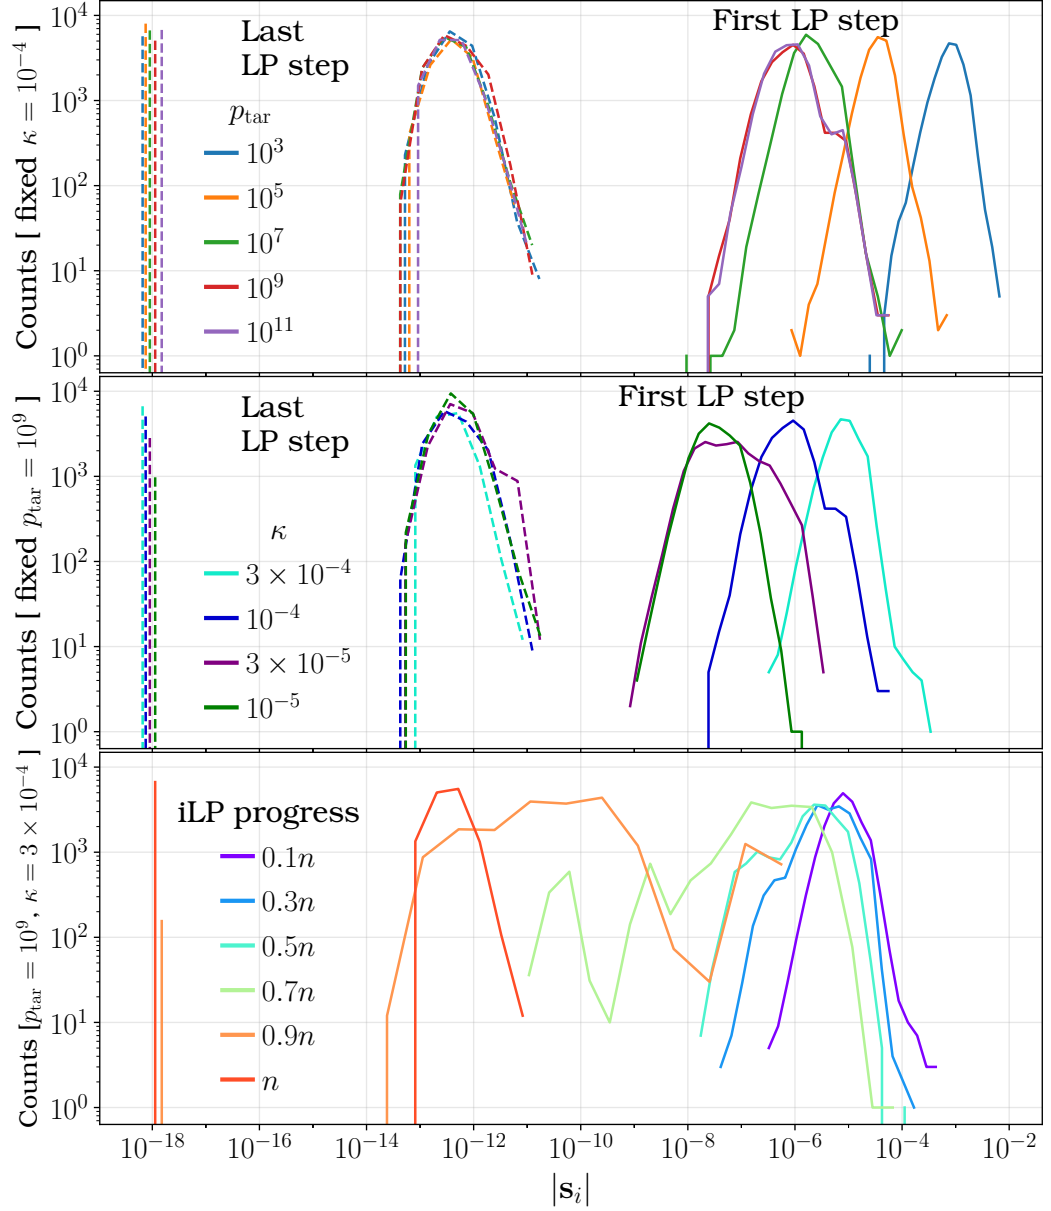


Figure 2.13. Histograms of individual particle displacements, $\{|\mathbf{s}_i|\}_{i=1}^{N'}$ (*i.e.* excluding rattlers), at different steps of our iLP algorithm. Results changing the target pressure (resp. growth rate) are shown in the top (resp. middle) panel, for the first (solid lines) and last (dashed) LP steps. In order to also include the case of $|\mathbf{s}_i| = 0$, such null displacements are drawn approximately at 10^{-18} . The bottom panel instead shows the change in the distributions of $|\mathbf{s}_i|$ at different steps of the iLP crunching (with κ and p_{tar} as indicated on the left). These results show that our LP algorithm brings a configuration to its closest minimum as discussed in the text.

the configuration remains fixed, while the final missing contacts are determined.

To further explore the mechanisms of the iLP algorithm during its convergence process, we can consider the configurations obtained with $\kappa = 3 \times 10^{-4}$, given that in

such case $\langle n \rangle \approx 10$ for the range of pressures $p_{\text{tar}} \geq 10^7$. The results are presented in the bottom panel of Fig. 2.13, whence we conclude that during the first half of iterations the displacements are practically of the same size, while it is only in the second half that a significant fraction of particles remain rather fixed. Besides, in the intermediate steps the distribution of $|\mathbf{s}_i|$ broadens across several decades, while during the initial and final steps it is mostly peaked. This provides evidence that the iLP algorithm works by first converging to a small region of a meta-basin (by relatively large displacements) and then realizing progressively the network of contacts (by small but broadly distributed displacements).




Figure 2.14. MSD for the iLP algorithm (defined in the text), for different growth rates and computed after the first, second, and last LP optimization step.

Finally, it is also interesting to keep track of the full configuration motion during the iLP process, so let me now analyse the MSD at each iteration, that is $\Delta_{\text{LP}}(k) := \frac{1}{N'} \sum_{i=1}^{N'} |\mathbf{r}_i(k) - \mathbf{r}_i(p_{\text{tar}})|^2$, where $k = 1, \dots, n$ and $\vec{\mathbf{r}}(p_{\text{tar}})$ is the configuration obtained from MD at a given target pressure. The values of this quantity for the different parameters considered here are reported in Fig. 2.14 for the first, second and last iterations. The data confirm the discussion of the previous paragraphs about the fact that for a given κ , there exists a target pressure above which the rearrangements of the configurations are essentially unchanged. Moreover, even for very small κ and very high p_{tar} , the first and last iLP steps have different roles: the former accommodates the particles, and the latter fixes the contacts. Put together, these two features suggest that for any finite growth rate and even in the $p_{\text{tar}} \rightarrow \infty$ limit of MD, some rearrangements are needed to obtain a proper jammed 1SS state. To recap: in combination with the results above, these findings further support the claim that iLP states reached from $p \geq p_{\text{tar}}^{(\text{th})}$ are identical, both from a macroscopic point of view (*i.e.* using φ_J) or from a microscopic one (characterized

through $\{|\mathbf{s}_i|\}_{i=1}^{N'}$ or $\overline{n_{ij}(p_{\text{tar}})n_{ij}(p_{\text{max}})}$). It is also worth recalling that $\varphi_J^{(MD)}$ apparently converges to the same limit value than iLP for sufficiently large pressure, which strongly suggests that iLP crunching is an efficient way to reach the infinite pressure limit; see Fig. 2.12. However, two sources of error should be kept in mind: (1) fitting the glass EOS to MD data yields an estimation of $\varphi_J^{(MD)}$ with a relatively large uncertainty; and (2) the configuration of the MD compression protocol in $p \rightarrow \infty$ limit might not be strictly jammed. Therefore, although we cannot ensure that both jammed states are equal, our findings suggest that they should be. In any case, a more extensive study with even slower compressions and higher target pressures should be useful to firmly confirm this picture.

As a last remark in this section, I would like to mention how iLP could be used in a future work to carefully explore the landscape structure. First, we could analyse if Δ_{LP} bears any relation with the hierarchy of meta-basins of the fractal FEL, itself characterized by the distribution of possible values of the MSD. A way to do this is to apply the MD compression up to a determined p_{tar} and then let the dynamics of the system continue its evolution but *without* compression. After a fixed number of collisions, the configurations are saved and crunched using iLP. By comparing the MSD of the actual dynamics, with the “distance” between jammed states, and also with Δ_{LP} we could further explore how iLP navigates the FEL towards a minimum. It should be interesting to find out if any of these quantities captures the jamming critical exponents. For instance, in Sec. 1.3.4 I mentioned that the plateau value of the MSD scales as $\Delta \sim p^{-\kappa}$. A similar relation could hold between Δ_{LP} and Γ . Nevertheless, note that the plateaus of Δ_{LP} that form for a given value of κ might indicate that configurations cannot probe the finer structure of a meta-basin, even as they go down (*i.e.* as p_{tar} increases) the landscape. So this proposal might not be straightforward to implement.

2.3.3 Complexity of iLP: Scaling with size

A final property I will briefly analyse is the size scaling of the iLP algorithm. I will mostly omit the MD compression part mainly because we used a well known protocol and did not implement any new features¹¹. But also because the LS compression and the iLP crunching benefit from two different and excluding computational approaches. In the former case, due to the deterministic dynamics and the serial nature of asynchronous event-driven algorithms¹², the MD compression of a single sample is not accelerated if more processors are available. (Clearly, if several systems are required, the LS compression can be trivially hastened by executing independent MD simulations in separate threads.) In contrast, the opposite is true for the iLP part, at least using the Gurobi solver^[199] as we did here¹³. The reason being that

¹¹Except possibly for the MD simulations of the MK model, but we will consider its full characterization in a future work

¹²Although see [203] and references therein for clever parallelised implementations, some allowing scaling of $\mathcal{O}(N)$.

¹³We decided to use Gurobi because a free academic license is available, which allows to solve large-scale optimization problems (other free licensed software limits the number of variables+constraints to about 3000) and is definitely faster than other freely available libraries, like GLPK or HiGHS; some benchmarks are available [here](#). The obvious drawback is that, being a proprietary software, it works as a “black-box” with few tunable parameters. That is, the user can

the optimization of each LP step in Algorithm 1 is done using a *non-deterministic* algorithm that runs multiple instances of the interior-point method^[190,191,193] with different random initial conditions. The solution reported is the first that converges within a user defined accuracy. Hence, the larger the amount of different initial conditions the faster the algorithm, so the solver works by implementing as many as available processors. (Unfortunately, however, the speed gain in increasing the number of initial conditions is not substantial: doubling the threads is far from halving the required computational time.) On the other hand, the alternative of using a “single-thread” solution technique, such as the simplex method¹⁴, in order to maintain the serial property in the full MD+iLP method, severely slows down the LP optimization. We therefore decided to focus mostly on the iLP part (which is the real novelty of our approach) and just consider briefly the execution times of the *full* MD+iLP procedure as I now discuss.

In tune with the results of the two previous parts, I will also explore the dependence of the running time (t) as a function of p_{tar} , for different values of κ and N . All the results that follow were obtained running the simulations in a 6 cores computer, with processor Intel Core i7-8700 at 3.2 GHz. In the iLP part, each LP optimization was solved using the barrier method as implemented in Gurobi version 8.1, with at most one initial condition per core. The statistics of t are summarized in Fig. 2.15, where mean running times of the full MD+iLP are shown (circular markers) and compared with the average running time of the iLP part only (crosses, dotted lines). The values reported correspond to the average over the same samples as above, and the error bars are their standard deviation. Left panels show the influence of the growth rate for fixed system size, $N = 1024$; while the graphs on the right explore the effects of the system size for fixed $\kappa = 10^{-4}$. Additionally, the bottom panels plot the same quantities, but considering that several MD processes can be executed simultaneously in a single computer. In particular, for these plots I assumed that 10 configurations are produced in parallel, which corresponds to the most common situation when these data were generated.

Let me begin by analysing the effects of the growth rate in the small systems. As expected, changing κ mostly affects the duration of MD, while its impact on the iLP times is very small, especially for $\kappa < 10^{-4}$. This feature is also reflected in the number of LP iterations (n) performed by our crunching algorithm as reported in Fig. 2.16 below. Yet, the most notable effect is that the time needed by iLP is roughly constant for $p_{\text{tar}} \geq 10^7$, as expected from the fact that n remains essentially unchanged for such large pressures, and confirms the features summarized above when discussing Fig. 2.12. Now, given that decreasing κ obviously lengthens the MD simulations, a trade-off between a large (but not huge) value of p_{tar} and a small growth rate is expected. This would be clearly signalled by a convex curve of the MD+iLP time, with the minimal value of t being rather unaltered by changing κ . However, this is not the scenario observed for the $N = 1024$ –where t increases monotonically with both p_{tar} and κ . This is easily explained because the

specify the desired accuracy for the solution (within a given range), the solving method, etc. but cannot access the real algorithmic implementation.

¹⁴Most accessible implementations of the simplex method –both in its primal or dual variants– are serial, although some parallelised versions have been recently developed^[204]. However, the most performant ones are proprietary and non-free; see footnote 13.

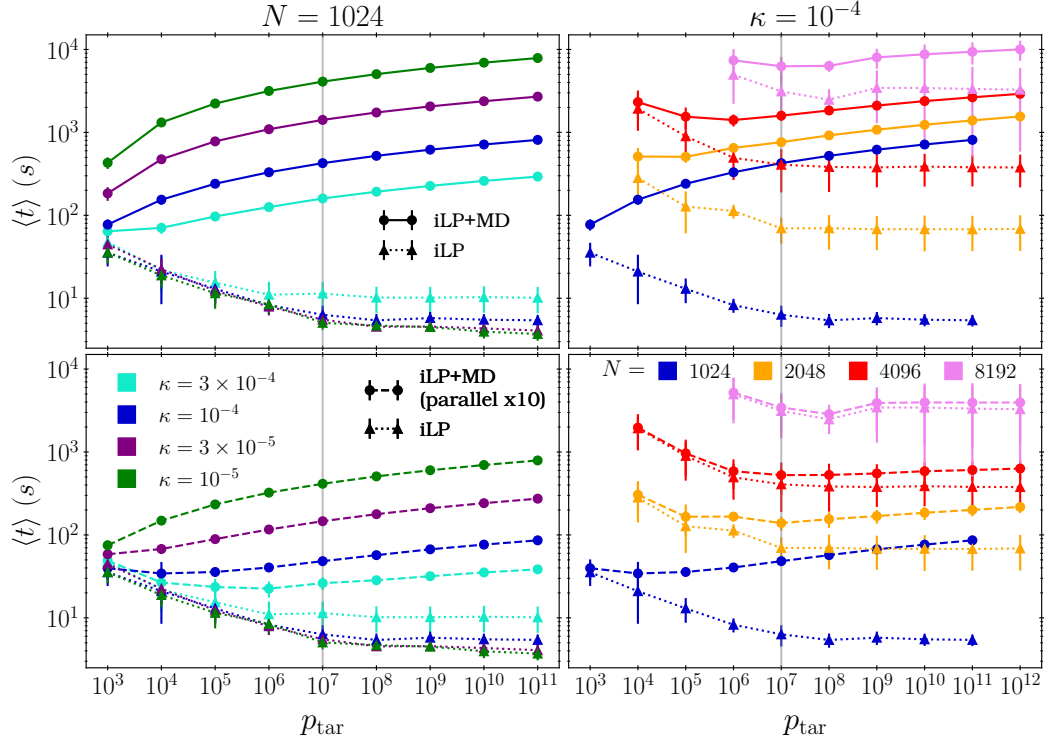


Figure 2.15. Execution times of the full MD+iLP process as a function of p_{tar} using $N = 1024$ and different growth rates (left panels) and for fixed $\kappa = 10^{-4}$ and varying system sizes (right). Upper panels show the total time of MD+iLP (solid lines), while the lower plots show the same quantity, but assuming that 10 MD samples can be generated simultaneously. In all cases, the iLP times have also been included (dotted, cross markers) for comparison. The vertical grey lines indicate the value of p_{tar} employed for most of the results presented in this section. For relatively small systems and $p_{\text{tar}} \geq 10^6$, it is clear that the most time consuming part is the MD compression. For larger systems instead, iLP always takes a considerable portion of the time. See text for more details.

iLP algorithm is very fast at crunching small systems, so extending the MD compression in order to reach a larger p_{tar} is in general a bad strategy. Of course, this conclusion only applies to the target pressures considered here: if iLP is used from a liquid configuration (*e.g.* using $p_{\text{tar}} \simeq 10$) this would not be true and the value of t would be significantly larger. Interestingly, for small systems iLP is fast enough that even when considering the simultaneous MD compression of samples, only a small dip in the curves of t can be appreciated. And it occurs at a relatively small target pressure, $p_{\text{tar}} \simeq 10^5$, and large values κ . Thus, moderate target pressures and fast compressions are optimal when jamming small systems. In turn, when dealing with large systems increasing p_{tar} is more beneficial the higher N is. Once again, however, the gain in speed of the iLP algorithm is capped for very large pressures, $p_{\text{tar}} \geq 10^9$, at least for the sizes considered here. Moreover, in the scenario of big configurations, it is clear that a significant amount of the computation time is caused by iLP. This is specially important when considering the possibility of simultaneous MD runs, as shown by the lower right panel of Fig. 2.15. This plot evinces that

for $N \geq 4096$ the major contribution to t is precisely the iLP crunching. Note also that, for fixed N , varying p_{tar} across 8 orders of magnitude rarely increases t by more than a decade. However, changing the system's size can lead to an increase in t (resp. t_{LP}) of almost two (resp. more than two) orders of magnitude, which means that reducing t of large samples should be favoured over small ones. Put it simply, it is unfeasible to choose values of (κ, p_{tar}) that are optimal for *any* given system size. This is the reason why in our analysis of the complexity of the iLP algorithm as a function of N we used a fixed value of $p_{\text{tar}} = 10^7$ (grey vertical lines) and just tried two different MD compression protocols as described below. This value of p_{tar} , although suboptimal for small samples, provides a good trade-off when using several values of N .

I will henceforth restrict the analysis of t , iterations, etc. to the iLP algorithm exclusively. From the considerations of the first paragraph of this subsection, it would seem that the characterization I will present here heavily depends on the specific solver we used. Fortunately, the situation I described there is rather general, *i.e.* interior point methods are usually faster than simplex-based ones (see, however, footnote 14). Additionally, the trend we found suggests that the complexity of the iLP jamming algorithm has a well defined dependence with the system size. A final technicality to consider when interpreting the results is that, when using an interior point method, the convergence rate to a solution within a given accuracy also depends on the system size^[190,193]. Here, we set the (absolute) tolerance for optimality and feasibility to 10^{-9} , the most accurate option available with Gurobi. It is also worth noticing that our algorithm is able to generate jammed packings with such relatively low accuracy, given that other methods require quad-precision to avoid over-shooting the jamming point^[150,151]. All in all, this means that the convergence times reported next and related quantities are good proxies of the real algorithmic complexity of iLP, but the precise values should be taken with a grain of salt.

The first feature to analyse is the number of LP steps (n) needed to reach the jamming point and its dependence on p_{tar} and κ . Analysing n complements the results discussed above regarding Fig. 2.15, as well as those of Secs. 2.3.1-2.3.2. The statistics of n are reported in Fig. 2.16, using data from the same configurations considered above. In agreement with the results presented previously, there is a negligible variation for $p_{\text{tar}} \geq 10^7$. In contrast, for fixed $N = 1024$, $\langle n \rangle$ barely changes for $\kappa \leq 10^{-4}$, indicating that the growth rate has a smaller influence on n than on other quantities. Instead, when different values of N are used, the results confirm the expected outcome that $\langle n \rangle$ increases monotonically with the system size. Note also that for sufficiently large p_{tar} , the number of LP steps remains practically constant and is never smaller than 3. When considered together with the results of Figs. 2.13 and 2.14, these data suggest that, for any $p_{\text{tar}} < \infty$ and finite κ , the iLP method requires at least two steps in order to reach the jamming point. In the first one, particles undergo relatively large displacements that define an inherent structure, while in the second step they are minimally displaced in order to turn linear (*i.e.* non tangent) contacts into physical ones (Sec. 2.1.3).

The next feature to consider is the size dependence of n and the convergence time τ . The results that follow were obtained fixing $p_{\text{tar}} = 10^7$ and testing two different compression strategies. In the first one, a constant growth rate $\kappa = 10^{-5}$ was used




Figure 2.16. Number of LP optimization steps required by the Algorithm 1 as a function of the target pressure, varying the compression rate (left panel; 20 samples per point) and the system size (right panel; 10 samples for $N \geq 2048$). Data obtained from the same configurations of Sec. 2.3.1.

for all values of $N = \{2^8, 2^9, \dots, 2^{14}\}$. Notice however that such procedure keeps the *radius*' growth rate constant, but not the density one. For the latter, we must keep in mind that $\dot{\varphi} \sim N\sigma^2\dot{\sigma} \sim (N\varphi^2)^{1/3}\kappa$. Hence, if for a given density φ we want configurations of different size to have the same compression rate, we should scale $\kappa_N \sim N^{-1/3}$. Thus, in the second strategy we used as reference $N_{\text{max}} = 2^{14}$ and then scaled $\kappa_N = (N_{\text{max}}/N)^{1/3} 10^{-5}$ for smaller sizes. To perform an accurate analysis, we generated $M = 100$ samples for each value of N and for each compression strategy.

The average values of n (upper panel) and τ (bottom panel) are reported in Fig. 2.17 for all the system sizes considered and both compression procedures. Error bars correspond to the standard error, although for τ the logarithmic scale renders them invisible. The first thing to notice is that both compression strategies produce very similar results, although less iterations are obtained in general when κ is fixed for all N , as expected. Such similarity is easy to understand because given the values of N considered here, we have that $\kappa_{N_{\text{min}}}/\kappa_{N_{\text{max}}} = (2^{14}/2^8)^{1/3} = 4$. In other words, the compression rate changes at most by a factor 4, which is not enough to compensate the influence of changing the system size; recall results from Fig. 2.16. Importantly, both series of results exhibit a well defined scaling with N , with $\langle n \rangle$ scaling as $N^{0.59 \pm 0.02}$ ($N^{0.58 \pm 0.01}$) for the fixed (scaled) compression rate, while $\tau \sim N^{2.75 \pm 0.08}$ in both cases. These exponents were found by least-squares fits, which are also included in Fig. 2.17 (dashed lines). Such fits were obtained only from data of $N \geq 1024$, because finite size effects are likely to influence the results for smaller N . The reason is that, given the particles' diameter for $N = 256$ and 512 , $\ell(\varphi)$ is about $1/3$ or more of the system's length, and thus the ratio of number of constraints/number of variables is considerably higher than for larger N . Furthermore, current data are not enough to rule out a faster scaling, such as




Figure 2.17. Average number of LP optimizations (upper panel) and convergence time (lower) as a function of the system size N . Results from a constant (resp. scaled) compression rate are indicated by the blue (resp. green) markers, while dashed lines are least-squares fits obtained from data with $N \geq 1024$. Each marker corresponds to the average over 100 samples, with the error bars being the associated standard error. The orange dotted line suggests that if even larger systems were considered, a cubic scaling with N is likely to result. For small N , the large proportion of constraints with respect to the number of variables explains the deviation from the large N trend.

N^3 (dotted, orange line) for $N \geq 2048$. Nevertheless, it is expected that the real exponent should be close to 3, although larger system sizes are needed to confirm this conjecture.

As a final feature, I will consider only the time spent in solving each LP step of the iLP algorithm, t_{LP} . This is an important variable because it measures how intrinsically complex the LP optimizations are, without taking into account the computational costs of identifying the relevant constraints, incorporating periodic effects, computing the force network, etc. Fig. 2.18 shows the results of $\langle t_{\text{LP}} \rangle$ obtained using a constant κ , and distinguishing between t_{LP} of the first and last




Figure 2.18. Main: Average time of LP optimization steps. The first and last steps are considered separately, since the corresponding displacements are very different; see Fig. 2.13. The behaviour of intermediate steps is also considered by including all the LP steps to estimate $\langle t_{LP} \rangle$. The dashed, blue line corresponds to a least-squares fit and shows that t_{LP} of the three sets scales as $N^{2.3 \pm 0.1}$. Inset: average ratio of t_{LP} between the first and last optimization. The dotted line signals that such ratio is roughly 1.5 for all system sizes, confirming that essentially the same scaling with N holds throughout the iLP algorithm

iterations, (circular and square markers, respectively). Recall that in Fig. 2.13 we have already seen that the first LP steps are considerably different from the last ones, therefore there is no reason *a priori* to assume that they scale equally with N and considering them separately is thus important. Nevertheless, our simulations demonstrate (inset) that their ratio is approximately constant for all sizes, implying that they scale with the same power of N . This is further confirmed by the values of exponents found by fitting the data of $N \geq 1024$ as above. The analysis yields that the time of the first (resp. last) scale as $N^{2.22 \pm 0.05}$ (resp. $N^{2.30 \pm 0.07}$). Moreover, if all the iterations are considered, it is found that $\langle t_{LP} \rangle \sim N^{2.27 \pm 0.07}$, which closely follows the other two cases. Hence, by simply assuming $N^{2.3}$ all the data for $N \geq 1024$ can be nicely captured, as evinced by the blue, dashed line.

Putting together these results, the first important thing that emerges is that the size scaling of τ is *different* from the one naively obtained by considering $n \times t_{LP}$, although barely within the margin of error: cf. $N^{2.9 \pm 0.1}$ and $N^{2.75 \pm 0.08}$. This implies that although the most costly part is the optimization itself, the additional operations –like computing distances between particle pairs, assigning constraints, and computing the structural variables– cannot be simply neglected. However, it is likely that as bigger systems are considered, the dominant contribution will be due to t_{LP} , increasing the size scaling of τ slightly. On the other hand, I want to emphasize the empirical character of the analysis I have presented here. That is, the

scalings reported were obtained only through curve-fitting (although showing very good agreement) and not from analysing the algorithmic complexity of each part of our iLP method. Even performing a more detailed analysis of the optimization part is not trivial, mainly, due to the proprietary code we used. But even the “turning the knobs of the black-box” approach is challenging because of the several secondary steps carried out by the solver, such as presolving, Cholesky factorization, barrier cross-over¹⁵. It is thus unsurprising that when combined, the size dependence of our iLP algorithm is not necessarily a rational power-law. A final remark is that the role of dimensionality has been completely neglected in this characterization, and although we tested our algorithm in up to 5 dimensions, we leave the corresponding complexity analysis as a topic for future research. Nevertheless, some important properties can be anticipated. First, the number of contacts increase linearly with d , as stated in Eq. (1.62) and discussed in Sec. 2.1.2, which means that complexity of LP optimizations is at least linear in d . However, not only touching particles are considered when introducing in the **Constrs.** step of Algorithm 1, but any potential contact within $\ell(\varphi)$ must be included. Unfortunately, the contribution of near contacts scales much faster with d , and results in $d = 4 - 6$ suggest that their amount is close to the kissing number^[6,94], which grows exponentially with dimensionality. It is thus likely that LP-based methods become very time consuming in higher dimensions.

2.4 Conclusions and future work

In this chapter I have presented in detail (Sec. 2.1) the iterative Linear Programming algorithm we developed to produce jammed packings of hard spheres. In contrast with other algorithms where an effective potential is introduced, or the impenetrability of particles is somewhat relaxed, our method does not use any of those techniques while maintaining the non-overlapping constraint between particle pairs at all times. I showed that upon convergence, the method always produces stable packings (Sec. 2.1.2), and it is empirically found that the vast majority of configurations have a single state of self-stress. It is a robust algorithm capable of producing packings of polydisperse particles in different dimensions without any major change. Besides, models with a modified distance function (like the MK one, Eq. (2.18)) can be trivially incorporated.

Furthermore in Sec. 2.3 I showed that our iLP method can be readily complemented by other compression protocols (such as the Lubachevsky–Stillinger protocol) that are faster for obtaining a very compressed glass, but fail in producing a proper jammed state. Under this scenario, I studied the scale with size of our algorithm (Sec. 2.3.3) and obtained that the time needed to generate a packing of size N scales as $\mathcal{O}(N^3)$ at most.

On the other hand, the results I have presented in this chapter motivate some interesting topics to investigate further. In the first place, the discussion of Sec. 2.3.1

¹⁵Each of these subroutines is almost independent from the others, so analysing their complexity is beyond the scope of this thesis, but an accessible exposition of each of them can be found in [193]. See also [this tutorial](#) for a simple assessment of their computational costs specific to the Gurobi solver.

about the fractal structure of the free energy landscape should be carried out in a more systematic way, for instance, using more samples and with a more careful evaluation of the effects of the compression rate. This analysis would be complementary to the one carried out in Ref. [111], with the benefit that our method allows us to precisely define the full network of contacts. Recall that the peculiar structure of the landscape is one of the most important predictions of Mean-field theory (Sec. 1.3.4), and recent results^[3,116] support the same picture for finite dimensional models. Hence, a thorough assessment of the energy landscape’s structure, as the one I am proposing here, will serve to dispel any remaining doubts about the range of validity of Mean-field theory.

Second, in Ref. [202] it was recently suggested the existence of a Gardner *algorithmic* transition when the jamming point was approached. Similarly to the case studied here, the authors also considered hard spheres models, but in several dimensions. However, their algorithm relied on an effective potential and the associated energy minimization. It would thus be interesting to explore if the same features can be reproduced with our method, since this would imply that such transition is an inherent property of hard spheres systems and not of the specific algorithm used.

Finally, using our algorithm we could also investigate the density of states (Sec. 1.4.5.1) of glass formers in the *under*-compressed phase. As discussed above, the spectrum of disordered solids is an important quantity because it is connected with several properties of interest, and it is known to display a critical behaviour as the jamming point is approached. However, current studies are restricted to the over-compressed regime because they usually rely on computing the Hessian, which in turn depends on the interaction potential. Because of the non-singular interaction of hard spheres, these systems have not been studied with as much detail, except possibly in [59], albeit also by introducing an effective potential. With our method both difficulties can be overcome by considering the set of *linear* contacts at each iteration of our algorithm. More precisely, recall that at each LP step we can obtain the active dual variables, and that they always satisfy the “mechanical equilibrium” condition (Sec. 2.1.2), implying that they act as “contact forces”. Strictly speaking, only at jamming these variables correspond to the real forces, however for slightly lower densities they can be used to construct a good approximation to the contact matrix, \mathcal{S} , whence an approximated Hessian can be obtained according to Eq. (1.65). Moreover, because of the equilibrium condition, we can guarantee that such Hessian corresponds to a minimum, and therefore its normal modes should be similar to the real ones. Additionally, it could also be checked that such contact forces follow the predicted critical distribution, Eq. (1.68). In this way, we could extend the analysis of the density of states to real hard spheres systems.

Chapter 3

Inferring the particle-wise dynamics near the jamming point

In this chapter I will address the issue of using the structural information contained in the network of contact forces at jamming to analyse the dynamics that takes place near such point. Most of the results have already appeared in Ref. [1]. As discussed in Sec. 1.1.1, finding a structure-dynamics connection in glassy systems is still an unsolved problem despite being a very active field of research. In fact, the intricacy of the problem is well exemplified by the ample variety of proposals of methods and structural variables, often producing results that are not consistent between one and other^[68,69]. In an attempt to get rid of several difficulties addressed in Sec. 1.1.1, we opted for two complimentary strategies. First, we make use of a simple glass former: a monodisperse system composed of frictionless, spherical particles. Second, and more important, we restrict our analysis to the short-time dynamics occurring very near the jamming point. It is worth mentioning that this is a regime that has received little attention from a theoretical and numerical point of view, (Refs. [51, 205] are notable exceptions, albeit in a somewhat different scenario from the one we will consider here). This fact is even more surprising considering that experimental techniques allow to precisely probe this type of dynamics^[206–212].

As discussed at length in the first chapter, at the jamming transition the dynamics is completely frozen due to geometric frustration. Unsurprisingly, it remains very sluggish near such critical point. Nevertheless, a jammed state is uniquely determined by the particles position¹, so we expect that any possible connection between structural properties and dynamics should be more noticeable in this scenario. Hence our choice to analyse such a constrained regime. As I will show, this approach allows us to obtain a particle-wise description, both of the relevant structural quantities and of the statistics of the dynamics. Our method is rooted on the network of contact forces that is formed at jamming, whereby we construct structural variables that allow us to characterize, individually, the displacement of particles. Our description identifies the most mobile particles but also if they exhibit any preferential direction of motion, a point that has been rarely addressed.

¹This is true, at least, for frictionless spheres which is the case I will focus on.

Note that by using the network of contacts, we dispel many of the issues mentioned in Sec. 1.1.1 because it is a well defined physical quantity. Moreover, the simplicity of the structural variables we consider –see Eqs. (3.1) and (3.2) below– makes our method versatile and robust enough to infer the particles motion in detail. For instance, we are able to identify any inherent anisotropy in the particles trajectory rather independently of the magnitude of their displacement.

Before concluding this introduction, I mention that our assumption that the statics has some bearing with the dynamics is motivated by the picture of the fractal free energy landscape (FEL) of structural glasses^[111] mentioned in Sec. 1.3.4. Recall that within this framework a jammed system can be thought of as being in one of the many possible minima of the FEL, while the dynamics takes place as the system explores the associated meta-basin and possibly the neighbouring ones. Hence, it is reasonable to expect that if one of these minima is used as initial condition for the dynamics, the trajectory of the configuration as it moves in phase space should be influenced by the specific jammed state from where it departed, at least for short times.

For reading ease I summarise here the main results: the starting point is the jammed configuration of spheres shown in Fig. 3.1 and the resulting network of contact forces. With this information we construct, at the single particle level, two quantities that are physically well defined –the sum of contact vectors (\mathbf{C} , see Eq. (3.1)) and sum of all pairs of dot products between them (S , see Eq. (3.2))– and that will be shown to be related to the dynamics. Some of their properties are analysed in Sec. 3.1.2 and later in Sec. 3.5. I then continue by analysing the dynamical variables introduced in Sec. 3.2 and studying the statistical properties of particles trajectories. I emphasize that this dynamical regime has been little studied so far, specially from a particle-wise perspective (see however [211]), which justifies the need of first finding a robust characterization of the particles motion. From the results of our numerical simulations, we gather that the statistical distribution of trajectories of each particle can be succinctly described by considering the first moment of its displacement (as a vectorial quantity indicating some anisotropy in the particles motion) and its norm squared (as a measure of its mobility). A sample of these distributions is depicted in Fig. 3.6, whence we deduce, on the one hand, that some particles have preferential directions of motion, while on the other hand, there is a broad distribution of mobility values, signalling that even close to jamming particles are constrained by their local environment in a heterogeneous way. I then turn to show that there is a strong link between these dynamical features and the aforementioned structural variables computed in terms of the contact vectors; see Fig. 3.8 (resp. 3.12) for the relation with the preferential directions in the Molecular Dynamics (resp. Monte Carlo) simulations, and Fig. 3.10 (resp. 3.14) for the connection with mobilities. To further test our approach, we verified that we can also make statistical predictions of single particle trajectories in a wide variety of systems by simply ranking the particles according to their value of \mathbf{C} and S (see Fig. 3.15). Interestingly, in Figs. 3.16 and 3.17 I provide evidence that the structural information contained in the jammed configuration gets lost rather slowly and *independently* of the dynamical protocols used for the simulations and the different parameters modelling their interactions. Some remarks about the physics behind this set of results are given in Sec. 3.5, with special attention to the fact

that ignoring the forces magnitude produces more informative structural variables. I anticipate that our guiding model will be the fractal FEL picture introduced in Sec. 1.3.4 and briefly studied numerically in Sec. 2.3.1. As a further benchmark, in Sec. 3.7 we use the vibrational modes obtained from the Hessian (at jamming) as an alternative scheme to analyse the short time dynamics, obtaining negligible correlations between them. This indicates that, in contrast with our method, the normal modes description fails to capture the statistics of the single-particle trajectories in the dynamical regime we consider here. These results show that the formalism we developed, based solely on exploiting the details of the network of contacts at jamming, provides a more effective statistical inferential technique than previous works. Finally, Sec. 3.8 includes our conclusions and gives some perspectives for future research.

3.1 Statics: Physical quantities at the jamming point

In Sec. 1.4 of the first chapter I explained that isostatic jammed packings are characterized customarily by their packing fraction, φ_J , given that such value shows a very weak dependence on the (possibly random) initial condition, changes of the system's parameters, and even the specific algorithm used^[84,94,136]. Hence, the most common situation is that random packings exhibit very similar densities^[133,136], although a broad range of values of φ_J can be accessed^[132,173,187,189] by tuning properties such as the average coordination number or by seeding some regularity in the particles position^[162,163,189]. In any case, in the unbiased scenario the value of φ_J is mainly determined by the dimensionality of the system^[16,94,187] and the distribution of the particles sizes. When $d = 3$ (the case I will be concerned with in this chapter), one can safely state that typical monodisperse isostatic jammed packings have $\varphi_J \approx 0.64$.

On the other hand, from a microscopic point of view, a jammed state is fully determined by the particles positions and radii^[150], $(\vec{\mathbf{r}}^{(J)}, \vec{R}_J) := (\{\mathbf{r}_i^{(J)}, R_{i,J}\}_{i=1}^N)$. (Note that I am following the convention used so far where \vec{X} denotes the set of N values of quantity X and boldface variables are d dimensional vectors. But, in contrast with such convention, I am parametrizing the particles' size in terms of the radius instead of the diameter because (i) the radius defines the units of time in the molecular dynamics simulations; and (ii) σ will be used in this chapter to denote the standard deviation of dynamical variables.) As explained in Sec. 1.4.3, in mechanically stable configurations these quantities suffice to determine the N_c contact forces, $\{\mathbf{f}_{[ij]}\}_{[ij]=1}^{N_c}$, where $[ij]$ with $i < j$ denotes an ordered pair and is used as a contact's index. In that same section I also mentioned that a requirement for mechanical stability is that the number of contacts is greater than the number of degrees of freedom², N_{dof} . To calculate N_{dof} (Sec. 1.4.3) we should consider that we are using a system with periodic boundary conditions, and that a small fraction of rattlers –particles that do not contribute to the rigidity of the network of contacts– are present. In d dimensions, this latter type of particles are characterised by having

²This requirement in the number of contacts is enough to guarantee that particles are in mechanical equilibrium, *i.e.* $\sum_{j \in \partial i} \mathbf{f}_{ij} = 0$, where ∂i is the set of all the neighbours of particle i . $\mathbf{f}_{ij} = \mathbf{f}_{[ij]}$ if $i < j$, or $\mathbf{f}_{ij} = -\mathbf{f}_{[ij]}$ otherwise. See Secs. 1.4.3 and 2.1.2 for more details.

less than $d + 1$ contacts. Hence, if $N' \lesssim N$ is the number of non-rattlers, we have that $N_{dof} = d(N' - 1)$. In order to generate a stable jammed packing we used the iterative Linear Programming (iLP) algorithm described in Chp. 2, but without the Lubachevsky–Stillinger compression protocol³. As discussed in Sec. 2.1, our iLP algorithm allows us to easily access the full network of contacts and always produces a stable jammed state. Moreover, the vast majority of times, such state has a single self-stress (1SS), implying that exactly $N_{1SS} := N_{dof} + 1$ contacts are present. Consequently, the configurations thus produced are marginally stable^[102,104,178]. (Strictly speaking, this means that we are dealing with an *hyperstatic* configuration, since there is an extra contact with respect to isostaticity; see Sec. 1.4.3 for a detailed discussion.)

We thus produced several monodisperse configurations of $N = 1024$ spheres. Fig. 3.1 illustrates^[213] a typical realization that we employed afterwards in the dynamical simulations. In the left panel, the spheres have been coloured according to their value of S_i , defined soon below in Eq. (3.2), while the right panel depicts the resulting contact vectors between neighbouring spheres in a small region of the same configuration. Importantly, our choice of using particles of the same size is grounded on their convenience for the statistical inference we aimed to achieve. Anticipating the features to be discussed in Secs. 3.2–3.4, notice that were we to use spheres of different radii, the smallest ones would be, presumably, the most mobile ones because they would experience less collisions with their neighbours. Yet, their relatively unimpeded motion would be a consequence exclusively of their size, and therefore *independent any structural property of their vicinity*. In order to avoid this sort of trivial inference and to really isolate the contribution coming solely from the structure, we restricted our analysis to monodisperse systems, for which I will henceforth use R_J to denote the particles' radius at jamming. The configuration used here and depicted in Fig. 3.1 has a packing fraction of $\varphi_J = \frac{4\pi NR_J^3}{3L^3} = 0.635$, with $R_J = 0.0529$ being the spheres' radius once the jamming point is reached. Only 1.3% of the $N = 1024$ particles are rattlers and none have any contact forces acting on them. As we will see later, this lack of constraints causes that most rattlers are able to move more freely than the rest of the particles.

3.1.1 Critical properties of the jammed configurations

As a security check, we verified that the packings obtained with our iLP algorithm exhibit all the properties encountered before using other algorithms^[167,169,174,177] and predicted by the best current available theory based on mean field (MF)^[16,104,150]: the critical distributions of (i) forces and (ii) gaps; and (iii) a constant density of states (DOS) for $\omega \ll 1$. The importance of these properties has been introduced in Secs. 1.4.4 and 1.4.5.1. Therefore, here I will only mention that MF theory predicts, in the $d \rightarrow \infty$ limit, that small contact forces and interparticle gaps are distributed as f^{θ_e} and $h^{-\gamma}$, respectively. The respective exponents are highly non-trivial: $\theta_e = 0.42311\dots$ and $\gamma = 0.41269\dots$. Similarly, marginal stability considerations and numerical results indicate that $D(\omega) \rightarrow \text{constant}$ as $\omega \rightarrow 0$.

³The reason for not including the molecular dynamics part described in Sec. 2.2 is that at the moment the characterization of the MD+iLP algorithms was not available.




Figure 3.1. Top: Jammed configuration of $N = 1024$ spheres used to perform the dynamics simulations. Its packing fraction is $\varphi_J = 0.635$ and each sphere is coloured according to its value of S_i , defined in Eq. (3.2), from light yellow (big values) to dark blue (small, negative values); see also the colour scale in Fig. 3.6b. Bottom: Zoom in a region of the same configuration (without colouring) showing the contact vectors obtained according to the method described in the previous chapter.

Fig. 3.2 shows that our configurations indeed fulfil these properties, by comparing the empirical cumulative distributions (markers) obtained from 20 independent packings⁴ (to improve the sampling), with the theoretical predictions (solid lines in

⁴The configurations employed here, either to perform the simulations reported in the main text

each panel). In the case of the forces, however, finite dimensional systems exhibit two separate contributions, due to extended and buckling excitations; see Eq. 1.68. Only the former follows the MF distribution, while no such prediction exists for the latter. As explained in Secs. 1.4.4-1.4.5, the two types of excitations can be differentiated by the number of contacts of their associated particles. Thus, bucklers are, with very high probability, associated with particles having $q_i = d + 1 = 4$ kissing neighbours, *i.e.* the minimum number of contacts to guarantee stability⁵. Nevertheless, the associated distribution in the leftmost panel of Fig. 3.2 (identified by the blue circles) was obtained computing instead the smallest contact force per particle. This distinction is not important –because bucklers are, in any case, usually associated with small forces– but helped in our case to obtain a cleaner separation of the two contributions⁶. At any rate, the agreement between theory and numerical results is remarkable, specially considering that our systems are rather far from the limit of infinite dimensions. Moreover, several of the observed deviations from the predicted power-laws are due to finite size corrections as analysed in detail in Chapter 4.

3.1.2 Structural variables of interest

I now address one of the main goals of this chapter: how to use the network of contacts formed at jamming to construct well defined physical quantities that can be used as predictors of the dynamics that takes place close to such point. Yet, before answering this question, I want to remark why this approach differs from previous ones. Most importantly, note that by using a jammed state as the initial configuration for the dynamics our knowledge is augmented in comparison with the scenario where the initial state is only *close* to a jammed configuration. In this latter case we would be missing data about the true contacts between particles, and thus our description would be limited to the usage of “coarse-grained” variables, for instance, the local density or a pair correlation function defined within some small vicinity of each particle, not unlike methods introduced in Sec. 1.1.1. In contrast, here I will be considering a scenario where the dynamics departs from a configuration in which we can identify the actual neighbours of each particle and, consistently, only include such relevant particles in our description. This distinction is crucial for uncovering preferential directions in the particles’ displacements, as we investigate in the next sections. Moreover, because in our dynamical simulations we only consider contact potentials between the spheres, a “coarse-grained” description would presumably fail to provide a realistic picture of the main interactions driving the particles trajectories. In other words, given that we explore the particles’ motion in the vicinity of a jamming point, it is to be expected that the correlation between structure and dynamics should be stronger than in other scenarios. Hence, if we

or the ones reported in the Appendix 3.A, are two independent systems selected at random from this set.

⁵I mention in passing that although not all particles with $q_i = d + 1$ contacts are strictly bucklers, with very high probability any buckler will have $d + 1$ contacts; see Ref. [150]. I will thus abuse on terminology and probability by simply referring to any particle with 4 contacts as a buckler.

⁶In the next chapter, where these critical distributions are studied in great detail I have used the criterion in terms of the coordination number to identify the bucklers and show that they do follow the previously known power-law. See Figs. 4.3 and 4.8.




Figure 3.2. Cumulative distribution functions (cdf) of contact forces (left), interparticle gaps (central), and frequencies (right). In all cases, the black lines are comparisons with the power laws predicted theoretically or found in previous studies. For the data of panel (a), I have split the forces in the ones associated to localized (blue) and extended modes (red), using the smallest force per particle for the former type and the criterion of Ref. [150] for the latter one. The green markers are all the forces put together and the exponent of the associated fit was taken from the same reference. In panel (b) I show the cdf of the dimensionless interparticle gaps (defined in Eq. (1.49)) and that it closely follows the expected scaling, Eq. (1.67). Finally, in the last panel I used the fact that the density of states is predicted to be constant for small values of ω and hence its cdf (the integral of $D(\omega)$) should be proportional to the frequency. The spectrum, $\{\omega_i\}_{i=1}^{dN'}$, was obtained by diagonalizing the Hessian *at* jamming, following Eq. (1.65) with the harmonic contact matrix. Deviations in all of these cases are due to finite size effects as studied in the next chapter.

have at hand more accurate information about the structure, we will be able to achieve a better and longer lasting statistical inference of the dynamics.

With this in mind, I will use the network of contacts to introduce some quantities that will be used in the rest of this chapter to describe the statistics of the particles' trajectories. Following the notation from Sec. 1.4.3, $\mathbf{n}_{ij} = \frac{\mathbf{r}_i^{(J)} - \mathbf{r}_j^{(J)}}{|\mathbf{r}_i^{(J)} - \mathbf{r}_j^{(J)}|} = \frac{\mathbf{r}_i^{(J)} - \mathbf{r}_j^{(J)}}{2R_J}$ will

denote the unit contact vector, pointing from particle j towards i (assumed to be in contact). Clearly, the set of such vectors defines the edges of the network depicted in the lower panel of Fig. 3.1. Using the contact vectors, we can easily construct the following two quantities: (i) the vectorial sum of all the contact vectors acting on the i -th particle,

$$\mathbf{C}_i := \sum_{j \in \partial i} \mathbf{n}_{ij}, \quad (3.1)$$

and, (ii) the sum of all pairs of scalar products of the contacts acting on the same particle,

$$S_i := \sum_{j < k \in \partial i} \mathbf{n}_{ij} \cdot \mathbf{n}_{ik} = \frac{1}{2} (|\mathbf{C}_i|^2 - q_i), \quad (3.2)$$

where q_i is the coordination number⁷ at jamming of particle i . The colours used in the upper panel of Fig. 3.1 indicate the value of S_i for each particle of the system, with a bigger (smaller) value corresponding to a colour on the yellow (dark blue) part of the scale reported in Fig. 3.6. The complex distribution of colours displayed by the configuration resembles the ones found using other structural variables and order parameters, which in turn have been linked with the heterogeneous dynamics observed in glassy systems^[42,44,46–49,74,205].

As I will argue in the next sections, it is the contact *vectors* and not the *forces* ($\mathbf{f}_{ij} = f_{ij}\mathbf{n}_{ij}$) what actually convey more information about the dynamics of the configuration near its jamming point. But before exploring such connection, it is convenient to analyse their differences in the packings we produced. First, note that the analogous of Eq. (3.1) for contact forces trivially vanishes – *i.e.* $\sum_{j \in \partial i} \mathbf{f}_{ij} = 0$ – given the mechanical equilibrium property of jammed states. That is, it cannot provide information about the dynamics because it has the same value for all particles. The case of the scalar structural variable is different and a more careful analysis should be carried out. In this part I will show that, despite being closely related quantities, the dot products between pairs of contact forces and contact vectors have strikingly different probability distributions (pdf). In later sections I will compare their performance as predictors of the dynamics. The claim about their pdf’s is summarized by Fig. 3.3, where the data of all 20 configurations is included. Left panels depict distributions of quantities involving only contact vectors, while the right ones show the distributions of the analogous quantities using contact forces. On the other hand, the upper panels show the distribution of all dot products between pairs of particles in contact, and the lower ones their sum per particle (thus the lower left panel is the distribution of S_i). Additionally, the contributions of bucklers and particles associated to an extended response (see Secs. 1.4.4–1.4.5) have been plotted separately. Let me first consider the pdf’s in the upper panels. At a first glance, the distribution of $\mathbf{f}_{ij} \cdot \mathbf{f}_{jk}$ suggests that the vast majority of contact forces are nearly orthogonal (notice the log-scale in the vertical axis). But when compared with pdf of $\mathbf{n}_{ij} \cdot \mathbf{n}_{jk}$ we see that the likeliest arrangement is the one consisting in three spheres touching each other (hence the peak at $\cos \pi/3 = 0.5$), while any other arrangement is roughly equally likely. This in turn implies that the noticeable peak at zero in the first case is due to the existence of many small forces whose effect is to blur any intrinsic geometrical feature of the configuration’s structure. Given the few amount of bucklers (see inset in the upper left panel for the pdf of q_i), the distributions of all the particles and the ones associated to extended excitations are virtually identical. Considering the distribution of $\mathbf{n}_{ij} \cdot \mathbf{n}_{jk}$ restricted to bucklers we see a relative deficit of three mutually touching spheres, and a slight excess of triplets spanning an angle *close to* $\frac{2\pi}{3} = \arccos\left(-\frac{1}{2}\right)$. In combination with the low fraction of particles almost perfectly aligned ($\mathbf{n}_{ij} \cdot \mathbf{n}_{jk} \simeq -1$), these results support the hypothesis that bucklers are very likely particles with d almost coplanar contacts. However, if this simplified picture is entirely valid, there is a peak “missing” at $\mathbf{n}_{ij} \cdot \mathbf{n}_{jk} \simeq 0$ corresponding to the almost orthogonal force-balancing contact. The absence of such peak is likely due to the fact that the remaining contact does not always bear a small force, and

⁷In other parts of this thesis I have used z_i for the same purpose, but to avoid confusion with the z component of a particle’s trajectory, in this chapter I will use q_i .

therefore is not necessarily orthogonal. See more details in [150, SI, Sec. II].

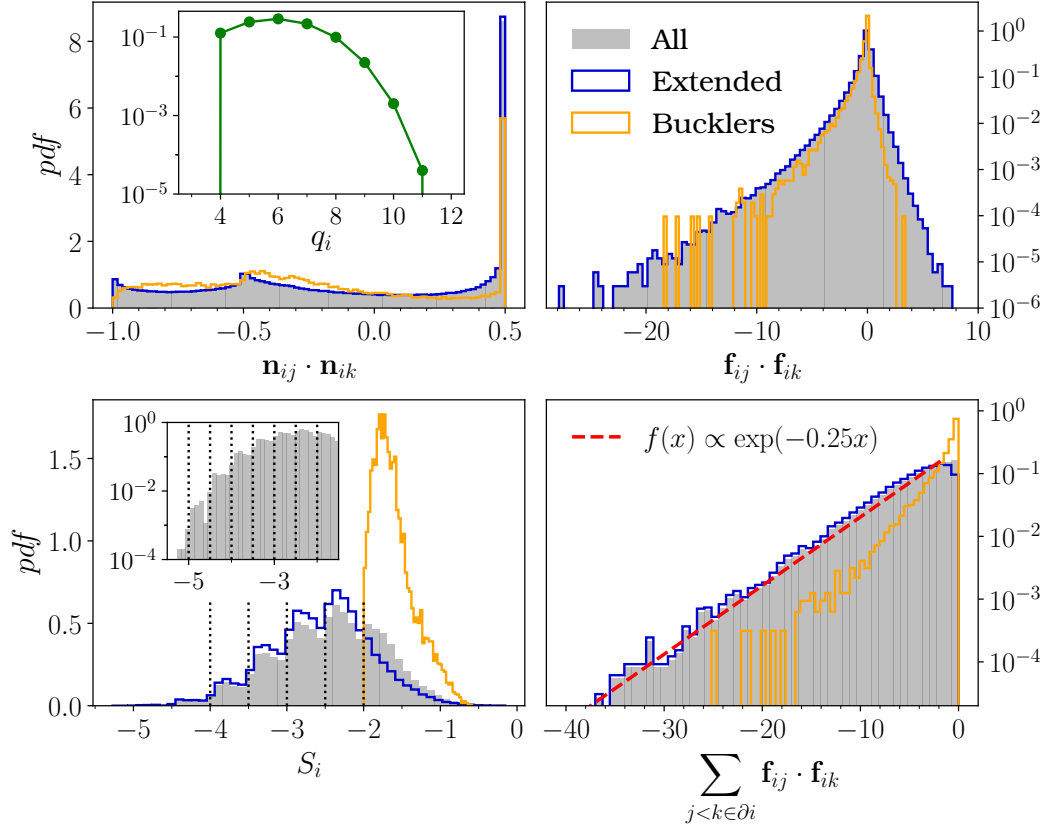


Figure 3.3. Distribution of dot products of all contact vectors (upper left panel) and contacts forces (upper right) acting on the particles and their respective sum per particle (lower panels). Rattlers have been excluded, and the separate contributions of bucklers (orange) and the rest of the particles (blue) have been also included. The inset in the pdf of S_i –defined in Eq. (3.2)– depicts the same distribution in a logarithmic scale to show that the influence of the bound $S_i \geq -q_i/2$ (dotted lines) is persistent for all possible values of q_i ; see text for more details. The pdf of q_i is plotted in the inset of the upper left panel.

Now, the sums of these dot products, *i.e.* S_i and $\sum_{j < k \in \partial i} \mathbf{f}_{ij} \cdot \mathbf{f}_{ik}$, are also distributed in a very different fashion, as shown in the lower panels of Fig. 3.3. For instance, the pdf of S_i displays an irregular shape because this structural variable is lower bounded by $-q_i/2$, as can be easily derived from its definition. This causes its value to change rather abruptly in regular intervals of $1/2$ for $S_i \leq -2$ (because the smallest number of neighbours is 4), as shown by the dotted lines. As the inset shows, this feature is preserved all the way down to $q_i = 11$, the largest connectivity observed. Importantly, in Ref. [1] this distribution was incorrectly compared to a Gaussian, because the lack of resolution (due to the small number of samples considered) prevented the discontinuities at $-q_i/2$ to be observed. In any case, the analogous sum using contact forces clearly follows a different distribution, if only by its monotonic form. This latter case shows that, not only there is an excess of

very small forces, but there are particles whose full set of contact forces is given by forces of minute magnitude. Even more, the probability of having a different value for this variable decreases exponentially, $p(x) \sim e^{0.25x}$, as indicated by the red dashed line. Finally, notice that the distribution being upper bounded at zero is a straightforward consequence of the identity

$$\sum_{j < k \in \partial i} \mathbf{f}_{ij} \cdot \mathbf{f}_{ik} = \frac{1}{2} \left(\sum_{j \in \partial i} \mathbf{f}_{ij} \cdot \sum_{k \in \partial i} \mathbf{f}_{ik} - \sum_{j \in \partial i} |\mathbf{f}_{ij}|^2 \right) = -\frac{1}{2} \sum_{j \in \partial i} |\mathbf{f}_{ij}|^2.$$

To close this section and to put into context the main results of this chapter, I emphasize that both \mathbf{C}_i and S_i –introduced in Eqs. (3.1) and (3.2)– are well defined physical variables for each particle and, consequently, our approach will provide a description of the system’s dynamics at the single particle level. We thus extend other techniques where the characterization of the particles’ displacement was done in terms of clusters or mesoscopic regions within the system, as exemplified in Sec. 1.1.1.

3.2 Dynamical variables of interest

Even though the jamming regime is by itself interesting and complicated enough, many of the most salient physical properties of amorphous solids are linked with the dynamical slowing down that takes place when, say, their density increases. So I now turn to the main part of this chapter in which we first characterise the dynamics of our configuration of spheres as it moves away from its jamming point, *e.g.*, by reducing the system’s packing fraction by a small amount and providing the particles with momenta. Then, in a second stage, we establish a connection between the sluggish motion of the configuration and its static properties computed at the jamming point.

At first sight, it might seem that making such two-step division of the analysis is somewhat redundant, since we could proceed instead by directly trying to construct some quantity using the network of contact forces and then relate it with certain dynamical features. This is the standard methodology in studies exploring the link between the local environment of particles and their square displacement, where this last quantity is always taken as a measure of the “mobility” of a particle. As mentioned in Sec. 1.1.1, this technique has been used to relate structural features of several model systems with their dynamics. Nevertheless, given that here we are considering the dynamics in a different and rather unexplored regime, we opted for first trying to answer the question of how particles move near the jamming point. In particular, do they have any preferential direction of motion? And if they do, how much do they move along it? To show that these attributes are not necessarily determined by the mobility of a particle, consider the case in which its closest neighbours are distributed uniformly around it forming a relatively large cage. It is clear that the particle will be more mobile than if the cage was small, rendering its motion very constrained. And yet, in neither of these two cases the particle’s displacement would exhibit a preferred direction. Conversely, if the arrangement of neighbours is quite irregular, the motion of the particle should be facilitated towards

the direction where fewer obstacles are present, however, this favoured path will only be discernible if the particle is also allowed to move enough, because otherwise the preferred direction will be washed out, for example, by thermal noise. These elementary examples illustrate that even a very simple study of the dynamics near jamming should take into account these two properties in order to yield a sufficiently comprehensive picture. Therefore, our approach is based on analysing the particles' displacement both as a vectorial quantity (that is, component-wise) and as a scalar one (using its norm squared, henceforth also termed “mobility”):

$$\delta \mathbf{r}_i(t) = \mathbf{r}_i(t) - \mathbf{r}_i(0) = (\delta x_i(t), \delta y_i(t), \delta z_i(t)) \quad (3.3a)$$

$$|\delta \mathbf{r}_i(t)|^2 = \delta x_i^2(t) + \delta y_i^2(t) + \delta z_i^2(t); \quad (3.3b)$$

where $\mathbf{r}_i(t)$ is the position of the i -th particle at time t . Expectedly, another variable of interest will be the mean squared displacement (MSD), introduced previously but whose expression I reproduce here for convenience:

$$\Delta(t) = \left\langle \frac{1}{N} \sum_{i=1}^N |\mathbf{r}_i(t) - \mathbf{r}_i(0)|^2 \right\rangle. \quad (3.4)$$

When it was considered in Sec. 1.1, the meaning of $\langle \bullet \rangle$ was a thermal average. In this chapter however, it will denote an average over the different trajectories generated, as explained next.

We want to probe how the configuration evolves as it explores the phase space close to the FEL minimum identified by the jammed state. It is then natural to use the jammed configuration as initial condition, *i.e.* $\bar{\mathbf{r}}(0) = \bar{\mathbf{r}}^{(J)}$, and then generate many independent trajectories. This approach is equivalent to sampling from the isoconfigurational ensemble^[62–67] (ICE), since the initial positions are kept fixed while the configuration evolves according to some prescribed dynamical protocol, following a different trajectory each time. Sampling from the ICE thus allows to uncover the contribution of the network of contacts to the dynamics because any stochastic contribution to the dynamics is expected to cancel out if the number of trajectories is large enough. With this in mind, the meaning of $\langle \bullet \rangle$ is simply the ICE average, whose value we estimate through several independent simulations of the dynamics.

In the next two sections I present the statistics of the particles trajectories in the ICE, sampled using two different dynamical schemes, namely, Molecular Dynamics (MD) and Monte Carlo (MC) simulations. In the former type, the infinitely hard sphere (HS) model is kept and the system's density is reduced with respect to φ_J . For the MC simulations we considered instead soft spheres (SS) with different interaction potentials, but with keeping the system's packing fraction constant at φ_J . The reason for using these two types of simulations is that we want to investigate the applicability of our method as we change the model system and its parameters. This is specially important because other studies^[69] using the ICE have reported that the level of correlation between local structure and dynamics is very sensitive to the glass former model and the type of interaction.

3.3 Results with MD simulations

I begin by showing the results obtained using MD simulations, for which the radius of all spheres was reduced by the same factor in order to reach a new density $0.9 \leq \varphi/\varphi_J \leq 0.999$, that was kept constant afterwards. For each value of φ we performed $M_{MD} = 5000$ independent simulations of event-driven MD as described in Sec. 2.2 and using the implementation of Ref. [94]. As mentioned above, the initial conditions were set to $\bar{\mathbf{r}}^{(J)}$ in all cases, while the initial velocities were assigned randomly according to a Maxwell–Boltzmann distribution at inverse temperature $\beta = 10$ (in the reduced units of the algorithm). As discussed in the previous chapter, this MD code uses an event-driven scheme, which implies that the ballistic regime will be mostly absent from all of our data. Additionally, in event-driven type of algorithms the natural units for the time step of the dynamics are the number of collisions per particle that have occurred. In order to make a clear distinction between this time unit and the “physical time” t in Eqs. (3.3) and (3.4) above, I will denote as τ_{MD} the number of events per particle. However, Fig. 3.4 in Sec. 3.3.1 shows that once we have fixed φ , there is a well defined relation between τ_{MD} and t . This incoming part is devoted to further details of the MD simulations that are not essential for the rest of the chapter, and can be skipped. Then, in Sec. 3.3.2 I show the statistics of the particles trajectories and how they can be captured by their moments. Finally, in Sec. 3.3.3 I show that these statistics are highly correlated with the structural variables introduced in 3.1.2.

3.3.1 Further details of the MD simulations

In Sec. 2.2 I mentioned that event-driven MD algorithms are well suited for simulating systems of HS given their trivial dynamics. With this method, the state of the system is advanced until a collision takes place, where the particles exchange momenta, and the next collision is predicted with the new velocities. In such a way, the evolution of the system is updated each time a collision event occurs and the global time t of the simulation is just the sum of the elapsed times between events. Another consequence of using the HS potential is that the only characteristic length of the system is the particles size, which in turn fixes the time units as $\sqrt{\beta m R^2}$, where m and R are the spheres mass and radius, respectively. Without loss of generality, we set $m = 1$, $\beta = 10$, while R is determined by the value of φ we chose when running the simulations. In Fig. 3.4 I show that for a fixed value of the packing fraction, there is well defined relation between the t , averaged over all the trajectories, and the value of τ_{MD} , which measures the number of collisions that have taken place. The very small error bars suggest that the distribution of times is considerably peaked around $\langle t \rangle$, so I will henceforth use t to denote its sample average.

To better characterize the dynamical regime studied here, we performed additional MD simulations (with respect to the ones used for our main results) letting the configuration evolve for 100 times as many collision events, using configurations of packing fractions $\varphi/\varphi_J = \{0.96, 0.99, 0.98, 0.999\}$ and a smaller number of trajectories, namely $M_{long} = 1000$. In Fig. 3.5 we present the results of the Δ as a function of time, for several packing fractions using the data of the main simulations




Figure 3.4. Relation of the (average) global time of the MD simulations as a function of the number of collision events, for different packing fractions (different colours). Each marker correspond to an instant at which the particles positions where stored and the solid lines are just a guide to the eye. Errorbars indicate the value of the standard deviation of the values.

(circular markers) and from the longer ones (solid lines). For reference, I included the expected trend from a purely ballistic growth (black dashed line). As explained in Chapter 1, most of the unique features of glassy systems occur after the configuration has abandoned such regime and its MSD has reached its characteristic plateau. But as this figure shows, we are exploring the dynamical evolution well before the MSD attain its asymptotic value. As a final remark, I want to stress that the ballistic regime is, expectedly, almost entirely lost due to the event-driven nature of our simulations.

3.3.2 Particle trajectories and their statistics

After generating $M_{MD} \gg 1$ trajectories in the ICE for a fixed packing fraction, we can access a statistical characterization of the particles' motion. For instance, the left panel of Fig. 3.6 shows the pdf of the $\hat{\mathbf{z}}$ component of the particles displacement (the distributions of the other components are very similar) at $\tau_{MD} = 20$ and using $\varphi = 0.995\varphi_J$. Keeping in mind that each curve is the pdf of a single particle, it is clear that all the particles have a well defined mean displacement, many of which are *different from zero*. This is the first of our main results, because it indicates that some spheres indeed have a preferred direction of motion and that it can be identified despite the statistical fluctuations coming from the thermal noise and sample-to-sample variations. More importantly, we can infer such direction using the sum of contact vectors defined in Eq. (3.1). This can be seen from the evident




Figure 3.5. Temporal evolution of the MSD for different packing fractions (different colours). The circular markers show the values of MSD obtained from the main simulations, while the solid lines correspond to simulations 100 times as longer, but for a reduced number of values of φ . The black dashed line is included as a reference for comparison with a ballistic behaviour. Note that all the results presented in the main text correspond to the evolution of the configuration much before the MSD reaches its plateau value, which is the regime that so far has received most attention.

division of colours to the left (right) corresponding to a negative (positive) value $C_i^{(z)} := \mathbf{C}_i \cdot \hat{\mathbf{z}}$, as indicated in the scale on top. Only 15% of the particles were used to generate the main figure, and for clarity reasons they were chosen as the 10% (5%) with the largest (smallest) absolute value of $\langle \delta z_i \rangle$, but no structural information nor other statistical property whatsoever was considered. Now, since all the distributions shown are unimodal and rather narrow, we can expect the value of $\langle \delta z_i \rangle$ to be descriptive enough. To test this idea, we compared the mean value and the median of each particle in the configuration and obtained the results included in the inset of the same figure, where data points of different colour correspond to the different spatial components as indicated. The fact that all of them lie very close to the identity line (black dotted curves) confirms that $\langle \delta \mathbf{r} \rangle$ can be used to discern the existence of preferential directions in the possible trajectories of the configuration. I note in passing that the presence of anisotropy in the particles' motion has only been recently studied in few works, *e.g.* [54, 214–217].

Analogous results for the particles' square displacement are shown in the right panel of Fig. 3.6, also selecting 15% of the particles according to their mobility using the same criterion as before. Correspondingly, the colour scale indicates the value of S_i defined in Eq. (3.2). Importantly, in this case the shape of the distributions is

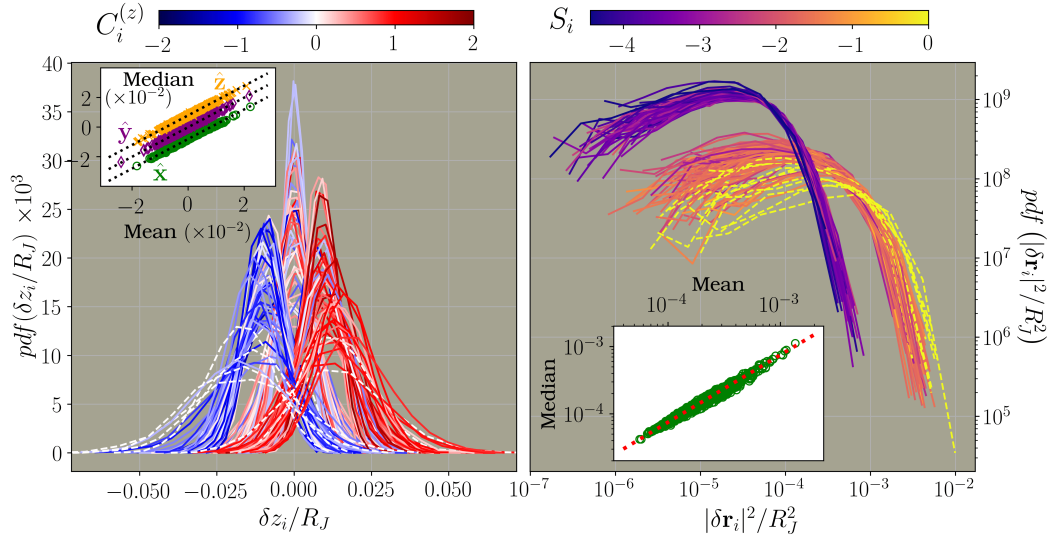


Figure 3.6. Main figures: Probability distributions (pdf) of the variables of Eqs. (3.3) at a fixed time $\tau_{MD} = 20$ and packing fraction $\varphi/\varphi_J = 0.995$, obtained from $M_{MD} = 5000$ independent trajectories sampled from the ICE. Each curve corresponds to the pdf of a single particle, but only 15% of the configuration is shown as described in the main text. Left panel: pdf of particles' displacement along the $\hat{\mathbf{z}}$ direction (similar results hold for the other directions). Right panel: pdf of their square displacement. Curves colours correspond to the particles value of $C_i^{(z)}$ (left) or S_i (right) as indicated in the colour scales on top, while dashed lines are used to identify if the curve is associated with a rattler. Insets: Correlation between mean and median of the distributions. In the left one, the points of each set have been displaced vertically by a fixed amount for clarity reasons and the identity lines (also displaced by the same amount) are shown in black dotted, while in the right panel's inset the red dotted line corresponds to $f(x) = 0.75x$.

notoriously different with respect to the ones of $\delta \vec{\mathbf{r}}$. This is exemplified by the fact that the mean and median of the mobility do not coincide in this case (see the inset) because the mean is a statistic more sensible to extreme events, such as trajectories where the mobility can be about 10 times larger than the most typical ones. Nevertheless, there is a simple linear relation between both statistics, as illustrated in the inset where the red dotted line indicates the curve $f(x) = \frac{3}{4}x$. In summary, the distributions of $\{\delta \mathbf{r}_i\}_{i=1}^N$ and $\{|\delta \mathbf{r}_i|^2\}_{i=1}^N$ have clearly a different shape, with unimodal, well peaked curves for the former, and broader, asymmetric probability densities in the latter case. Such difference evinces that these two variables have different behaviours and should therefore be studied separately as we do here. I also remark that, by the same token, it is not obvious that the ICE mean of a particle's mobility, *i.e.* $\langle |\delta \mathbf{r}_i|^2 \rangle$, would be a good descriptor of its full distribution. Therefore, to justify the usage of such average to characterize the statistics of $|\delta \mathbf{r}_i|^2$, I provide the following reasons: i) the support of the distributions shown in the figure differs, roughly, by an order of magnitude, which is enough to tell apart the least mobile particle from the most mobile ones; ii) the inset shows that there is still a linear relation between the median and the mean, implying that a higher value of the latter is accompanied by a proportional translation of the full distribution also to higher

values; iii) other studies using the ICE have found that several dynamical features worth studying in glassy systems are indeed captured by $\langle |\delta \mathbf{r}_i|^2 \rangle$. Additionally, a word of caution is in place to avoid possible confusions regarding how to interpret the mobility distributions. First of all, the first moment of $|\delta \mathbf{r}_i|^2$ in general can be different to its typical values, as has already been reported in [68]. Additionally, the grouping of the curves depicted in the right panel of Fig. 3.6 in two different sets should not be understood as if there were two limiting distributions of the mobilities, but instead as a consequence of the fact that we selected only 15% of the particles trajectories; that is, the remaining 85% of the curves fill the gap between the apparently two different sets.




Figure 3.7. Comparing the average displacement and its standard deviation along the $\hat{\mathbf{z}}$ component of the full configuration, at four different times as indicated by the label on the bottom of each panel. The colour of each point corresponds once again to the scale of $C_i^{(z)}$ as in the previous figure, while rattlers are identified by crosses. In cyan I have included the \pm identity lines for a better comparison, so that whenever a point lies beneath them it means that the mean dominates over sample-to-sample fluctuations. These results also demonstrate the clear correlation between $\langle \delta \mathbf{r} \rangle$ and $\bar{\mathbf{C}}$, but the lack of it for $\sigma[\delta \mathbf{r}]$. This is analysed in detail in later sections.

It is worth analysing more closely the variation around the mean displacements with the purpose of verifying that preferential directions can be identified despite sample to sample fluctuations. To this end, in Fig. 3.7 I compare the expected value of the displacements, $\langle \delta \vec{\mathbf{r}} \rangle$, to their standard deviation, $\sigma[\delta \vec{\mathbf{r}}]$, at different times (as indicated by the label of each panel). Without loss of generality, I once again restrict the analysis to the $\hat{\mathbf{z}}$ component. These plots include the statistics of the full configuration, with the rattlers identified by crosses. Anticipating the results of the next section, the colour of each point is associated to its value of $C_i^{(z)}$, using the same scale as in Fig. 3.6. The first thing to notice is the very high correlation between $\vec{\mathbf{C}}$ and $\langle \delta \vec{\mathbf{r}} \rangle$ at short times (upper panels). However, no such correlation is present for $\sigma[\delta \vec{\mathbf{r}}]$, as will be further analysed in the next part. More importantly for the purposes of this part, except for the cases for which $\langle \delta \mathbf{r}_i \rangle \approx 0$, the mean displacement is of the same order (and sometimes even larger) than the corresponding $\sigma[\delta \mathbf{r}_i]$; *cf.* with the \pm identity lines included in all the panels. This implies that if particles move, on average, predominantly along a given direction, different realizations of the dynamics will exhibit a bias for similar trajectories. These results also show that for longer times (lower panels) preferential directions are better defined, as evinced by the larger fraction of points below the \pm identity lines. Unfortunately, this temporal regime coincides with the one where the dynamics has lost most of its initial structural information. This is exemplified by the mixture of colours, but below I provide a more quantitative characterization.

The results of this part constitute our first main finding. Namely, that near the jamming point, the trajectories of individual particles show preferential directions of motion. During the very initial times, such preferred directions might be hard to spot due sample to sample fluctuation as well as the small value of $\langle \delta \vec{\mathbf{r}} \rangle$. But as time passes by, an increasing number of particles have a mean displacement such that $|\langle \delta \vec{\mathbf{r}}_i \rangle| > \sigma[\delta \mathbf{r}_i]$, implying that the mean value becomes the dominant one and most of the trajectories are similar to the preferred one. Similarly, statistics of the particles' mobility are well described by their mean and, actually, in this case it is customary that $\langle |\delta \mathbf{r}_i|^2 \rangle > \sigma[|\delta \mathbf{r}_i|^2]$ even for short times. Hence, the characteristic mobility of particles is a persistent feature across different samples, despite the fact that its average might not coincide precisely with a typical value.

3.3.3 Correlation with structure

I now turn to the question of finding how much information the jamming point contains about the dynamics close to it. This means that the correspondence “similar local environment leads to similar dynamics”, illustrated in Fig. 3.6 for few particles (by the clustering of distributions of similar colours), must now be extended to the full configuration. A convenient way to achieve this is to use just few statistics of the distributions, instead of the complete pdf. As argued above, the first moment comes across as a quantity descriptive enough, and the results of Fig. 3.7 support this approach. However, to perform a more quantitative analysis and assess whether there is an underlying connection between $\langle \delta \vec{\mathbf{r}} \rangle$ and $\vec{\mathbf{C}}$, Fig. 3.8 shows the scatter plots of these two quantities for the N particles. Each of the three rows in the figure corresponds to a single direction, while the different columns refer to the different times of the dynamics at which the mean displacement was calculated. Once

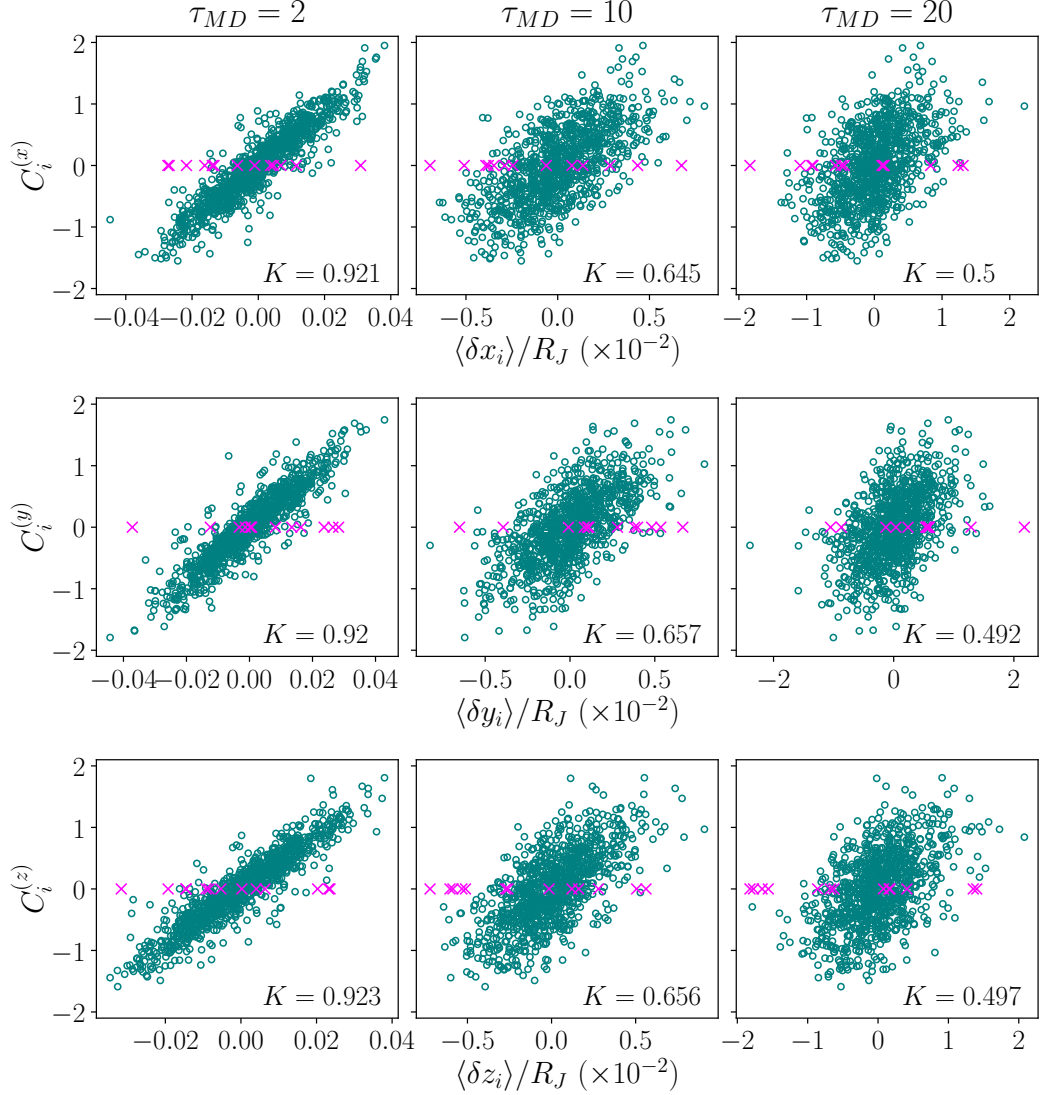


Figure 3.8. Scatter plots of the particles' mean displacement (scaled by the jamming radius R_J) and the sum of contact vectors, \mathbf{C}_i , defined in Eq. (3.1). The reported values are the ICE averages with $\varphi = 0.995\varphi_J$. Each row corresponds to one spatial dimension, while each column depicts the values of the mean displacement at different times (measured in collision events, τ_{MD}). Rattlers are identified by pink crosses and the value of the Spearman's rank correlation coefficient is indicated in the lower right part of each panel.

again, I mention that the distribution of values of the mean displacement is broad enough to conclude that many of the particles have a preferential motion direction, *i.e.* $\langle \delta \mathbf{r} \rangle \neq \mathbf{0}$. It is worth emphasizing that, since we used a large set of trajectories, this latter fact cannot be attributed to statistical noise. Another relevant feature that validates our hypothesis is that the rattlers (shown with pink crosses in Fig. 3.8 and whose pdf are drawn with dashed lines in Fig. 3.6) generically exhibit a larger value of mean displacement, which is consistent with the expectation that their tra-

jectories are less constrained because they are not part of the network of contacts; cf. Fig. 2.3. Similarly, it is also normal that initially the dynamics of each particle is highly correlated with its value of \mathbf{C} , but as more collisions occur this relation eventually gets lost.

These results confirm that there is indeed a correlation between $\langle \delta \mathbf{r} \rangle$ and the sum of unit contact vectors, which means that we can use our knowledge about $\vec{\mathbf{C}}$ to identify preferential directions in the motion of individual particles. To measure the quality of such inference we used the Spearman’s rank correlation coefficient, K , whose value is reported in the lower-right part of each panel. The advantages of using K for studying the correlation between local structure and dynamics in disordered systems are twofold: for one part, this coefficient has proven to be very sensitive to changes of the different parameters of a glass former model^[49,69]; and on the other, it naturally reflects our inferential approach based on the hypothesis that the particles ranking by a suitably chosen static variable (in this case \mathbf{C}_i) should have a direct connection with the ranking obtained via a dynamical variable (here $\langle \delta \mathbf{r}_i \rangle$).

Intuitively, the influence of a particle’s set of contact vectors in establishing its preferred direction of motion can be understood because $\vec{\mathbf{C}}$ provide a good description of where the particles’ neighbours are located and, therefore, which are the directions that will be favoured by possible collisions. Thus, a large value of, say, $C_i^{(x)}$ indicates that the particle’s nearest neighbours are arranged in such a way that their net effect is to predominantly push it in that direction. In contrast, a mostly uniform arrangement will lead to a very small value of \mathbf{C}_i and thus the particle would mainly remain in a reduced vicinity.

On the other hand, we do not expect \mathbf{C}_i to appropriately describe the fluctuations of a particle’s motion around its mean displacement mainly because the information about the directionality is lost when computing the second moment of its displacement. Fig. 3.9 corroborates this prediction by a component-wise comparison of $\text{Var}[\delta \mathbf{r}]$ with the absolute value of the sum of contact vectors. (The data were obtained with the same parameters as in Fig. 3.8, and rattlers are also identified by pink crosses). As anticipated from the results of the last part, there are no correlations between these two quantities and the values of K found in this case are notably small, reflecting the fact that contact vectors only provide information about preferential directions in the particle’s motion, but not about the their corresponding deviations.

To continue, I will study the correlation with the particles’ square displacement. $\langle |\delta \mathbf{r}_i|^2 \rangle$ is an important quantity to take into account in the high density regime because it provides an estimate of the “cage” size in which the particle is moving. As anticipated in Fig. 3.6, there is a close correlation between $\{\langle |\delta \mathbf{r}_i|^2 \rangle\}_{i=1}^N$ and \vec{S} . However, here there is no reason *a priori* why this variable should be used instead of other observables that can be obtained from the network of contacts. In particular, the role of the forces magnitudes has been so far ignored. Given that $\sum_{j \in \partial i} \mathbf{f}_{ij} = 0$, their role as analogous variables of $\vec{\mathbf{C}}$ can be immediately discarded. But this is not the case when scalar combinations are considered. Therefore, I will now show that it is more convenient to include only the contacts direction and not their magnitude in order to obtain a better predictor of the particles’ mobility. To do so, I will consider

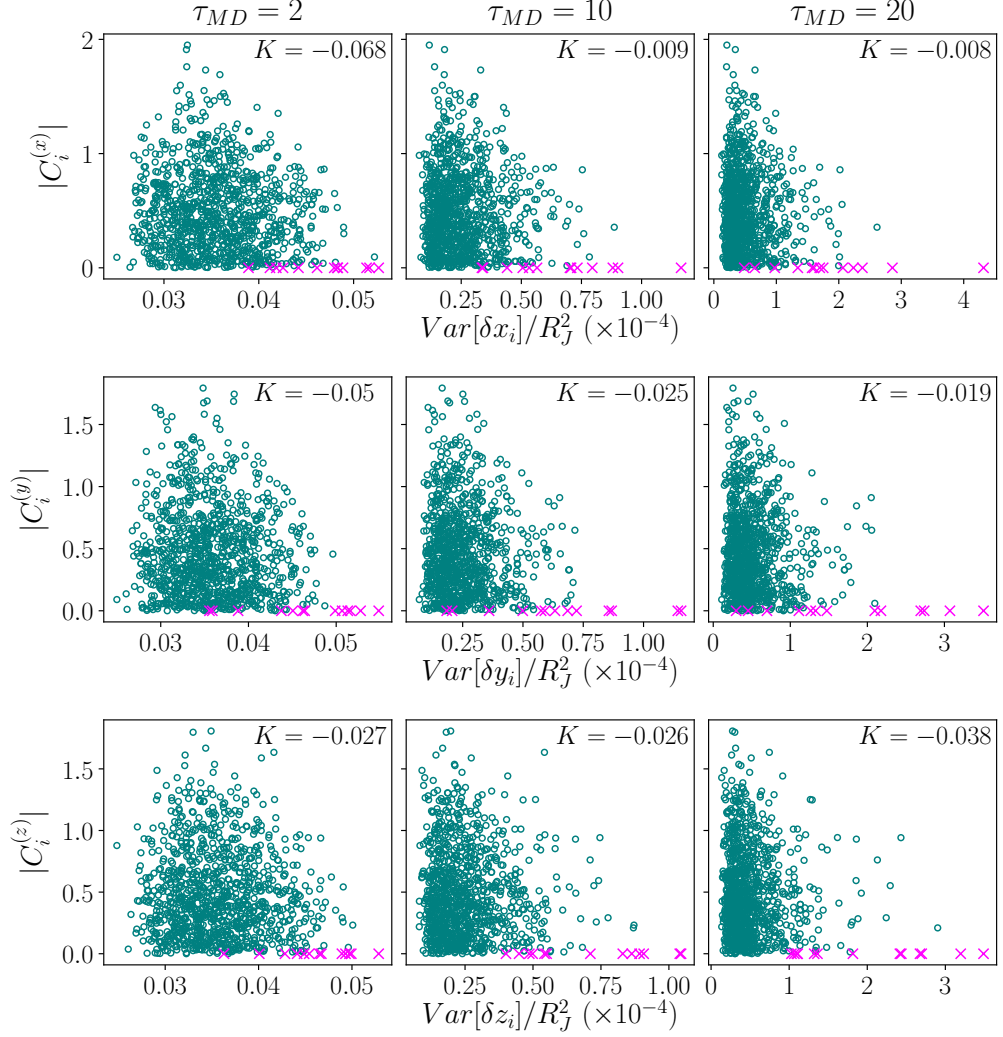


Figure 3.9. Scatter plot of the components (one in each row) of $Var[\delta \mathbf{r}_i]$ and the absolute value of the components of \mathbf{C}_i in the MD simulations at different times (different columns). Rattlers are identified by pink crosses as before.

here two other scalar quantities that are, both, physically well defined (in terms of the network of contacts at jamming), and intuitively related to the statistics of $|\delta \mathbf{r}|^2$. The first one is the sum of all the contact forces magnitudes, $\sum_{j \in \partial i} |\mathbf{f}_{ij}|$; while the

second one consists in the sum of all the pair of dot products between contact forces⁸, $\sum_{j < k \in \partial i} \mathbf{f}_{ij} \cdot \mathbf{f}_{ik}$. Fig. 3.10 compares the correlation between $\{|\delta \mathbf{r}_i|^2\}_{i=1}^N$ and these

two observables (first two columns), as well as \vec{S} (right-most column). (The values reported correspond to $\tau_{MD} = 10$.) Upper (resp. lower) panels show the results from the MD trajectories with a density of $\varphi/\varphi_J = 0.999$ (resp. $\varphi/\varphi_J = 0.99$). As before,

⁸Note that because there are many very small forces and only few large ones, despite the discussion at the end of Sec. 3.1.2, these two variables will have different distributions.

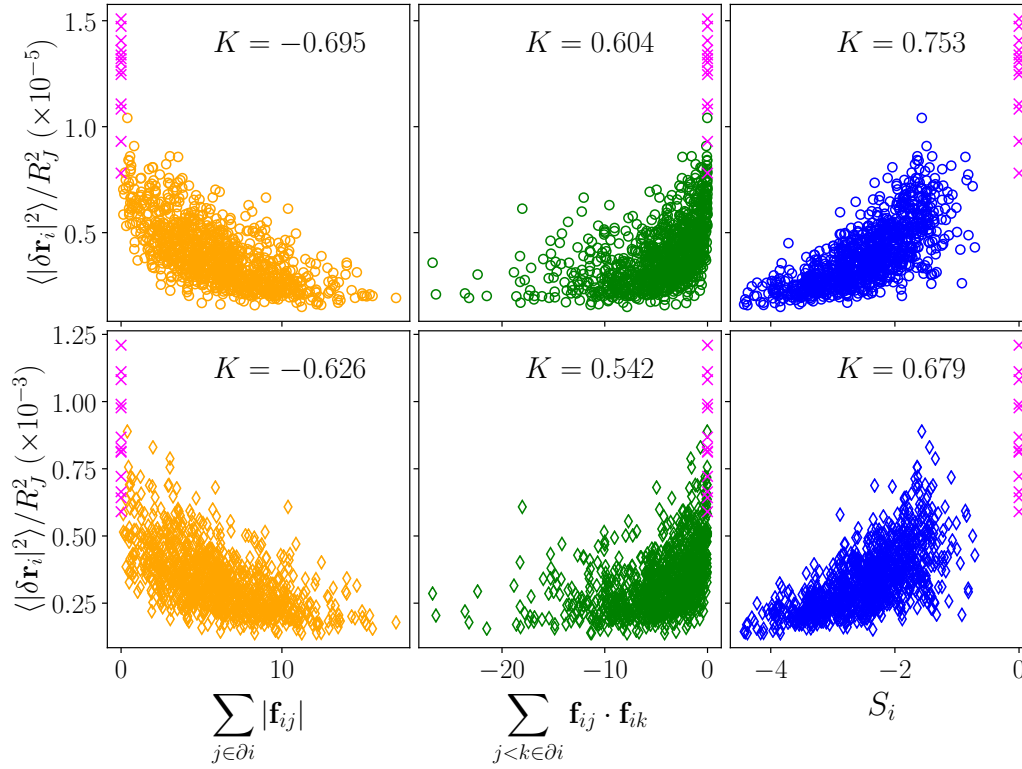


Figure 3.10. Scatter plots of the particles’ average mobility and three possible scalar quantities that can be constructed using the network of contact forces at the jamming point: the sum of forces magnitudes (first column), the sum of dot product between pairs of contact *forces* (second column), or using the contact *vectors*, S_i , (third column) as defined in Eq. (3.2). In the upper panels we report the results from simulations with density $\varphi/\varphi_J = 0.999$, while the lower ones show the analogous results using $\varphi/\varphi_J = 0.99$. As in the previous figure, rattlers are also identified with pink crosses and the Spearman’s rank correlation coefficient is also included. Note that only S_i cleanly separates this type of particles.

data associated with rattlers are indicated with pink crosses and the corresponding values of K are also included for comparison. Note that rattlers are distinguished much more clearly through S_i . As expected, in all cases the correlations attained in the system of higher density are larger because the local structure resembles more the one of the jammed state. The dynamics is therefore considerably influenced by the collisions with the closest neighbours at such state. Yet, at first sight it seems paradoxical that including more information, namely, the magnitudes of the forces, effectively reduces the predictive power. However, before explaining why this happens, I will present the analogous results from the MC simulations, where we obtained very similar findings. Hence, a common explanation applies to both types of simulations as discussed later in Sec. 3.5. To close this section, I point out that we tested this same methodology on another independent jammed configuration, as described in the Appendix 3.A, obtaining identical results; see Figs. 3.21 and 3.22.

3.4 Results with MC simulations




Figure 3.11. Average energy per particle as a function of the number of MC steps (τ_{MC}) for the three potentials used: $\alpha = 3/2$ or sub-harmonic (blue), $\alpha = 2$ or harmonic (brown), and $\alpha = 5/2$ or Hertzian (green). Different line styles correspond to different temperatures ($T_1 < T_2 < T_3$), as indicated in the right legends.

In contrast with the MD simulations, for the MC ones we fixed the spheres' radius at R_J and soften the interaction between them by introducing a contact pair potential as the one introduced in Sec. 1.4.2, but without scaling by the exponent, *i.e.*

$$U(\mathbf{r}_i, \mathbf{r}_j) = |2R_J - |\mathbf{r}_{ij}||^\alpha \Theta(2R_J - |\mathbf{r}_{ij}|); \quad (3.5)$$

where $\mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j$ and Θ is the Heaviside step function. α plays the role of a “stiffness” parameter that we varied to explore the effect of different types of interactions. Because we generated trajectories using different temperatures, introducing the proportionality constant of $1/\alpha$ simply rescales the temperature, T . Analogously to the MD simulations, we used the jammed configuration as initial condition to generate a trajectory using the Metropolis-Hastings algorithm at a fixed T . Hence, for this type of dynamics, the natural time unit is the number of MC steps performed, which I will denote as τ_{MC} to distinguish it from the time scale of the MD simulations and the physical time. We tried different interaction potentials by setting $\alpha = \{3/2, 2, 5/2\}$ (henceforth referred as sub-harmonic, harmonic, and Hertzian interaction, respectively), performing $M_{MC} = 1000$ MC simulations for each type of potential and temperature. The different values of T were selected in such a way that the samplings were done with a similar acceptance rate. This is shown in Fig. 3.11 by plotting the temporal evolution of the average energy per particle

scaled by the inverse temperature, $\beta \langle E \rangle / N$. The figure includes the results of the three interactions potentials and the different temperatures used for each one. For a fixed value of α , the energies of the different simulations are comparable between them, once the scaling with β is considered, indicating that a similar sampling was carried out. We can then conclude that the trajectories generated, using different values of T and α , belong to the same dynamical regime given that all the values of $\beta \langle E \rangle$ lie within a small range of each other.




Figure 3.12. Scatter plots of $\langle \delta \vec{r} \rangle$ and \vec{C} (each row corresponds to one direction), and for the three types of interaction considered: sub-harmonic (first column), harmonic (second), and Hertzian (third); different colours are included just for clarity sake. The temperature used to generate the MC trajectories for each potential is indicated at the top, and the values of $\langle \delta \vec{r} \rangle$ reported correspond to $\tau_{MC} = 250$ MC steps. For each value of α , the corresponding values of temperature were chosen so that the average energy of the system divided by T remained within a small range; see Fig. 3.11. As in the analogous plots of MD simulations, the quantities associated with rattlers are identified by pink crosses and the value of K is also included.

I now show that with this different dynamical protocol we obtain results in

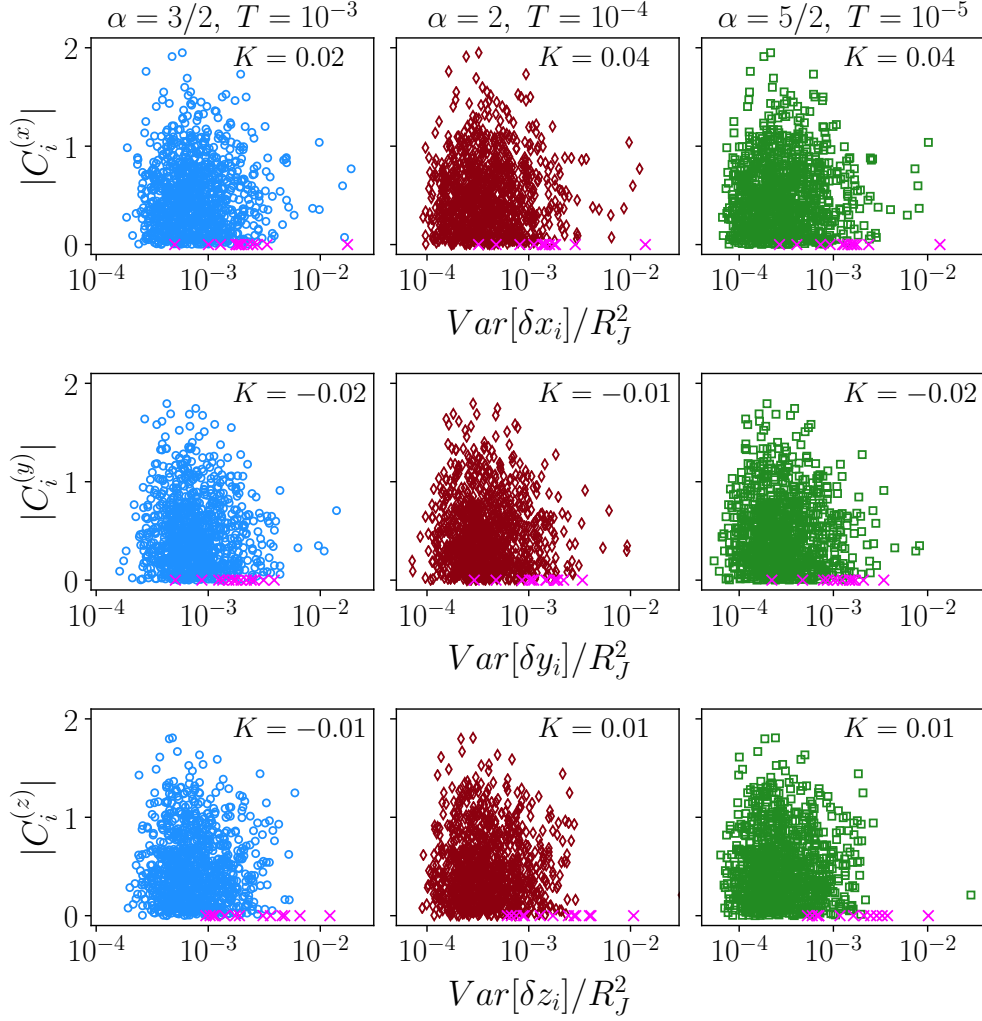


Figure 3.13. Scatter plot of the components (one in each row) of $Var[\delta \vec{r}]$ and the absolute value of the components of \vec{C} in the MC simulations for the three different potentials considered in this work (different columns). All the data reported here correspond to $\tau_{MC} = 250$ and the value of T used throughout the simulations of each potential are reported above the corresponding column. Different colours are used for the sake of clarity, except for the rattlers data, which as usual are depicted with the pink crosses. The very small values of K obtained are also included and confirm that these two variables are uncorrelated.

excellent agreement with those of MD simulations. In Fig. 3.12 I compare the correlations between $\langle \delta \vec{r} \rangle$ and \vec{C} for the three types of interactions (one in each column), and with the temperatures indicated on top. As mentioned above, the value of T was selected in such a way that a similar acceptance rate was obtained for all potentials. Each row corresponds to a component of $\langle \delta \vec{r}_i \rangle$, computed at $\tau_{MC} = 250$ in all cases, and with a different colour for each value of α . Additionally, rattlers are identified in all panels rattlers by pink crosses as in previous figures. The values of K are comparable for all the interaction potentials tested and also




Figure 3.14. Scatter plots of each particle's average mobility at $\tau_{MC} = 250$ with the same interaction parameters and temperatures of Fig. 3.12. Each row consists on the same values of $\langle |\delta \mathbf{r}_i|^2 \rangle$ as function of different scalar quantities that can be constructed using the network of contact forces at the jamming point, like the sum of forces magnitudes (first column), the sum of dot product between pairs of contact forces (second column), or using the contact *vectors* (third column) as defined in Eq. (3.2). As in the previous figures, points associated with rattlers are identified with pink crosses and the Spearman's rank correlation coefficient is also included.

with the ones obtained in the MD simulations, thus signalling that our approach is robust enough to be applied to both types of dynamics. It is also worth mentioning that there is a systematic increase of the correlation with α , *i.e.* the smallest (largest) value of K was obtained with the sub-harmonic (Hertzian) interaction, independently on the components. However, as I discuss later, this last result is related to the fact that different values of α produce different decorrelation rates, if measured by τ_{MC} . In Sec. 3.6.2 I show that if the MSD is used instead as a measure of the evolution of the configuration, the decrease of K follows a universal

behaviour. On the other hand, in Fig. 3.13 I present the analogous results of the correlation between $Var[\delta\vec{r}]$ and the absolute value of \vec{C} , using the same format and parameters than in the previous figure. In agreement with the MD results, these data demonstrates that the sum of contact vectors is a very poor indicator of the single-particle displacement's variance, even when the interaction potential has been softened. This claim is quantified by the almost vanishing values of K reported in each panel.

More importantly, when considering the mean mobility of particles and its correlation with the scalar observables constructed from the network of contacts, we also found that \vec{S} is the most informative variable in comparison with the ones that include the magnitudes of the contact forces. This finding holds independently of the type of interaction as shown in Fig. 3.14, where each column compares the correlation between $\langle|\delta\mathbf{r}_i|^2\rangle$ and the three scalar quantities considered before, while different rows (from top to bottom) correspond to the different interaction types (from $\alpha = 3/2$ to $\alpha = 5/2$). Once again, we obtained values of K comparable to the ones of the MD simulations, although this time the highest value of K was obtained with the harmonic interaction between particles. But this is due to the fact that there is a different decorrelation rate between different statistics, as I argue in Sec. 3.6.2.

To close this section I want to stress that the structural variables considered here, namely \vec{C} and \vec{S} , are computed exclusively from the properties of the configuration at jamming. Consequently their value remains fixed and is independent of the type of dynamical protocol employed. Yet, our results demonstrate that they can be utilized as proxies of the statistics of single-particle trajectories in a rather broad range of circumstances.

3.5 Contact vectors *vs* contact forces

Let me now go back to the question of why the contact *vectors* (\mathbf{n}_{ij}) are variables more informative about the dynamics near jamming than the contact *forces* (\mathbf{f}_{ij}). In other words, $\vec{C} = \{\mathbf{C}_i\}_{i=1}^N$ and $\vec{S} = \{S_i\}_{i=1}^N$, defined respectively in Eqs. (3.1) and (3.2), are predictive enough to allow us to identify a preferential direction of motion for each sphere, as well as to spot the most mobile particles. However, the value of these observables does not depend on the magnitude of the contact forces but only on their direction. And the results of Secs. 3.3 and 3.4 show that when the forces magnitude is included the quality of the prediction *worsens*. Moreover, this is the case independently of the model, *i.e.* either HS or SS, and even the interaction potential in the latter case. It is thus worthwhile to analyse more closely this counterintuitive finding.

Let me begin by considering the same problem from a purely static point of view. As derived in detail in Sec. 1.4.3, in configurations having 1SS –as we know ours do– the forces magnitude can be obtained as the unique zero mode of the *contact matrix*, \mathcal{S} defined in Eq. (1.59). Therefore, magnitudes do not contain any extra information than what is already present in the spheres position. Second, considering for the time being only the MD dynamical protocol, it is clear that as soon as the spheres are no longer touching each other, the magnitude of a contact force loses its

physical meaning, while the true force is related to the momenta exchange during the collisions. In contrast, the corresponding contact vector still contains information about which directions constrain a particle's motion the most. Similarly, in the case of SS the value of the contact force at jamming is not the relevant physical quantity either, because the real forces driving the dynamics in the MC protocol depend explicitly on both, their mutual overlap and the type of interaction potential, which we have parametrized here through α . Yet, independently of the choice of α , and hence the value of the force magnitude, the contact vectors can be used as a proxy of the force field experienced by each particle, at least for the initial part of the dynamics. So, in either case, the set of contact vectors provide an estimation of the optimal direction in which a particle should move in order to minimize the free energy of the system. Additionally, as can be noted from the rightmost equality of Eq. (3.2), \vec{S} also incorporate the influence of the number of neighbours of each particle, q_i . This feature intuitively explains that regardless of the vectorial contribution of the contacts, particles with more (less) neighbours are expected to be more (less) constrained. Besides, it also explains why a significant correlation was found when considering $\sum_{j \in \partial i} |\mathbf{f}_{ij}|$ and $\sum_{j < k \in \partial i} \mathbf{f}_{ij} \cdot \mathbf{f}_{ik}$ as a predictors, the reason being that they are essentially a weighted sum of the number of contacts. Another fact to consider is that the distribution of dot products between contact forces and the analogous one between contact vectors are significantly different, as discussed in Sec. 3.1.2; that is, Fig. 3.3 provides further evidence that these two quantities convey different information.

Even more importantly, another physical mechanism is suggested by the fractal FEL picture; see Secs. 1.3.4 and 2.3.1. The argument goes as follows: in Ref. [111] the authors measured numerically the overlap of contact *vectors* between pairs of jammed states (*i.e.* pairs of minima in the landscape), defined by a target pressure p^* , as a function of a common initial pressure of the configurations, p_0 . What they found is that as p_0 was closer to p^* , the overlap increased *gradually*. This implies that, although jammed configurations belonging to a same (fractal) meta-basin share a similar structure, the specific neighbours with which a particle will eventually be in contact, once in the jammed state, are not defined unequivocally by its local environment above jamming. Instead, they are determined progressively as such state is approached. By the same token it also indicates that as a configuration goes up in the landscape, the jammed state it departed from will still encode useful structural information, since a significant amount of the network of contacts is shared by the nearby minima. Now, note that we can think of the short time dynamics studied here as several realizations of trajectories followed by a configuration in its way to explore a small neighbourhood above its initial minimum. Under these assumptions, the original meta-basin should contain all the phase space available as this brief exploration takes place, which in turn implies that many of the interactions between nearest neighbours influencing the dynamics are going to remain roughly unchanged, and therefore we can expect major correlations with the actual network of contact vectors.

3.6 Further statistical properties

3.6.1 Predicting the mean displacement and mobility

In Sec. 1.1.1 in the first chapter it was mentioned that there have been many proposals to tackle the problem of the connection between local structure of disordered solids and their dynamics. However, one of the most pressing issues is not only to find such correlations between statical and dynamical properties, but instead using the former to *predict* the latter. To the best of our knowledge, just a handful of works have dealt in detail with this problem using well established physical variables^[41,67,218], although more recently in^[70–75, 219] researchers addressed similar questions employing machine learning methods. As I have commented before, this latter techniques have an outstanding prediction power, but they rely on an artificial representation of the particles' environment at the *mesoscopic* scale, which is an important difference with the approach we have adopted here. At any rate, as a further benchmark of our method we now address the predictability issue.

First, I should clarify that this is the inverse scenario of the one explored so far. Consider in particular Figs. 3.6 and 3.7, where we began by distinguishing the particles with some prior information about the dynamics (*e.g.* the ranking based on the first moments of $\delta\mathbf{r}_i$ and $|\delta\mathbf{r}_i|^2$) and showed that such distinction carried over to the structural features we here study (as exemplified by the clustering of data with similar colours). In contrast, I now focus on the reciprocal problem where our starting point is the static information (namely, the contact vectors) which will then be used to infer a property about the dynamics of the configuration. I will show that within our approach, prediction is indeed possible inasmuch as we can use $\tilde{\mathbf{C}}$ to reveal preferential direction in the particles' motion, while S_i can be utilized to identify the most mobile ones, *without any prior information about the dynamics*. As expected, the validity of such predictions is highest at the initial part of the trajectories and later it decays as time goes by and the system evolves.

To be more precise about how the initial correlation is lost, it is important to find a way to consistently compare the evolution of the configuration in both types of dynamical protocols and for different values of the parameters. A natural candidate is the (ICE averaged) MSD, defined in Eq. (3.4). Intuitively, it is a measure of how much the configuration has changed from its initial state, $\tilde{\mathbf{r}}^{(J)}$. Because we focused on the short-time dynamics, the values attained by the MSD during our simulations are well below the one of its characteristic long time plateau –identified with the onset of the Debye behaviour in the density of states^[51]– as I have shown in Sec. 3.3.1. Furthermore, Fig. 3.5 there shows that all our simulations were done in the same dynamical regime, where the MSD have a similar behaviour for the different packing fractions we used. Thus, I will henceforth report the evolution of the quantities we analyse as a function of $\langle\Delta\rangle$ instead of as a function of time, τ_{MD} , or τ_{MC} .

We can pose the problem of predicting the preferred directions in configuration space as investigating whether particles for which the observable \mathbf{C}_i has a large absolute value in one of its components also exhibit trajectories that move predominantly along that same component. We can answer in the affirmative this question following a straightforward method: we rank the particles according to their value




Figure 3.15. Evolution of the \hat{x} component of the mean displacement (upper panels) and the average squared displacement (lower panels) of 20% of the particles for the two type of dynamics utilized here (MD on left panels and MC on the right ones), with the parameters specified on top. The trajectories on the upper (resp. lower) panes were selected by choosing the top and bottom 10% of the particles ranked according to their value of C_i (resp. S_i); rattlers were also included in the first case. Trajectories are coloured according to the scale shown in the right. Thus, the division of colours demonstrates that we can identify both: which particles are the most mobile ones and if they have a preferred direction of motion. We used the mean squared displacement, Δ , to measure the dynamical evolution for the reasons explained in the main text.

of $C_i^{(x)}$, then select the top and bottom 10%, and plot the evolution of their trajectories. This results in the plots in the upper panels of Fig. 3.15 for the \hat{x} component of the trajectories generated via the MD simulations with $\varphi/\varphi_J = 0.99$ (left panel) or the MC algorithm with harmonic interaction and $T = 10^{-4}$ (right panel). In both cases, the lines colour indicates the value of the sum of contact vectors according to the scale on the right. These figures show that most particles which share a close value of C_i , and thus whose respective curves are shown in a similar colour, move in the same direction. Trajectories of rattlers are also included (as dashed lines) and even though we can not predict any preferred direction in their motion, we can nevertheless identify them as highly mobile particles and thus can be used as an estimate of bounds delimiting the maximum displacement of individual particles. Analogously, if we rank the spheres according to their value of S_i

we have enough information to identify the most mobile particles, as demonstrated in the lower panels of Fig. 3.15. In this case, the effect of the different protocols for the simulations is more evident for the initial part, since the MD simulations show a much narrower range of values of the squared displacement, while the MC ones yield values of $\langle |\delta \mathbf{r}_i|^2 \rangle$ that differ by more two orders of magnitude. However, the grouping of trajectories with similar colours is still preserved, confirming that S_i can be used to distinguish between mostly mobile and mostly blocked particles. Note that in this case, rattlers also provide a good estimate for an upper bound on the value of $\langle |\delta \mathbf{r}_i|^2 \rangle$ and that only for these particles S_i is identically zero, while for the others $S_i \simeq -1$ at most. As anticipated, the clear division of colours, reflecting an almost perfect correspondence between contact vectors and the statistical properties of the dynamics, is subsequently lost as the configuration moves away from its initial state, or, in other words, as Δ gets larger.

3.6.2 “Universal” decorrelation rate




Figure 3.16. Evolution of the Spearman correlation between preferential directions and total contact vector (upper panels) and between mobility and the sum of dot products of contact vectors pairs (lower panels), for several values of φ (each one in a different line colour). The right panels show that if the MSD is rescaled by a factor that measures the distance from the jamming point, simulations with different parameters follow the same behaviour. See main text for a detailed discussion.

Results so far show that significant correlations exist between the structural properties at jamming and the dynamics near the such point. Nonetheless, it is clear that as time evolves such initial correlation should eventually disappear. Therefore a crucial question is how fast and in which way this initial information is lost. To answer it, we computed the value of $K[\mathbf{C}_i, \langle \delta \mathbf{r}_i \rangle]$ and $K[S_i, \langle |\delta \mathbf{r}_i|^2 \rangle]$ at different times. (For brevity, I introduced the notation $K[X, Y]$ to denote the Spearman's rank correlation coefficient between variables X and Y .) Expectedly, different systems and different parameters yield diverse decorrelation rates, using Δ to measure the temporal evolution as above. The results obtained with the different parameters using the MD and MC simulations are reported in the left panels of Figs. 3.16 and 3.17, respectively, confirming this behaviour. However, a surprising feature is revealed when Δ is rescaled in terms of the “distance to the jamming point”: all the curves can be collapsed on top of each other for a significant fraction of the dynamics. This finding signals the existence of a common process by which the configuration decorrelates from its initial state. For instance, considering first the correlation between the total contact vector and the mean displacement in the MD simulations, it is natural to use φ to parametrize the distance from jamming. In fact, we found the appropriate rescaling factor to be $(\varphi_J - \varphi)^{1.8}$ and, as the upper right panel of Fig. 3.16 shows, when the MSD is divided by such factor the different curves fall on top of each other almost two orders of magnitude. Very similar findings hold for $K[S_i, \langle |\delta \mathbf{r}_i|^2 \rangle]$ as reported in the lower right panel of Fig. 3.16, except that in this case the scaling factor turns out to be $(\varphi_J - \varphi)^{1.65}$.

When considering the dynamics of SS, it is reassuring that we found completely analogous results using the MC dynamics. More precisely, each value of α fixes an intrinsic softness to the spheres, while the temperature determines how much they are allowed to overlap on average. Hence, once the interaction parameter α is set, the rescaling of the MSD should only depend on T . With this in mind, we can define a “thermal scale” for the MSD as $\tilde{\Delta}_\alpha(T) = c_\alpha T^{\nu_\alpha}$, where c_α and ν_α are constants that only depend on the interaction type; we report the latter in the first row of Table 3.1. By measuring Δ in terms of the corresponding thermal scale, all the different curves of $K[\mathbf{C}_i, \langle \delta \mathbf{r}_i \rangle]$ can be collapsed; see upper right panel of Fig. 3.17. For the case of the mobility, the same type of rescaling applies, albeit with different constants, *i.e.* $\tilde{\Delta}'_\alpha(T) = c'_\alpha T^{\nu'_\alpha}$; the values of these exponents are reported in the last row of Table 3.1.

	$\alpha = 3/2$	$\alpha = 2$	$\alpha = 5/2$
ν_α	0.84	0.64	0.51
ν'_α	0.45	0.40	0.35

Table 3.1. Values of the exponents used for the rescaling of the MSD with the temperature in the MC simulations, producing the curves shown in Fig. 3.17. Values in the first row correspond to the results of $K[\mathbf{C}_i, \langle \delta \mathbf{r}_i \rangle]$ while in the second row we report the associated values for $K[S_i, \langle |\delta \mathbf{r}_i|^2 \rangle]$.

To better comprehend these findings, some remarks are useful. First, it is clear that we obtained considerably high values of K for both types of structural variables




Figure 3.17. Analogous to Fig. 3.16 but for MC dynamics, modelling SS with different contact potentials (different colours) and at different temperatures (different markers shape). Note that once again, if the MSD is rescaled by a function of the distance to the jamming point (in this case determined by T), the different curves can be collapsed into a single one (right panels).

and dynamical protocols. Additionally, the fact that during the intervals considered here the value of K decreases by roughly an order of magnitude, while the (rescaled) MSD increases by at least two decades, reveals that the decorrelations are actually rather slow. It is also worth emphasizing that, in all cases, the rescaling factors essentially measure how far away from the jamming point the dynamics takes place: either by reducing its packing fraction (in the HS scenario) or by providing some thermal energy (in the case of SS with a contact potential). Unfortunately, no theory exists for predicting the form of these scaling variables, and here we had to determine them by inspection. Nevertheless, if we resort to the fractal FEL picture, the curves collapse suggests that, independently of the type of dynamical protocol or even the interaction potential, the same mechanism drives the loss of correlation as the system explores its meta-basin. It is tempting to rationalize this feature in terms of the jamming universality (see Sec. 1.4.1), but I consider that a careful, dedicated study, with a more systematic characterization of the dynamics and other physical variables, is needed for such purpose. Finally, we can understand that the correlations with $\tilde{\mathbf{C}}$ and $\tilde{\mathbf{S}}$ yield different scalings exponents because these two

observables provides information about different features of the dynamics, namely, the directions in which the particles' displacements are facilitated and how much they are expected to move, respectively. I finish this section mentioning that we also verified that the same results hold for the other configuration we tested, as shown in Fig. 3.23 in the Appendix 3.A.

3.7 Comparison with normal modes and other works

In the Sec. 1.1.1 it was mentioned that the dynamical properties of systems near jamming have been investigated predominantly in experiments, and for the scope of this chapter it is relevant to consider Ref. [211]. In this study, the authors found a clear connection between the dynamics of their samples, which consisted in a bidisperse mixture of photoelastic soft disks, and the network of contacts formed among them. They studied both the under- and over-compressed parts of the dynamics, although an important difference with our work is that in their case the dynamics was driven by an external vibrating apparatus rather than by thermal fluctuations as we considered here. At any rate, our results should be considered as complementary to each other because while we focused on the statistics of the trajectories for a given configuration, their results are based on several realizations on the network of contacts, *i.e.* over several jammed states. The fact that during our simulations we found displacements of the same magnitude as the ones reported experimentally is reassuring.

On the other hand, this reference is also important because their experiments managed to probe several of the dynamical features analysed theoretically by Ikeda *et al.* in Ref. [51]. In this latter work, the authors identify a “critical regime” characterized by strong anharmonicities that make a vibrational description unavailable. As I will argue below, our simulations fall precisely within this same critical region. To provide the necessary context, I briefly summarize the results obtained by Ikeda *et al.* that are most important for this thesis. Their main finding is that the changes in the behaviour of the MSD can be associated with corresponding changes in the DOS, $D(\omega)$, where this latter quantity is calculated via the Fourier transform of the velocity autocorrelation function of the particles, $d(t) = \frac{1}{3N'T} \sum_{i=1}^{N'} \langle \mathbf{v}_i(t) \mathbf{v}_i(0) \rangle$ (as above, N' is the number of non-rattlers). For instance, they showed that the time at which the MSD deviates from the ballistic regime can be associated with a frequency for which the rapid decay of the DOS at large ω occurs. Similarly, for the time at which the MSD reaches its plateau value at high densities, there is a corresponding frequency which indicates the onset of the Debye behaviour for small frequencies, $D(\omega) \sim \omega^{d-1}$. And although this identification is blurred out as the systems move away from their jamming point, say, by raising the temperature or changing the packing fraction, the authors were able to define a region where the harmonic description is always valid. Indeed, they showed that there is a temperature, $T^* \sim |\varphi - \varphi_J|^\alpha$, such that whenever $T < T^*$, then the collective properties of the dynamics are well described by the vibrational modes. In turn, if for a given φ the temperature of the system is large in comparison with T^* , then the dynamics is dominated by anharmonic effects and no vibrational description is available and the system is in the so called critical regime.

Importantly, the instances considered in our work lie precisely in this latter regime as I will now show. To begin with, recall that the MC simulations were performed exactly at $\varphi = \varphi_J$, meaning that for any finite temperature, the dynamics are not be purely vibrational. In fact, Ikeda *et al.* state that in this scenario, the jamming DOS cannot be used to infer properties of the vibrational dynamics for any $T > 0$. Additionally, the temperature range that we considered for the harmonic interaction –the only case analysed numerically by the authors– is several orders of magnitude above T^* . So even if the harmonic description was appropriate at the packing fraction we utilized, all its effects in our simulations would have been washed out due to thermal noise.

The case of the MD simulations using HS is more subtle, but in Ref. [211] it was pointed out as well that the experimental configurations analysed in [207, 218, 220], which consisted of hard brass disks, also belong to the anharmonic regime. Clearly, these are empirical realizations that closely resemble the infinitely rigid spheres of our simulations, and suggest that our findings with $\varphi < \varphi_J$ should lay beyond an interpretation in terms of vibrational modes. Nevertheless, we deem indispensable to stringently test that normal modes are unable to described the single-particle dynamics in the regime considered here, so I will present some results obtained following such method. I should emphasize that our procedure cannot mimic exactly the one of Ref. [51] because, as mentioned above, the authors used the velocity autocorrelation function in order to compute $D(\omega)$. However, note that $d(t)$ is intrinsically a dynamical property, while our approach relies on exclusively using knowledge about the structure of the configuration. This limitation is, fortunately, only apparent because using the formalism developed in Refs. [150, 167–169] and reproduced in detail in Sec. 1.4.3, we can compute the Hessian (H) at the jamming point and then obtain the *exact* DOS by diagonalizing it. I think it is worth remarking that, formally, this is the procedure that should be followed to compute $D(\omega)$, although under certain conditions it has been shown^[52] to coincide with the DOS obtained from the Fourier transform of $d(t)$, which is the technique most commonly utilized.

In the case of the harmonic potential, the Hessian is given by $H = \mathcal{S}^T \mathcal{S}$, where \mathcal{S} is the $N_c \times dN'$ contact matrix mentioned above an whose entries are spelled out in Eq. (1.59). As a side note, I bring the attention to the fact that the vibrational formulation is unable to take the rattlers into account, since they are omitted from \mathcal{S} by construction; in contrast, our method incorporates them readily. Now, when dealing with a different type of interaction, universality ensures that a jammed state obtained with a given potential is also a valid jammed configuration for any other potential (see Secs. 1.4.2, 1.4.3 and Refs. [104, 150]). In the case of HS, the entries of the contact matrix should be rescaled by the corresponding force magnitude (see Eq. (1.64)), whence a similar expression for H is recovered, as stated in Eq. (1.65). Explicitly, its entries are

$$H_{i\alpha}^{j\alpha'} = \delta_{ij} \sum_{k \in \partial i} f_{ki}^2 n_{ki}^\alpha n_{ki}^{\alpha'} - \delta([ij]) f_{ij}^2 n_{ij}^\alpha n_{ij}^{\alpha'} \quad (3.6)$$

where $\delta([ij])$ is a function that is one only if particles i and j are in contact and zero otherwise.

To obtain the normal modes of the system, we need to obtain the eigenvalues of the Hessian, which in turn determine the normal frequencies of the system and, thus, the DOS: if $\{\lambda_n\}_{n=1}^{3N'}$ are the eigenvalues of H , then the associated frequency is $\omega_n = \sqrt{\lambda_n}$. The $3N'$ eigenvectors⁹, $\{\vec{v}_n\}_{n=1}^{3N'}$, correspond precisely to the vibrational modes, each of which describe a unique oscillatory mode of the configuration around the jamming energy minimum. For further use, I introduce here the Inverse Participation Ratio (IPR), which can be calculated in terms of the eigenvectors as^[58],

$$Y(\omega) = \frac{\sum_{i=1}^{N'} |\mathbf{v}_i(\omega)|^4}{\left(\sum_{i=1}^{N'} |\mathbf{v}_i(\omega)|^2\right)^2} \quad (3.7)$$

where $\mathbf{v}_i(\omega)$ is a $d = 3$ dimensional vector obtained from the eigenvector \vec{v} with corresponding frequency ω by taking the components associated with particle i . Intuitively, $Y(\omega)$ quantifies the localization of the mode with frequency ω . That is, for very small values ($Y \sim 1/N'$) of the IPR, the associated mode is essentially extended, meaning that many particles are collectively displaced, all of them by roughly the same amount. In contrast, for large values ($Y \sim 1$) few particles participate in the displacement excitation, so the mode is localized.

Now, if the particles were *only vibrating* around an energy minimum their displacements would be easily captured by the normal modes formalism. Moreover, if this was the case, we could expect the lowest energy modes to carry most of the weight of the configuration's displacement. To put to test this assumption, we measured the projections of the configurational vector of mean displacement along the different modes, *i.e.* $\langle \widehat{\delta \vec{r}} \rangle \cdot \vec{v}_n$, where only non-rattlers components are included and $\langle \widehat{\delta \vec{r}} \rangle$ indicates that the average displacement vector has been normalized. Such normalization is performed at each time, so the different values of the projections remain comparable to each other and a fair comparison is possible as time evolves. The resulting values from the MD simulations with $\varphi/\varphi_J = 0.995$ are reported in Fig. 3.18, where the Hessian was obtained with both the harmonic (upper panel) and hard sphere (lower panel) interaction potentials, while different colours correspond to different times at which $\langle \widehat{\delta \vec{r}} \rangle$ was calculated. The scatter of the points hints that, for either type of potential, the collective particles' displacement is not concentrated in few modes, specially during the initial times. For longer times, more weight is allocated in the first modes, but still there are many other low frequency modes with negligible weight. We consider that such lack of criterion to select the most relevant modes exhibits one important drawback of the normal mode description at the single-particle level, and, unfortunately, it carries over to the IPR as I discuss next.

Considering the IPR is important because, independently of their frequency, we expect that not all the modes are equally important, due to the fact that the trajectories we studied here consist in displacements of all the particles of the system. Hence, a more reasonable property to study is whether the extended modes possess a higher weight than the localized ones. Were this to be the case, then choosing

⁹I will also use the notation $\vec{\bullet}$ to denote the $3N'$ dimensional vectors of the configuration space, in order to highlight the fact that any displacement $\delta \vec{r}$ can be expressed as a linear combination of the eigenvector. Of course, such decomposition assumes that *rattlers are excluded* from the components of $\delta \vec{r}$.




Figure 3.18. Value of the projection of $\widehat{\langle \delta \vec{r} \rangle}(t)$ along each of the normal modes, $\{\vec{v}_q\}_{q=1}^{N'}$, obtained from the Hessian using the harmonic (upper panel) or hard-sphere interaction potential (lower panel). The different times at which the dot products were computed are identified by different colours and markers, according to the legend. Values of $\widehat{\langle \delta \vec{r} \rangle}$ were obtained from the simulations with $\varphi/\varphi_J = 0.995$. It is only for later times that more weight is assigned to few of the low frequency modes, but there is no way of knowing, *a priori*, how to select them.

the modes with the lowest value of $Y(\omega)$ could be taken as a criterion for selecting the vibrational modes that influence the dynamics the most. Yet, as Fig. 3.19 demonstrates, even using the IPR we are unable to identify the correct normal modes to account for the configuration's displacement statistics. That is, even when the IPR is very small, *i.e.* $Y(\omega) \in (2 \times 10^{-3}, 10^{-2})$, the corresponding values of the projections lie within a range of 2-3 orders of magnitude. Such a broad distribution exemplifies that the collective dynamics cannot be described as if it consisted of few extended excitations. Note that, on the one hand, these results are independent of the type of potential, while on the other, allowing the dynamics to evolve for longer times only affects the localized modes, the vast majority of which loses most of its weight. A more quantitative comparison is at hand by considering the three most extended modes (the ones lying on the vertical line $1/N'$) corresponding to the zero modes of H , which in turn are associated with the translational invariance

of our system induced by its periodic boundary conditions. These modes determine directions along which the system could be collectively displaced without any energy cost, but since we know that during the dynamics the systems is *not* being uniformly translated, we can use their weight as a lower threshold to distinguish relevant modes for the dynamics. What we observe from Fig. 3.19 is that there are several modes (also considerably extended) that have an even smaller weight, while we still lack a guideline to find the ones more aligned with $\widehat{\langle \delta \mathbf{r} \rangle}$. Furthermore, note that the sign of the projection is particularly important if we really want to be able to identify preferential directions. But still, the vibrational scheme fails to provide this information because it is assumed that the motion mainly consists in oscillations around a given minimum, so the direction of motion becomes irrelevant. Nonetheless, our results of the previous sections prove that this is not the case for the dynamical regime we are studying. For completeness, I mention that we verified that our results are unmodified if use $K[\widehat{\langle \delta \mathbf{r} \rangle}, \vec{\mathbf{v}}_n]$ instead of the projection for measuring the correlation between the configurational displacement and the normal modes.

As I argued before, the lack of connection between the mean configurational displacement and the distinct vibrational modes can be explained by the fact that these modes are not suitable to examine the trajectories generated with our simulations, given that they belong to the critical region where the harmonic approximation is not valid. And even though the results presented so far support this hypothesis, we have not yet verified that the dynamics is mainly driven by anharmonic effects. To do so, we computed the value of $\widehat{\langle \delta \mathbf{r} \rangle} \cdot H \cdot \widehat{\langle \delta \mathbf{r} \rangle}$, which should provide an estimate of the “curvature” induced by the dynamics at the FEL’s minimum. The idea is that if the harmonic description was valid, this curvature should be small given that particles move predominantly along the lower-energy (*i.e.* flatter) directions. Conversely, a large value signals that the forces driving the dynamics are not captured by the linear assumption implicit in the harmonic approximation. To test these hypotheses, in the upper (resp. lower) panel of Fig. 3.20 we plot the cdf of the frequencies for the harmonic (resp. hard-sphere) potential and we indicate the values of $\widehat{\langle \delta \mathbf{r} \rangle} \cdot H \cdot \widehat{\langle \delta \mathbf{r} \rangle}$ with vertical dashed lines, using different colours to distinguish different times. Recall we have normalised the average displacement at each time, so the intersection of each vertical line with the curve of $\text{cdf}(\omega)$ determines the fraction of frequencies smaller than the corresponding value of the curvature. Clearly, during the initial times the value of the curvature is quite large, *i.e.* greater than 40 – 60% of the frequencies, specially considering the hard-sphere potential. But also for longer times this fraction remains around 20% (*i.e.* an extensive number of modes), corroborating that the preferential directions of motion, inferred via the vibrational modes of low energy, do not correspond to the real trajectories followed by the particles. This finding establishes that our systems were in the anharmonic region, according to the classification of Ref. [51].

3.8 Conclusions

I presented a simple and robust method to connect the structural information present at the jamming point of a configuration of N spherical particles and the




Figure 3.19. Same projections as in Fig. 3.18 but plotted as a function the IPR defined in Eq. (3.7), using the harmonic and hard-sphere potentials, panels (a) and (b) respectively. The black vertical dashed line at $1/N'$ is the lowest value that can be attained by $Y(\omega)$ and correspond to maximally extended modes, while the dotted line at 1 identifies the most localized ones. The few points with $Y(\omega) \approx 1/N'$ are associated with the $d = 3$ zero modes caused by the translational invariance of our systems and their respective weight can be used as a lower threshold for determining when a mode can be considered relevant (see main text).

dynamics that happens close to such point. I first analysed the statistics of the spheres' displacement after producing several trajectories originating from the positions at the jammed state. The resulting distributions show, on the one hand, the existence of preferential directions of motion, and on the other, the presence of heterogeneity in the particles' mobility, reflected by the fact that some spheres remain mostly fixed, while others are more mobile by several orders of magnitude, see Fig. 3.6. I also argued that these features are captured by their respective first moments, *i.e.* $\langle \delta \mathbf{r}_i \rangle$ and $\langle |\delta \mathbf{r}_i|^2 \rangle$ for $i = 1, \dots, N$, and thus we used these statistics to succinctly represent the full distributions.

The main result is that we found a straightforward and significant connection, at the single particle level, between these statistics and the network of contact vectors that is formed at the jammed state. In particular, we found that the sum of contact vectors acting on a given particle, \mathbf{C}_i (defined in Eq. (3.1)), is a good predictor of




Figure 3.20. Comparison of the “curvature”, measured by $\widehat{\langle \delta \vec{r} \rangle} \cdot H \cdot \widehat{\langle \delta \vec{r} \rangle}$, with the spectrum of H computed using the harmonic (top) and hard-sphere (bottom) interaction. Values of this curvature at different times are indicated by the vertical dashed lines, and their intersection with $\text{cdf}(\omega)$ (black curve) equals the fraction of frequencies smaller than ω . The fact that the fraction of frequencies smaller than a given value of curvature is always extensive, even for long times, reveals that the dynamics probed during our simulations are not captured by the normal modes approach.

its preferred displacement direction, if any; while the sum of dot products between all pairs of contact vectors, S_i (defined in Eq. (3.2)), is highly correlated with the particle’s mobility. Our approach was tested using two different types of dynamical protocols, namely Molecular Dynamics and Monte Carlo simulations, and verified that it remains valid in both cases and for the different parameters we studied; cf. Figs. 3.8-3.15. The statistical correlation, measured by the Spearman coefficient, K , decays rather slowly as the dynamics evolves, thus showing that the jammed state used as the initial condition of the trajectories has a persistent influence on the particles displacements. More precisely, using the MSD as a measure of how far the configuration is from its initial state, we found that while K decreases its value by an order of magnitude or less, the MSD has increased by at least two decades. Because we are not probing the ballistic regime, the change in the MSD is only due to interaction between particles, and thus implies that \vec{C} and \vec{S} are relevant after a

small, but definitely measurable, rearrangement of the whole configuration has taken place. Furthermore, in Figs. 3.16 and 3.17 I showed that when Δ is rescaled by a factor that depends on how far from the jamming point the dynamics occurs, the curves of the loss of correlation follow a single master curve. This suggests that for a given dynamical protocol the decorrelation rate is, surprisingly, independent of the distinct parameters used in our simulations. We therefore conclude that the network of contact vectors formed at a jamming point contains significant information to infer the short-time features of the particles' motion near such point.

Importantly, our results show that in the dynamical regime we studied, and for the timescales of our simulations, the particles' motion does not consist of simple vibrations around a reference state, which in our case is naturally identified with the initial jammed configuration. This novel finding is supported by the fact that the vibrational modes fail to capture the real mean displacement of the configuration, as argued in the previous section. The negligible correlations obtained in this case are reported in Figs. 3.18 and 3.19. I should emphasize that the scope of our work is restricted to finding a connection between dynamical and structural (and hence, static) variables, therefore we only used data about the contact forces as input for our statistical inference. These data can be used to compute the Hessian at the jamming point, whence the normal modes can be easily obtained. This shows that, utilising the same structural information as input, our approach yields a better description of the statistics of the single-particle dynamics than the normal modes approach. Moreover, the level of statistical correlations we obtained is comparable to the ones derived from more sophisticated, state of the art methods.

Additionally, I want to mention that I am optimistic that our method could be applied, given that it is rather simple and straightforward, to other type of systems having jamming as a critical point (see Sec. 1.4). In other words, it is likely that the structural variables we identified here, namely \vec{C} and \vec{S} , also convey information about the particles trajectories when the dynamics is driven by other physical variables, *e.g.*, the application of a load or a stress. A related feature is that because jamming criticality is noteworthy robust to changes in dimensionality and interaction types, our approach presumably has the same range of validity. Besides, recall that the jamming universality class is composed not only by amorphous solids and glass formers but also by several constraint satisfaction problems and learning algorithms. Thus, once the analogous structural variables are identified in those cases, our method could also be used to extract new insights about the *algorithmic* dynamics of that sort of problems.

Now, some remarks about the regime of validity of our results are also in order. For this purpose, I will once again follow the picture of the fractal FEL, in particular the results of Ref. [111] and Sec. 2.3.1. There it was shown that a jammed state is realized gradually as the configuration goes down towards a minimum and the contact forces are concomitantly determined. Therefore, similar jammed states exhibit similar network of contacts and, more importantly, different realizations of jammed states, originated from a single initial configuration, would have a significant fraction of common contacts. Of course, the closer this initial configuration is to an energy minimum the higher the fraction of common contacts in the final jammed packings. Now, assuming that the short time dynamics essentially explores a small vicinity around the reference jammed state, these findings set the Gardner phase as

a rough limit beyond which we do not expect our method to remain valid. That is as long as the dynamics departs from a state (in parameter space) within this marginal phase, it is likely that we find non-negligible correlations between our structural variables and the single-particle displacement statistics. To see why this is the case, let us imagine that we use a jammed configuration as the initial condition of a given dynamical protocol. After a short time we stop the evolution, then bring the configuration to a new jammed state and, finally, we restart the dynamics using this second state as initial condition. Since the two jammed states are similar, we expect that the second trajectory to be correlated, at least to some extent, with the structure of the first jammed state. Yet, note that for this scenario to be true it is crucial that the dynamics occurs inside the Gardner phase because only in this way we can guarantee that the structural variables from the original configuration are going to be similar to the ones of the nearby sub-basins explored during the dynamics. Nevertheless, this argument does not provide a criterion for determining the rate and sharpness of the loss of correlation as the stable glass phase is approached. This and other open questions are described in more detail next.

3.8.1 Future work

I finish by pointing out several routes of improvement that can be considered. The first one is to construct other variables that incorporate information of the “second nearest neighbours”. In other words, as it is now, our method only considers the role of a particle’s first neighbours, but it is expected that a more complete description can be attained when a larger vicinity is considered. An independent possibility to explore is that, with enough data, we could aim to infer the full distribution of the particles displacements and not only their moments as we did here. This would provide a complete description of the trajectories followed by amorphous solids near their jamming point. Clearly, both approaches can be combined, where the degree of statistical agreement of the full distribution expectedly increases as more complicated models are taken into account, which in turn would include observables constructed from an even larger set of nearby particles and more parameters. This opens the way to the application of Bayesian model selection, which has proven to be very fruitful in a surprisingly broad range of topics.

On the other hand, as I stressed through this chapter, we have used a particle-wise approach in which the most relevant structural features are captured by the position of a particle’s nearest neighbours. But one could also think on implementing a more “coarse-grained” description in which the dynamics of a group of particles (say their mean mobility) is characterized by a structural quantity of the corresponding region (for instance the average value of S_i in the cluster). This clusterised approach would be similar to the more common methods that explore the influence of the local environment on the dynamics of glassy system. Hence we expect that the more we adopt a cluster description, the longer our inferences should remain valid. In this way, one could extend the duration of high correlations by including the information of more neighbours and analysing their statistical properties.

The possibilities I just mentioned are concerned with developing techniques to

improve the quality and duration of the statistical inference, but our results also hint at some properties that deserve to be explored, from a physical point of view, with greater detail. First of all, the spatial correlations of \vec{C} and \vec{S} . That is, based on the fact that particles are dynamically correlated at the large scale, while we have shown that dynamics and local structure are connected, we can anticipate that \vec{C} and \vec{S} (or other similar structural variables) should also exhibit signatures of correlations on longer length scales. Unfortunately, due to high fluctuations in these structural variables, a very large number of independent jammed configurations should be used in order to extract a meaningful signal and thus identify the presence of truly extensive structural correlations. Nonetheless, we deem that putting this hypothesis to test is particularly relevant since it would provide stronger evidence of the influence of the local structure on the dynamics.

Additionally, a promising path for extension, also related to long range correlations, is how the presence of system-wide order or regularity would influence our results. Candidates for such studies could be, for instance, either low density systems –such as tunnelled crystals^[189]– or, on the contrary, high density ones –formed by slightly polydisperse crystals^[162,163]. Besides the usual marginal stability of common jammed packings, which leads to structural correlations lengths comparable to the system size (see Chp. 4), these systems are characterised by periodic ordering of their particles. Therefore, two different types of structural correlations coexist across the whole configurations. It is clear that our method can be directly applied to such systems because the network of contacts is equally well defined in all of them, but the extent to which their inherent order increases the relevance of the second nearest neighbours in the dynamics, and thus the level of correlation our approach would yield, is something that needs to be explored.

Another independent feature to investigate is how the dynamics of the configuration in its original basin affects the loss of correlation in time. In other words, it would be interesting to study if there is a connection between this decorrelation and the distance from the original minimum to nearest one at a given state of the configuration. This would provide evidence that the decorrelation is strongly influenced by the overlap between the many jammed states in a basin of the FEL and, moreover, it could suggest how to use the structure of these other free energy minima to enhance the quality and the duration of the statistical inferences we presented here or other similar ones. Additionally, if a stronger link between such statistical correlations and the overlap of jammed states within a meta-basin is found, it could be used to set a more accurate limit for the range of applicability of the sort of inference tools we developed here.

Appendix

3.A Reduced analysis on an independent configuration

To verify that our methodology is system independent, we performed the same analysis on another configuration selected at random from the set of configurations analysed in Sec. 3.1.2. It also consists of $N = 1024$, with 1.4% of rattlers, and packing fraction $\varphi_J = 0.6350$. We used the same MD algorithm to explore its

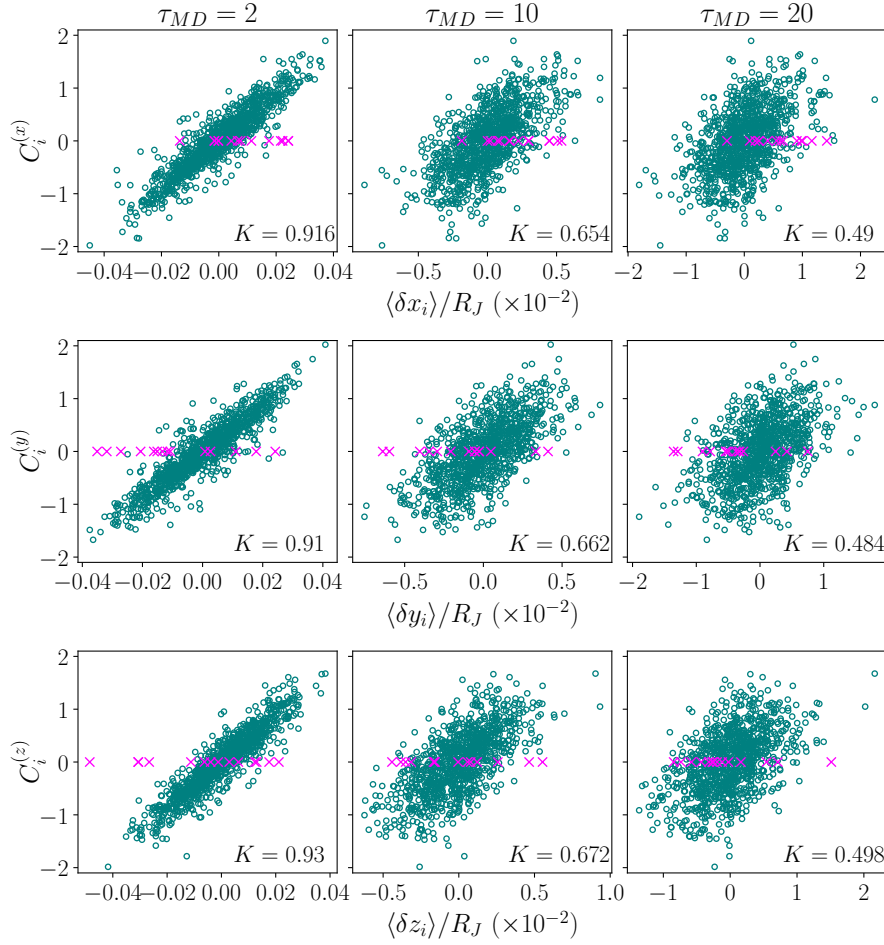


Figure 3.21. Correlations of mean displacements and \mathbf{C}_i for a second independent configuration. Note the close resemblance to the results reported in Fig. 3.8. The packing fraction used is $\varphi/\varphi_J = 0.995$.

dynamics, although in a reduced set of packing fractions, obtaining very similar results and values of correlations with the two structural variables of interest. For instance, the scatter plots of Fig. 3.21 show the correlations between $\langle \delta \mathbf{r} \rangle$ and $\vec{\mathbf{C}}$, analogous to Fig. 3.8. Likewise, we tested the three scalar variables considered in Figs. 3.10 and 3.14 and once again found that the value of \vec{S} is a better predictor of the particles mobility. We thus confirm that, compared to contact forces, contact vectors are more useful to infer dynamical statistical properties.

Finally, we also verified that the decorrelation from the initial configuration follows the same universal behaviour as the ones reported in Fig. 3.16. To do so, we computed the values of $K[\langle \delta \mathbf{r}_i \rangle, \mathbf{C}_i]$ and $K[\langle |\delta \mathbf{r}_i|^2 \rangle, S_i]$ as they evolved in time, and compared them with the ones of the original configuration at the same packing fraction. The comparison is shown in Fig. 3.23. The fact that the two curves resemble each other very closely shows that our approach is robust enough to be applied to any jammed configuration of spherical particles.

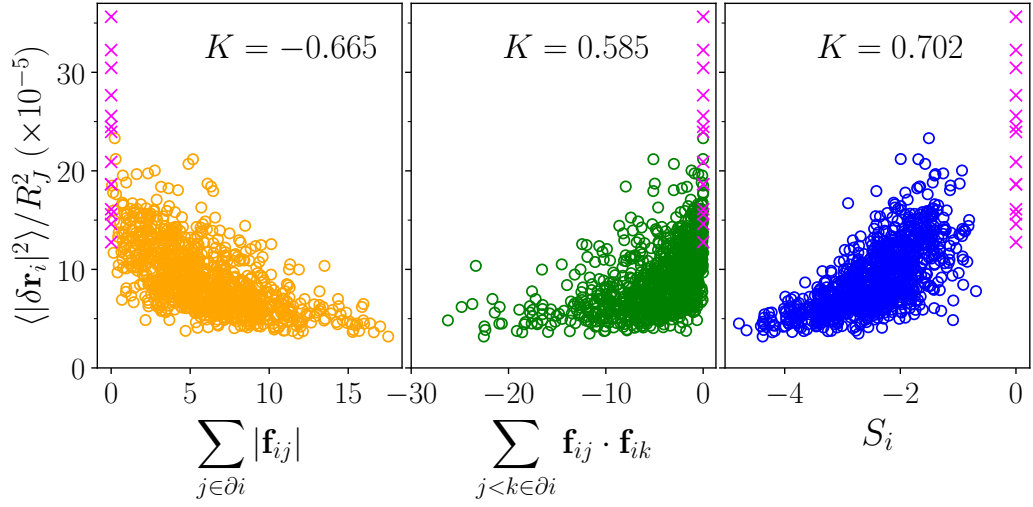


Figure 3.22. Correlations of particles' mobility and three different structural variables obtained from the network of contact forces. Note that these results reproduce the ones reported in Fig. 3.10. The packing fraction used is $\varphi/\varphi_J = 0.995$.




Figure 3.23. Comparison of the temporal evolution of the statistical correlations between structure and dynamics in two independent configurations.

Chapter 4

Finite size scaling of microscopic structural variables

4.1 Introduction

This chapter contains a detail study of the finite size effects in the critical distributions of contact forces and interparticle gaps; see Secs. 1.4.4-1.4.5 and Eqs. (1.67)-(1.68) in particular. For convenience, I reproduce here such equations:

$$g(h) \sim h^{-\gamma}, \quad \text{with } \gamma = 0.41269 \dots ; \quad (1.67)$$

$$p_\ell(f_\ell) \sim f_\ell^{\theta_\ell}, \quad \text{with } \theta_\ell \simeq 0.17, \quad (1.68a)$$

$$p_e(f_e) \sim f_e^{\theta_e}, \quad \text{with } \theta_e = 0.42311 \dots ; \quad (1.68b)$$

Many of the results of this chapter have already been presented in Ref. [2], in collaboration with several colleagues to whom I am deeply grateful. There, we studied said distributions in four models: monodisperse spheres, the Mari–Kurchan (MK) model (see Secs. 1.4.1 and 2.2) in $3d$, polydisperse disks, and minimally polydisperse FCC crystals. However, here I will only present results from the first two, because those are the ones to which I contributed more actively. Nevertheless, in the case of the monodisperse spheres, I will also make use of the numerical data of our colleagues to exemplify the very nice match obtained when the jamming point is approached from either side of the transition using hard spheres (HS) or soft spheres (SS); see Figs. 4.3 and 4.6 below. Actually, this last point is but a reason for the need of further testing the non-trivial power law distributions of forces and gaps, despite what has been said in Sec. 1.4 about the universality and robust criticality of the jamming transition. To explain why, I will briefly come back to some of the points already discussed in that section to give some perspective of how they affect the topics of this chapter. Specifically, a precise determination of the exponents of the distribution of forces, $\{\theta_e, \theta_\ell\}$, and gaps, γ .

Recall that in Sec. 1.4.2 I mentioned that several thermodynamic variables in configurations of spheres exhibit scalings in the vicinity of the jamming point but that, importantly, these relations differ if jamming is approached from the *over-compressed* (OC) phase –see Eqs. (1.51)– or from the *under-compressed* (UC) one¹

¹Throughout this chapter, I will only use physics-based terminology. But it is important to keep

–Eqs. (1.52). The most salient example is the (reduced) pressure that, assuming an harmonic interaction in the OC regime, scales as $p \sim |\varphi - \varphi_J|^{\pm 1}$, as $\varphi \rightarrow \varphi_J^{\pm}$, where once again φ_J is the system’s density at jamming. Such power laws behaviours of thermodynamic variables have been amply validated both by numerical results^[51,136,137] and by detailed scaling analyses^[154–156]. Additionally, we now know that these features are essentially unmodified by changes in dimension, polydispersity, etc. And even though they are strictly valid only in the $T \rightarrow 0$ limit, the role of temperature has been partially incorporated in some studies^[51,152]. This behaviour is not exclusive of thermodynamic variables, but is also observed in quantities such as the coordination number, z , for which numerical simulations show that its mean value, $\bar{z} = \frac{1}{N} \sum_{i=1}^N z_i$, exhibits a discontinuity exactly as $\phi \rightarrow \phi_J^-$, and then grows as $\bar{z}(\phi) - \bar{z}(\phi_J) \sim (\phi - \phi_J)^{1/2}$ for $\phi > \phi_J$ ^[136,153].

Hence, we should not immediately take for granted that the exponents of the critical distributions have the same value if jamming is approached from the high or the low density phase. In fact, to the best of my knowledge, the only work analysing this question in detail is Ref. [153] and they obtained *different* exponents of the forces distributions in the OC regime ($\theta_{SS} \approx 0.42$) and in the UC one ($\theta_{HS} \approx 0.28$). The situation was further complicated because shortly afterwards, the authors of Ref. [169] estimated $\theta_{SS} \approx 0.17$ after improving on the statistics (although they also reported $\theta'_{HS} \approx 0.23$ in $3d$ monodisperse packings). Besides, some dependence on the dimensionality has also been observed^[167,169]. All these works however reported a value of $\gamma \approx 0.4$. To add even more confusion, the mean-field (MF) theory –developed almost simultaneously^[84,111,112,117,118]– predicted *exactly*, in the limit of very high dimensions, the values $\gamma = 0.41269$ and $\theta = 0.42311$, as stated in Eqs. (1.67) and (1.68). So how to reconcile all the numerical results between themselves and explain the occasional agreement with MF theory? Fortunately, an elegant way out of this dilemma was soon devised^[150,169]: the MF predictions are correct also for low dimensional systems, and the apparent contradictions observed in the value of θ are caused by the presence of two different types of contact forces. The difference is that opening a contact associated to one or other type gives rise to either an extended or to a localized excitation; see Fig. 1.20. (For brevity, I will henceforth use “localized” and “extended” to distinguish the corresponding contact *forces*, although they obviously are only present at the contact point.) When the two contributions are considered separately, it is *consistently* found that localized (extended) forces are distributed according to $\theta_\ell \approx 0.17$ ($\theta_e \approx 0.42$) for any dimension $d \geq 2$. However, the relative fraction of localized and extended modes strongly depends on the dimensionality and, possibly, also on the preparation protocol (and maybe even on polydispersity in $d = 3$). This is caused by the fact that localized forces are predominantly associated with buckling excitations which occur, with very high probability, in particles with $z_\ell \equiv d + 1$ contacts. The reason is that buckling forces occur mainly because d contacts are approximately coplanar and thus, the remaining force should be almost orthogonal and of small magnitude. Note that as z_i increases it is ever more unlikely to achieve a similar nearly copla-

in mind that many of the arguments considered here also apply from the constraint satisfiability perspective of jamming (see Secs. 1.4.1 and 2.1). For instance, the UC regime (or when jamming is approached from below) corresponds to the SAT phase, while the OC one (or approaching from above), is equivalent to the UNSAT phase, where at least one constraint is not satisfied.

nar arrangement, and therefore the vast majority of bucklers have z_ℓ contacts; see details in [150, SI. II]. Importantly, the amount of buckling particles decreases very rapidly with dimension, because fluctuations away from the isostatic value $z_c = 2d$ are exponentially suppressed^[150]. Putting all of these arguments together we can satisfactorily explain, on the one hand, why MF is unable to provide a description of localized modes (namely, because they are an exclusive low d phenomenon) and, on the other, the discrepancy of values of θ . In particular, the results of Ref. [150] reveal that in $d = 2$ localized forces are dominant, and hence if the two forces are considered together it is found that $\theta \approx 0.19$; in $d = 3$, the same analysis yields $\theta \approx 0.25$ (cf. θ_{HS} and θ'_{HS} above), and $\theta \approx 0.30$ in $d = 4$. But these dissimilar values are actually simply caused by the dependence of the amount of bucklers with the dimension. Even more, when localized and extended forces are considered separately, in all these cases their probability function (pdf) follow very closely $p(f_\ell) \sim f_\ell^{\theta_\ell}$ and $p(f_e) \sim f_e^{\theta_e}$, where $\theta_\ell = 0.17462$ and $\theta_e = 0.42311$, in agreement with the MF prediction for the latter.

Unfortunately, once this matter was solved, no further test was made to carefully verify that the three values $(\gamma, \theta_e, \theta_\ell)$ match as $\varphi \rightarrow \varphi_J^\pm$. That is, the available numerical evidence strongly supports that this is the case but, as with many other phase transitions, there is no reason *a priori* to suppose that the jamming criticality is unaffected by the direction of the limit. Moreover, a precise estimation of their value is needed because the three exponents are related by mechanical stability bounds as derived in Sec. 1.4.5. MF predictions of θ_e and γ saturate the corresponding bound, Eq. (1.74a), and current estimates of θ_ℓ indicate that the same is true for the other relation, Eq. (1.74b). Hence, it is fundamental to perform a careful analysis of the values of $(\gamma, \theta_e, \theta_\ell)$ inferred from numerical simulations in order to assess the stability, if only marginal, of jammed configurations. Conducting such an analysis is especially important considering that packings of slightly polydisperse crystals are reported to exhibit a microstructure characterized by exponents that significantly differ from those of Eqs. (1.67) and (1.68)^[162,163]. Additionally, recent works have shown that many of the salient features of spherical packings depend sensitively on particle shape. For instance, introducing even an infinitesimal amount of asphericity changes the universality class^[165,166], in which the isostatic condition no longer holds. The relevance of investigating the extent of the jamming universality class and thoroughly testing its many theoretical predictions is therefore needed^[102], specially considering that jamming criticality is, up to date, the most precise prediction reaped from the replica method applied to realistic materials models.

With this in mind, in this chapter I will present a detailed analysis of the finite-size scaling of the distributions of interparticle gaps and contact forces. These distributions are one of the fundamental consequences of the presumed non-trivial critical behaviour of jammed packings, hence their testing is a key step toward rigorously validating a whole set of critical properties. Although a similar analysis has been carried out for the perceptron^[140] and for the gaps distribution of a two-dimensional binary mixture^[165], no systematic result exists for jammed packings of spherical particles, nor for amorphous packings with other sources of disorder. As I mentioned above, in Ref. [2] we considered several models, both in $d = 2$ and $d = 3$, with the aim to define precisely which are the most robust features of jamming criti-

cality, and thus better demarcate its physical universality. For instance, to properly account for size effects in $2d$ systems logarithmic corrections^[155,221–223] are important, thus confirming that $d = 2$ is the upper critical dimension^[150,155,167]. (Further evidence comes from recent studies^[114,224] where a different criticality is observed in quasi-one-dimensional systems.) Similarly, we also found that crystalline ordering breaks the jamming universality and renders jammed configurations unstable according to the criterion of the usual bounds, Eqs. (1.74). Nevertheless, I will not consider such systems here and instead only focus in $3d$ systems of monodisperse spheres, and the analogous MK model.

This chapter is organized as follows, in Sec. 4.2 I provide some details about the generation of jammed packings and in Sec. 4.3 I explain how to study the size effects we are interested in. Then in Sec. 4.4 I present a meticulous analysis of the finite size effects in the critical distributions of packings of spherical particles. An unexpected outcome is the striking contrast of such effects on the distributions f_e and h . Then, in Sec. 4.5 I carry out a similar study for the MK configurations, which shows that size effects are much pronounced in this model and that localized forces exhibit a different distribution. A tentative explanation is advanced for both features in terms of the particular properties of the MK model. Finally, a discussion and brief conclusion are given in Sec. 4.6.

4.2 Methods for generating jammed packings

I will first consider the case of three-dimensional configurations of N spheres of equal diameter σ in a cubic box under periodic boundary conditions. In a certain sense, this choice is the minimal model for producing jammed packings. Lower-dimensionality systems inevitably crystallize unless polydisperse mixtures are used, but ordering can be avoided for monodisperse spheres in $d \geq 3$. Sphere positions then serve as the only source of disorder. For configurations initially in the UC regime, a hard-sphere potential is used and a combination of molecular dynamics (MD) and linear optimization algorithms are employed to reach ϕ_J . More precisely, we start from a low-density configuration ($\varphi = 0.2$) with particles' position assigned randomly and uniformly in the box volume. Then, we use event-driven MD with a Lubachevsky–Stillinger growth protocol^[94] (see Sec. 2.2 for details) to increase the (reduced) pressure up to $p = 500$. This first step is performed with a fast compression rate ($\dot{\sigma} = 10^{-2}$) in order to avoid any partial crystallization and is then followed by a second, much slower ($\dot{\sigma} \leq 10^{-3}$), growth protocol until $p \gtrsim 10^7$. (Details about the influence of $\dot{\sigma}$ and the final p on the jammed configuration thus produced are reported in Sec. 2.3.1.) At this point the configuration is used as input for the iterative Linear Programming (iLP) algorithm described in Sec. 2.1. As discussed there, at each step, the LP algorithm finds the optimal rearrangement of particles that allows to maximize their radius, considering a linearised version of the non-overlapping constraint between all pairs of particles. Upon convergence, this algorithm produces a jammed configuration, because neither particle displacements nor size increases are possible. With this we can easily reconstruct the full network of contacts at jamming, because genuine contact forces can be extracted from the active dual variables associated to the non-overlapping constraints. As I proved in

Sec. 2.1.2, our iLP algorithm always produces a mechanically stable configuration and, the majority of times, it has a single state of self-stress (1SS). As discussed in Sec. 1.4.3, such a state is characterized for having one contact above isostaticity, *i.e.* when the total number of constraints in a system (N_c) matches its number of degrees of freedom (N_{dof}), and determines where the jamming criticality comes about. The extra contact with respect to isostaticity is required in order to achieve a finite bulk modulus^[154,174]. Put differently, the system density is an additional variable that needs to be fixed, and thus requires one additional contact above isostaticity^[172] (Sec. 2.1.2).

On the other hand, when in the OC phase² and given the position vectors, $\{\mathbf{r}_i\}_{i=1}^N$, we considered an harmonic contact potential,

$$U(\{\mathbf{r}_i\}_{i=1}^N) = \frac{\epsilon}{2} \sum_{i,j} (\sigma - |\mathbf{r}_i - \mathbf{r}_j|)^2 \Theta(\sigma - |\mathbf{r}_i - \mathbf{r}_j|), \quad (4.2)$$

where ϵ is a constant that defines the energy scale, and Θ is the Heaviside step function. Hence, a pair of particles only interacts if there is an overlap between them. Starting with $\varphi > \varphi_J$ and a Poisson random distribution of spheres, a series of energy minimization steps and packing fraction reduction steps are performed until the system has 1SS. More precisely, at a given density the FIRE algorithm^[151] is used to achieve force balance in the configuration. The energy of the configuration is then calculated and the known scaling relation^[150], $U \propto (\phi - \phi_J)^2$, is used to determine by how much the sphere radii should be uniformly decreased to reduce the system energy by a fixed fraction. Thus, after several iterations of this procedure, the packing has precisely $N_c = N_{dof} + 1 \equiv N_{1SS}$ contacts. After removing rattlers (see below), the dynamical matrix^[4] is used to ensure that the packing is jammed. This algorithm is implemented in the pyCudaPacking software using general purpose graphical processing units and quad-precision calculations, see details in Refs. [58, 225, 226].

Once the contact forces have been computed for a given configuration, they are normalized so that their mean value equals 1. We then proceeded to classify them as localized or extended. The criterion we used is based on the coordination number of each particle, z_i . Thus, if a particle has $z_\ell = 4$ contacts, it is automatically considered a buckler and all its contact forces are classified as localized. In turn, if a particle has $z_i > z_\ell$ all the related forces, except the ones possibly associated with bucklers, are classified as extended. Clearly, this criterion is far from perfect, because (i) not all particles with z_ℓ neighbours necessarily have a buckling contact and (ii) not all bucklers have z_ℓ contacts. However, the fact that a buckling force is associated with a specific geometric configuration as discussed above and that relatively few particles have less than $z_c = 6$ contacts justifies that by doing so, we are capturing most of the relevant localized modes. Of course, more accurate methods have been devised^[162,167,169], but as shown in Ref. [150] this criterion often suffices. Besides, another way to extract the distribution of localized and extended forces is to consider the smallest and median force *per particle*. When this is done, it is found that $p(f_{small}) \sim f_{small}^{\theta_\ell}$ and $p(f_{med}) \sim f_{med}^{\theta_e}$. We verified that our results

²All the simulations in the OC regime were performed by our collaborators, so I just describe them briefly.

hold with either methodology, so I will only present the data obtained from the connectivity criterion.

On the other hand, computing the dimensionless gaps, h (defined in Eq. (1.49)), in each configuration is straightforward once the jamming point has been reached: it is enough to compute the distance between *non-touching* particles, rescale it by σ_J and subtract 1. However, given that the jamming power-law only applies for small gaps, it is convenient to introduce a cut-off to the distances considered. Here, we set it to $2\sigma_J$, which amounts to $h_{max} = 1$, a value well beyond the critical regime (see Figs. 4.6 and 4.14). As a side comment, note that when the iLP algorithm is employed, another advantage of using active dual variables to define contacts between particle pairs is that such pairs can be directly excluded from the list of distances to be considered.

Finally, for the MK model we restricted our analysis to configurations of HS, and hence, approaching jamming from below. As described in previous chapters, this model is characterized by particles interacting through a randomly shifted distance, as in Eq. (2.18). For our simulations, we assumed the random shifts to be quenched vectors uniformly distributed in space, corresponding to the MF limit. To generate the packings we followed essentially the same procedure relying on the MD+iLP algorithm. With respect to HS packings, there were only a couple of differences in the MD compression part of MK (the actual protocol is illustrated in Fig. 2.6). First, that a single value of $10^{-3} \leq \dot{\sigma} \leq 2 \times 10^{-3}$ was used, because there is no risk of crystallization in this model, so there is no need to begin with a fast compression. Second, the initial configuration for the dynamics was a planted configuration with density $\varphi = 2.5$. In this way, we drastically reduced the time of the MD simulations and guaranteed to go beyond the dynamical transition point^[109,110] (see Fig. 2.6). However, I should mention that a disadvantage of doing so is that the fraction of rattler increases considerably to about 15%–20%. Nevertheless, we always obtained a configuration having 1SS with the help of our iLP algorithm. As a final remark I mention that generating a jammed packing of N particles following this procedure takes considerably longer than for a usual HS configuration of the same size. This is due to the fact that any given particle is surrounded by many more particles than in the case with Euclidean distance. Besides, the neighbours-list approach^[194] for the MD compression is also significantly slower than the cell-based one^[94]; more details are given in Sec. 2.2.

4.3 Methods for studying finite size effects

We used the methodology just described to produce several jammed configurations for each system size, parametrized by N . For systems of spheres with the usual Euclidean (resp. MK) distance, we tried $N \in \{2^8, 2^{14}\}$ (resp. $N \in \{2^8, 2^{12}\}$). Moreover, to ensure that we sampled all the systems with the same accuracy, for a given value of N , M_N independent packings were produced in such a way that $N \times M_N \simeq 2.2 \times 10^6$ (5.5×10^6) for standard HS (SS) configurations, while $N \times M_N \simeq 10^6$ for the MK ones. The reason of the smaller sizes and reduced dataset in this latter model is that both the MD and iLP algorithms perform much slower than in the common case as mentioned above. These data provide enough accuracy to study

the forces and gaps across many orders of magnitude, and finite-size corrections can be systematically identified.

It is known that using a logarithmic binning when testing for power-laws in a pdf often leads to poor comparisons as a consequence of the loss of resolution when grouping enough data in a single bin^[227]. To tackle this difficulty, we instead utilized the cumulative distribution function (cdf). Note that if a random variable, x , is distributed according to a pdf of the form $\rho(x) \sim x^\alpha$ for $\alpha > -1$, then its cdf follows $c(x) \sim x^{1+\alpha}$. The cdf allows the accuracy of the data acquired to be preserved and, in the case of independent random variables, maximum likelihood methods can be used to obtain a more precise estimation of α than by least-squares fit^[227]. However, for the cases I consider here, a size scaling analysis leads to better results (see discussion in Sec. 4.5.2.2).

To empirically estimate the cumulative distributions, we opted to put together the data of all the M_N samples for each system size. Doing so improves the statistics of the smallest values of the structural variables we investigate here, so we are able to probe with some detail the left tail of all the distributions. Such analysis is an important part of our results as I describe below. However, given the M_N independent samples for each system size, this is not the only way in which one could estimate the corresponding cdf. Nevertheless in Appendix 4.A I show that this approach produces essentially the same results than when the cdf of single samples are considered.

Now, when fitting a distribution to empirical data it should be considered that, even if x strictly follows such a distribution all the way down to $x = 0$, finite sampling inevitably leads to deviations. Here, the situation is further complicated by our consideration of marginals of *correlated variables*. Gaps and forces distributions of finite N configurations are indeed prone to exhibit deviations from their expected form due to both finite sampling and system-wide correlations. As I will argue in the remaining part of this chapter, introducing a scaling function (as is usually done in the study of critical phenomena^[175,228]) can account for both effects, allowing the N dependence of the cdf to be carefully analysed. To derive the size scaling of the distributions of x , we first note that in a sample of size $N \gg 1$, we can estimate the order of the smallest value observed in the data, x_{\min} from the probability mass assigned to the extremes of the distribution:

$$\int_0^{x_{\min}} \rho(x) dx \sim x_{\min}^{1+\alpha} \sim \frac{1}{N}. \quad (4.3)$$

It then follows that $x_{\min} \sim N^{-1/(1+\alpha)}$. Note that strictly speaking in this last equation N should be replaced by N_c when analysing, for instance, the distribution of contact forces. However, given that $N_c \sim dN$ and that we are mostly concerned with the scaling exponent, we can safely neglect the associated proportionality constants. The behaviour of the gaps distribution is expected to be similar, in the sense that particles almost in contact should be self-averaging. Following the traditional path for analysing size scalings, we can write the pdf as

$$\rho(x) \sim N^\beta \tilde{\rho}\left(xN^{\frac{1}{1+\alpha}}\right) \quad (4.4)$$

where $\tilde{\rho}$ is the scaling function of the pdf such that $\tilde{\rho}(x) \sim x^\alpha$ for $x \gtrsim 1$. The exponent β can be easily determined by requiring that $\rho(x)$ exhibit no N dependence.

We thus get that $\beta = -\frac{\alpha}{1+\alpha}$, whence the expressions used for the scalings studied in Ref. [140] are recovered. For the cumulative distributions, repeating the above analysis for $c(x) \sim N^{-\delta} \tilde{c}(xN^{\frac{1}{1+\alpha}})$ gives $\tilde{c}(x) \sim x^{1+\alpha}$, and it immediately follows that $\delta = 1$, whence the relevant scaling relation is

$$c(x) \sim N^{-1} \tilde{c}(xN^{\frac{1}{1+\alpha}}). \quad (4.5)$$

Using the correct α should remove any dependence on N in this last equation. That is, data for different system sizes can be rescaled in such a way that they follow a common curve, \tilde{c} . Finding a good collapse of the curves for different N thus indicates that deviations from the expected algebraic behaviour are caused by our systems necessarily being off the thermodynamic limit, but not because variables follow a different power-law scaling. Additionally, showing that the system size influences the cdf of a given variable strongly evinces that such a variable is correlated across the whole system. Hence, an upper bound to the correlation length can then be estimated.

Notice that Eq. (4.3) suggests another way to estimate the value of α from the datasets obtained from different system sizes. And in fact, this was the method employed in Ref. [140] for a similar study in the perceptron. The idea is the following. For a fixed N , instead of computing the cdf of a given variable by putting together the data of the M_N samples, the typical x_{\min} is approximated by its sample average, that is,

$$\overline{x_{\min}} \equiv \frac{1}{M} \sum_{m=1}^M x_{\min}^{(m)} \sim N^{-\frac{1}{1+\alpha}}, \quad (4.6)$$

where $x_{\min}^{(m)}$ is the smallest value of the random variable x in the m -th sample. In practice, this expression can be tested by showing that $N^{\frac{1}{1+\alpha}} \overline{x_{\min}} \approx \text{constant}$ for different values of N . Unfortunately, this method cannot be used in the cases we study here due to corrections to Eq. (4.4) for very small values, as I now describe.

Given that our jammed configurations have one extra contact than N_{dof} (see Secs. 1.4.3, 2.1.2), theoretical predictions^[141,165,229] establish that the microstructural critical variables should exhibit a different scaling in their left tails. For instance, MF theory predicts that interparticle gaps are distributed as $h^{-\gamma}$ only for values larger than a cut-off $h^* \sim \delta z^{\frac{1}{1-\gamma}}$, where δz is the excess of contacts in a system. In our case, $\delta z \sim 1/N$, so instead of Eq. (1.67) the pdf describing the distribution of h reads,

$$g(h) \sim \begin{cases} N^{\frac{\gamma}{1-\gamma}} g_0(hN^{\frac{1}{1-\gamma}}), & hN^{\frac{1}{1-\gamma}} \ll 1 \\ h^{-\gamma}, & hN^{\frac{1}{1-\gamma}} \gtrsim 1 \end{cases}; \quad (4.7)$$

where $g_0(x) \sim 1$ for $x \ll 1$ ^[165]. Analogously, for extended forces Eq. (1.68b) should be replaced by

$$p(f) \sim \begin{cases} N^{\frac{-\theta_e}{1+\theta_e}} p_0(fN^{\frac{1}{1+\theta_e}}), & fN^{\frac{1}{1+\theta_e}} \ll 1 \\ f^{\theta_e}, & fN^{\frac{1}{1+\theta_e}} \gtrsim 1 \end{cases}; \quad (4.8)$$

where $p_0(x) \sim 1$ for very small values is to be expected. Equations (4.7) and (4.8) are indeed consistent with Eq. (4.4), and, repeating the same arguments as above, it is straightforward to derive that both regimes can be captured by Eq. (4.5) using a single scaling function, such that

$$\tilde{c}(x) \sim \begin{cases} x, & x \ll 1 \\ x^{1+\alpha}, & x \gg 1 \end{cases}. \quad (4.9)$$

That is, using the correct α in Eq. (4.5) accounts for size effects that give rise to deviations from the main power-law scaling as well as the appearance of the linear regime in the left tails. By plotting Nc as a function of $N^{\frac{1}{1+\alpha}}x$ both corrections can thus be tested from a single scaling collapse.

Importantly, from this last equation and the preceding arguments, it is easy to conclude that the left-tail corrections become important in the same regime than the finite N effects, thus what we will see in a realistic scenario is a crossover from the MF algebraic distributions to a flat one, given the form of g_0 and p_0 . This is the reason why Eq. (4.6) yields a very poor estimation of the critical exponents.

Before concluding this section I want to emphasize that for microscopic variables of jammed configurations the situation is conceptually different from that of standard critical phenomena, because the systems are already *at* the critical point. That is, all the results that I will present here do not deal with the question of how the distributions of contact forces and gaps converge to the expected ones as we move away from φ_J . Instead, I am going to analyse how the system size affects the range over which power-law scalings are followed. Equation (4.5) can nevertheless be used to estimate the scaling functions of the cdf of gaps and forces, obtained by integrating Eqs. (1.67) and (1.68b), respectively. I should mention as well that most of our results are concerned with extended forces and gaps since theory and previous numerical studies suggest that these variables are critically correlated across the whole system. In contrast, the distribution of buckling forces is expected to be independent of the system size, given the localized nature of their corresponding floppy modes. As I show in the next sections, our findings verify this picture.

4.4 Finite size effects in spheres systems

4.4.1 Some structural properties

In this section, I briefly analyse two quantities of the HS jammed packings, namely, the distribution of coordination number, z , and the radial distribution function (RDF), $g(r)$, both for all the values of N tested. The idea is to verify that the system's size does not affect the expected properties of configurations with at the jamming point. Hence, the results here are not concerned with the size scaling, but are more of a “security check” for the scalings presented in the next part. The pdf of z is shown in Fig. 4.1 for all values of N . Clearly, no signatures of size effect are present, and in all the cases it is found that $z_c = 2d = 6$ is the likeliest value. This last feature is expected on the basis of the constraints counting argument at the jamming transition, introduced in Sec. 1.4.3. The very small amount of particles with $10 \leq z \leq 12$ (the latter being the $3d$ kissing number) suggest that no ordering is present in our configurations.




Figure 4.1. Probability density function of the coordination number for different system sizes. The mode is located at $z_c = 2d$ as expected from the marginal stability of our jammed packings. No visible size effects are found.




Figure 4.2. Radial distribution function of HS jammed packings (data average over at least 100 samples for each system size). Curves of different values of N have been displaced vertically for clarity, while the domain of each curve is limited by $r < L/2$ due to the periodic boundaries, with L being the size of the box containing the system. No influence of the system size is found. Grey dotted lines indicate the expected position of peaks or discontinuities in $g(r)$ present in *amorphous* packings, while red dashed lines are the analogous positions for *ordered* structures; cf. Fig. 2.10 and see text for details.

This last feature can be tested more stringently by looking at the peaks of the RDF's, computed according to Eq. (1.17) and presented in Fig. 4.2. (In this case curves associated to a different value of N have been displaced vertically for

clarity.) Notice first that no size dependence is found in this case either, except for the trivial fact that periodic boundary conditions constrain r to be restricted to less than half of the box. That is, given that $\sigma \sim N^{-1/3}$, the smaller N the faster periodic effects are important. But this does not affect the behaviour of the RDF. Now, the positions of the peaks commonly found in *disordered* packings are indicated by grey dotted lines, showing an excellent agreement. In contrast, the red dashed lines correspond to the positions of peaks associated with the formation of crystalline domains (see Fig. 2.10 and the related discussion). The absence of peaks at such locations confirms that no partial crystallization is present in our samples, for all values of N .

4.4.2 Finite size scaling of the critical distributions

4.4.2.1 Distribution of extended forces

Once we know that our packings do not present signatures of crystallization, let me begin by analysing the size effects in the cdf of the extended forces, $\{f_e\}$. Figure 4.3 shows the distributions of f_e obtained coming from below (UC, upper left panel) and from above (OC, upper right panel). Comparing the results with the theoretically predicted power-law (dashed red) reveals an outstanding agreement over at least three decades. More importantly, no visible signature of finite-size corrections can be detected over the range of N considered. To verify more stringently the absence of finite-size effects, we attempted to collapse the different curves by rescaling the extended forces and their cdf following Eq. (4.5), obtaining the curves reported in the bottom panel. This latter figure contains our first main result: strong evidence that the same critical distribution of forces is found independently of whether the jamming point is generated from the UC or OC regimes.

Yet, it is clear that our packings exhibit an excess of very small forces (an effect more noticeable when jamming is reached from below; cf. upper panels of Fig. 4.3), echoing earlier observations^[140,150,162,167]. Note that the scaling of Eq. (4.5) does not remove these deviations from the predicted distribution, and that they occur, roughly, at the same scaled force: $N^{\frac{1}{1+\theta_e}} f_e \lesssim 1$. It is therefore likely that forces are subject to size effects caused by the onset of a different behaviour, $p(f) \sim 1$ (see Eq. (4.8)). One way to analyse more carefully these two regimes, and in particular the linear left tail of the cdf, is to divide the cdf by the corresponding power law. That is, assuming $c(x) \sim x^{1+\alpha}$, plotting $x^{-1-\alpha}c(x)$ should result in a rather constant function. Performing such rescaling in the cdf of f_e yields the curves shown in Fig. 4.4, where the exponent for the rescaling factor used in each panel is, from top to bottom, $(1 + \theta_e, 1, 1 + \theta_\ell, 1.25)$. The reason for including $1 + \theta_\ell$ and 1.25 is to check whether the left tail is due only to localized modes, or, to the fact that the “ $z_i = z_\ell \implies$ localized” criterion fails for the smallest forces, respectively³. Thus, the topmost figure confirms, in the right part of the distribution, the MF prediction for the extended forces, while the panel below it tests the linear behaviour mentioned for the left tail. Notice that in this latter case the cdf has also been scaled by $N^{\frac{\theta_e}{1+\theta_e}}$,

³Recall that in the introduction I mentioned that in $d = 3$ the distribution of the two types of forces combined has an exponent of 0.25. This was shown in Ref. [150] and we verified that the same is true for our data.




Figure 4.3. Cumulative distributions of contact forces associated with extensive excitations in monodisperse configurations of frictionless spheres for different system sizes N . The upper left panel (resp. upper right panel) show the data as the jamming point is reached from below or UC regime (resp. above or OC regime). To better distinguish between the two different regimes, results belonging to the UC (OC) phase are identified by solid lines (circular markers). Bottom panel: Rescaling data from the upper panels according to Eq. (4.5) clearly collapses the data. The red dashed line corresponds to the power-law scaling of Eq. (1.74a), and shows an excellent agreement between the MF predictions and our numerical results. The coincidence of results from the UC phase and OC phase for various N confirms that θ_e is the same when jamming is reached from either direction. In the left tail of the distributions of the bottom panel I also include a comparison with the linear scaling expected for very small values, following Eq. (4.8).

as predicted by Eq. (4.8), in order to verify that the $f_e \rightarrow 0$ tail leads to a constant of order 1⁴. Data from the UC phase (lines) show a reasonable agreement, *i.e.* a slight plateau, for $f_e \in [10^{-4}, 10^{-3}]$ approximately. Nevertheless, the same comparison with data from the OC regime (circles) is slightly worse. Moreover, the second to last and bottom panels of the same figure show that, arguably, the cdf corresponding

⁴Clearly, this “extra” factor has the effect of splitting vertically the curves of different values of N . However, I verified that this effect can be removed by rescaling the *horizontal* axis by $N^{\frac{1}{1+\theta_e}}$, as expected.

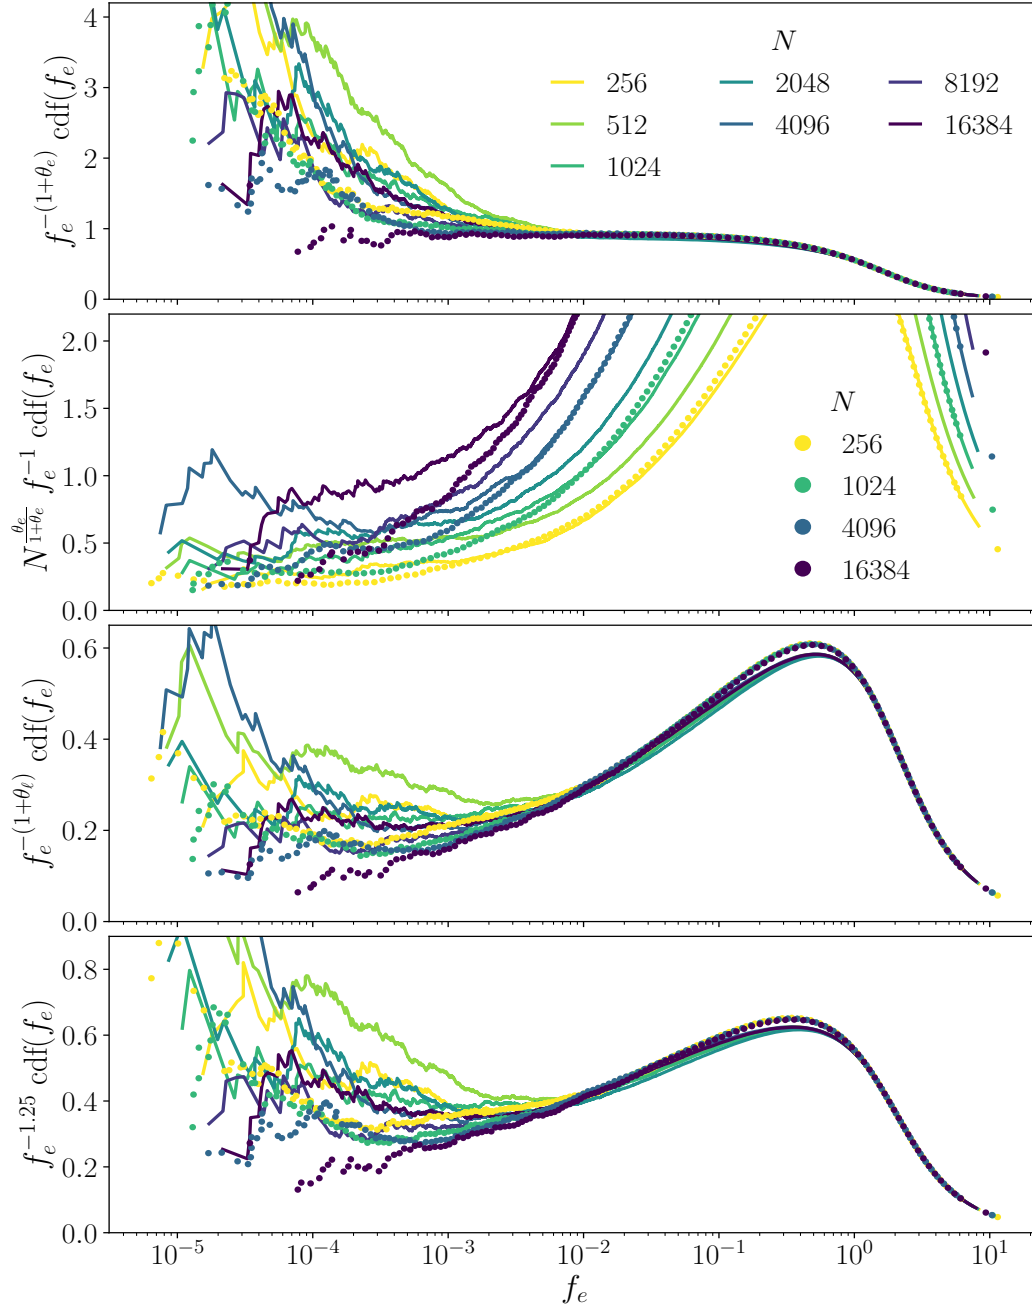


Figure 4.4. Same distributions of Fig. 4.3 but rescaled by f_e^α to examine more carefully the different scalings present in the cdf. The value of α in the different panels, from top to bottom, is $1 + \theta_e$, 1, $1 + \theta_\ell$, 1.25, for the reasons explained in the main text. Notice that the cdf values in the second panel (from top to bottom) have also been multiplied by $N^{\frac{\theta_e}{1+\theta_e}}$, following Eq. (4.8); this serves to verify that once rescaled, the behaviour of the left tail should be a constant of order 1. The data from the UC phase (lines) seem to be consistent with a linear scaling in the left tail, while OC data suggest that localized modes might also be present, inducing small deviations from the MF behaviour.




Figure 4.5. Same distributions of Fig. 4.3 but subtracting $1/N$ to the cdf. This is a simple strategy for isolating the main scaling domain, by removing the contribution of the extremes. This reduced cdf clearly shows (left panel) that finite N effects can be identified in the distribution of extended forces and, more importantly (right panel), accurately accounted for by the scaling function of Eq. (4.5), using the MF value of θ_e .

to SS data is better rescaled by assuming some influence of the localized modes. In contrast, this does not seem to be the case for HS data. Unfortunately, with the available data it is impossible to prove (or even to disprove) the linear form of the left tail, which calls for a more assiduous future study of the true behaviour of the extremes of $p(f_e)$.

As a final “trick” to isolate the main finite- N effects in the cdf notice that, according to Eq. (4.3), by subtracting $c(x) - a/N$ the influence of the left tails can be teased out, with a a constant of order 1 to be determined. In other words, given that the weight of the extremes is of order $1/N$, by removing their contribution to the cdf only the main, MF power-law should remain, albeit in a reduced domain. Hence, by applying the same scaling arguments in such main region a more rigorous analysis is obtained. In Fig. 4.5 I show the outcome of this technique when applied to the usual cdf (left panel), as well as to the distribution rescaled according to Eq. (4.5) with the MF value of θ_e (right panel). Notice that in the first case, the size effects are now evident, but the second figure shows a remarkable collapse onto the scaling function, confirming the MF prediction for θ_e . I assumed $a = 0.25$, given the value of the approximate plateaus in the second panel of Fig. 4.4, but the results are rather insensitive to this value.

4.4.2.2 Distribution of gaps

I will now analyse the distribution of interparticle gaps, $\{h\}$, and their associated size effects. Fig. 4.6 shows the corresponding cdf’s obtained from the UC (upper left




Figure 4.6. Cumulative distributions of interparticle gaps for the same configurations of Fig. 4.3, as their jamming point is reached from below or UC phase (upper left) and from above or OC phase (upper right). Lower panel: rescaling both datasets according to Eq. (4.5) shows that finite-size corrections can be accounted for in all cases. For comparison, the power-law scaling derived from MF theory, Eq. (1.67), is also shown (red dashed line). Once again, the fact that datasets from both phases, *i.e.* UC (markers) and OC (lines), neatly superimpose confirms that the exponents at the jamming point are the same, independently of how ϕ_J is approached. Additionally, the secondary scaling regime $g(h) \sim 1$ of Eq. (4.7), also predicted by MF theory, can be observed for very small values. Its associated linear cdf is shown (magenta dash-dotted line).

panel) and OC (upper right) from the same configurations as above. Once again, the data are in very good agreement with the predicted distribution, Eq. (1.67), independently of the direction in which jamming is approached. More interesting however, is the fact that the cdf(h) exhibits strong signatures of finite N effects, in contrast to $p(f_e)$. Note that the scaling correction given in Eq. (4.5) using the MF value of γ precisely corrects for such effects over almost *seven* orders of magnitude (lower panel of Fig. 4.6). We thus verified that the MF value of γ holds on both sides of the jamming transition as well.

On the other hand, the growing deficit of very small gaps as the system size decreases is, very likely, another manifestation of the cut-off in the main power




Figure 4.7. Same distributions of Fig. 4.6 but rescaled by $h^{-(1-\gamma)}$ (top plot) and $N^{-\frac{\gamma}{1-\gamma}} h^{-1}$ (bottom one) to examine more carefully the different scalings present in the cdf's. The N dependence of the rescaling factor in the second case follows from Eq. (4.7); see main text. As with the case of extended forces, configurations of HS (lines) show a better agreement with the linear dependence assumed, while the smallest packings generated from the OC phase using SS (circles) exhibit significant deviations in this extremal regime.

law of $g(h)$. It leads to a secondary linear regime, as stated in Eq. (4.7), that is in reasonable agreement with the numerical results shown in the lower panel of Fig. 4.6. Yet, to assess the presence of such regime more thoroughly we can once again resort to the method of dividing the cdf by h raised to the corresponding power of the regime we want to investigate. When doing so, we obtain the plots shown in Fig. 4.7, where the rescaling factor is $h^{-(1-\gamma)}$ in the upper panel (to verify the theoretical exponent) and $N^{-\frac{\gamma}{1-\gamma}} h^{-1}$ in the lower one (to test for the presence of the linear left tail, including the prefactor spelled out in Eq. (4.7)). Once again, by including such correction the $h \rightarrow 0$ behaviour of the rescaled tail approaches a constant value of order 1, for almost all values of N . In fact, the only discrepancies observed are for the two smallest sizes of the systems obtained from the OC regime. But larger systems of SS as well as all the cases with HS reproduce the expected plateau. I deem important to emphasize that to remove as much as possible any size dependence of this secondary regime the $N^{-\frac{\gamma}{1-\gamma}}$ term is needed. That is, if the cdf are rescaled only by h^{-1} an increase in the left tails is observed, *i.e.* an *spurious* sub-linear behaviour.

At any rate, the different role of the system's size in the distributions of forces

and gaps is, nonetheless, striking. Our numerical results indicate that distances between nearby spheres are significantly modified in finite N configurations and, consequently, so is the distribution of gaps. This is an intriguing finding that I will attempt to rationalize heuristically relying on the response to contact openings in jammed packings. The argument goes as follows. In the thermodynamic limit, a system has always enough space to relax any perturbation caused by a contact opening, and hence it is always able to re-accommodate enough particles –even if this requires bringing many of them infinitesimally close to each other– in order to guarantee stability. In a finite system, by contrast, no such unconstrained relaxation can take place. Rearranging an extensive fraction of particles necessarily influences the pair of spheres involved in the contact just opened. Therefore, there is a certain scale below which the occurrence of small gaps is disfavoured. If the system was further relaxed, then at least one extra contact would form. This is consistent with the uniform distribution hypothesis, since such distribution decreases the weight of small enough gaps in comparison with the MF one, that actually predicts a divergence as $h \rightarrow 0$.

At this point, I wish to stress that our results demonstrate the existence of two different types of finite-size corrections to the distributions of extended forces and gaps. The first is a consequence of large scale correlations and can thus be readily taken into account by the scaling of the cdf given in Eq. (4.5). Although this correction is practically absent in the forces distribution, for $g(h)$ it is the main source of deviation from the theoretical prediction. The second is a consequence of the critical scalings of Eqs. (1.67) and (1.68b) being cut off at very small values. This effect, which is likely related to the extra contact with respect to N_{dof} (see Sec. 4.3), affects both microstructural variables. I get back to this point in Sec. 4.5, after having considered its signature in the MK model.

4.4.2.3 Distribution of localized forces

The only critical distribution that remains to be analysed is the one associated to the set of localized forces, $\{f_\ell\}$. As I have anticipated above, no size effects are expected for this variable and this claim is nicely confirmed by our numerical results reported in Fig. 4.8. Note that in this case the estimation of $\text{cdf}(f_\ell)$ is very close to the predicted power-law for forces as small as 10^{-4} (cf. the case of extended forces in Fig. 4.3). Besides, for even smaller values the curves corresponding to different N are disperse both above and below the theoretical line, which indicates that no systematic size dependence is present.

Before concluding this section, it is worth emphasizing that our numerical results are in excellent agreement with the MF, $d \rightarrow \infty$ predictions for the power-law scaling of the distributions of both the extended forces and the interparticle gaps, as well as analogous previous estimates concerning the localized forces. These results confirm that the jamming criticality of these microstructural variables is also valid for low dimensional systems all the way down to $d = 3$, in agreement with earlier albeit less accurate studies^[153,167,169]. Because results from both OC and UC phases superimpose onto each other, we can further conclude that the critical behaviour is controlled by the same exponents on both sides of the jamming point.




Figure 4.8. Cumulative distribution function of localized forces, $\{f_\ell\}$, in configurations obtained from the UC (lines) and OC (circles) phases and for different system sizes, as indicated by the colours in the legend. No signatures of finite N dependence are observed.

4.5 Finite size effects in the MK model

I will now present the analogous results obtained from the jammed configurations in the MK model. Here however, packings were produced only from the UC phase, so only data for HS will be considered. In Sec. 4.5.1 I begin by analysing some properties of contact vectors which are distinctive of the MK configurations. Then, in Sec. 4.5.2 I present the main results of this section, concerning the size effects in the critical distributions.

4.5.1 Distribution of contact vectors

The first notable property of the MK packings is their connectivity. First, because the possible number of contacts per particle is no longer constrained by the respective kissing number. But also because $z_c = 2d$ is no longer the most common value of the coordination number. Instead, our results show that $z_\ell = 4$ and $z = 5$ are the likeliest values, at least in our $3d$ configurations. Both features are illustrated in Fig. 4.9 where it is also shown that these properties are independent of the systems' size. Note that the distribution of z is monotonically decreasing, in contrast with HS systems (see Fig. 4.1). Along the same lines, given that $\langle z \rangle = z_c + \mathcal{O}(1/N)$, the broader domain of the distribution shifts the mode to lower values than z_c , differing from previous results^[150] in $d = 2, \dots, 8$. This is an unexpected result given that all our packings have 1SS, and therefore their average connectivity is indeed z_c with corrections of order $1/N$.




Figure 4.9. Probability distribution of the coordination number per particle in MK configurations of different size, identified according to the legend. The expected value from isostaticity arguments (z_c) in $3d$, as well as the corresponding kissing number are highlighted. Cf. with Fig. 4.1 for the analogous results in HS.

For future use, I will now study the geometric arrangement of contacts, for which I will follow the same methodology than in Sec. 3.1.2. Hence, I will first analyse the distribution of dot product between pairs of contact vectors, $\mathbf{n}_{ij} \cdot \mathbf{n}_{ik}$, for $j, k \in \partial i$. The pdf obtained from the data of the $N = 2048$ configurations is shown in Fig. 4.10, where I have also included the corresponding distributions when localized and extended modes are considered separately. Comparing with Fig. 3.3, where the corresponding results in HS are presented, it is clear that random shifts in the MK configurations allow particles to form rather different structures. The most salient feature is that no peak is observed at 0.5, and actually the distribution continues all the way down to 1. This is expected, because many neighbours of a given particle can overlap (from the point of view of such particle), and these results show that they can even be concentric. As in the HS case, the distributions of the full set of particles and the ones associated to extended forces are very similar to each other, while the one of bucklers is in this case notoriously different. Note that for this latter type of particles more weight is given to contact vectors forming angles greater than $\pi/2$, which supports the near coplanarity assumption of bucklers. However, an accompanying peak around 0 is clearly missing, signalling that other configurations that fulfil the mechanical equilibrium condition are possible.

Finally, I consider the sum of such dot product per particle, *i.e.* the variable S_i defined in Eq. (3.2). Fig. 4.11 shows the resulting distributions in both MK and HS configurations, and considering different system sizes. These results show, once again, that size effects do not influence the structural variables considered in the previous chapter. More importantly, we can observe that even if S_i can take values greater than -0.5 in MK configurations, the effect on the distribution is rather small. In fact, bucklers in both models seem to have a very similar pdf, which is rather unexpected from the corresponding distributions shown in Figs. 3.3 and 4.10.




Figure 4.10. Probability distribution of dot products between contact vectors in MK configurations of $N = 2048$ particles. The contributions of bucklers (orange) and extended (blue) modes are also plotted separately. Cf. with Fig. 3.3 for analogous results in HS.

All these results will be important to understand some of the findings presented in the following part, but the comparison included in this last figure will have a special role in the discussion of the distribution of $\{f_\ell\}$ (see Sec. 4.5.2.1).




Figure 4.11. Pdf of S_i of bucklers in MK (left) and HS (right) jammed packings of different sizes. Inset: comparison of the same distributions in both models (using a linear scale); data from configurations of $N = 2048$ particles.

4.5.2 Finite size scaling of the critical distributions

4.5.2.1 Distribution of extended forces

I will now proceed to analyse the size effects in the critical distributions, following the same methodology than for HS (Sec. 4.4.2). Hence, I begin by the considering the distributions of extended forces for $N = 256, \dots, 4096$, shown in the left panel of Fig. 4.12. Notice that in contrast with the HS results, here the size effects are noticeable, even if small. Interestingly, this seems to be a feature exclusive of the MK packings, since the other models studied in Ref. [2] do not show such dependence on N . The influence of the size effects in $\text{cdf}(f_e)$ is clearer when the size scaling is carried out, following Eq. (4.5). The result is shown in the right panel of the same figure, where the collapse obtained using the MF value of θ_e is again outstanding. The reasons explaining the pronounced finite N effects in the MK systems however is postponed to Sec. (4.5.2.2), because similar findings hold for $g(h)$.




Figure 4.12. Cdf of the set of extended forces in MK configurations of different size (left) and the resulting scaling (right) following Eq. (4.5) with the MF value of θ_e . The red dashed line is the theoretical power-law, while the pink, dash-dotted one corresponds to the linear regime mentioned in Eq. (4.8).

On the other hand, an excess of very small forces is easy to observe, similar to the one found for HS (and the other systems of Ref. [2]). To analyse this left tail more carefully, I will make use of the same methodology than from HS, namely, rescale $\text{cdf}(f_e)$ by a power of f_e in an attempt to find a constant function. The outcome of trying $1 + \theta_e$ (to verify the main scaling) and 1 (to test the hypothesis of a linear left tail) is presented in the top and bottom panels, respectively, of Fig. 4.13. Note that, once more, these two exponents describe the respective two regimes within an acceptable degree, but more data is needed to confirm the power-law of the left tail beyond doubts. I mention in passing that in this case there is no need to try θ_ℓ or another power because for MK packings the localized forces actually follow a *different* power-law; see Sec. 4.5.2.3.




Figure 4.13. Same distributions of Fig. 4.12 but rescaled by $f_e^{1+\theta_e}$ (upper panel) and $N^{\frac{\theta_e}{1+\theta_e}} f_e^{-1}$ (bottom one) to examine more carefully the different scalings present. Also in this case the left tail of $\text{cdf}(f_e)$ agrees reasonably well with a linear behaviour. Note that, in analogy to the analysis of the previous section, the size dependence of the left tail has also been accounted for, thus leading to a rather constant value of order 1.

4.5.2.2 Distribution of gaps

To continue, I now analyse the cdf of the interparticle gaps. The distributions obtained from the same configurations as above, together with the size scaling of Eq. (4.5) using the MF value of γ , are shown in the left and right panels, respectively, of Fig. 4.14. The analysis of these results is more interesting, given that in comparison with the analogous results in HS, here the curves do not really seem to follow the power-law behaviour predicted by MF. In fact, we can note that individual distributions of h suggest that a smaller exponent would better fit the curves in the left panel of Fig. 4.14. Yet, using the theoretical value of γ produces an excellent collapse as shown in the right panel of the same figure, and trying different values significantly worsens the quality of such collapse. This is, in fact, a typical situation of many critical scalings in finite N systems^[175,228]. In these cases, the most reliable way to determine true critical exponents is from the finite-size scaling analysis, as we have done here.

Before further analysing the prominent size effects observed in this model, I test once again the hypothesis of a linear behaviour for very small values of gaps. Applying the technique used so far, in Fig. 4.15 I present the cdf rescaled by $h^{-(1-\gamma)}$ (upper panel) and $N^{\frac{-\gamma}{1-\gamma}} h^{-1}$ (bottom one). Notice that in MK packings, the curves do not really have a constant value the region associated to the main power law,




Figure 4.14. Cdf of the set of interparticle gaps in MK configurations of different size (left) and the resulting scaling (right) following Eq. (4.5) with the MF value of γ . The red dashed line is the theoretical power-law, while the pink, dash-dotted one corresponds to the linear regime mentioned in Eq. (4.7). Additionally, the dotted, cyan line is the result of fitting Eq. (4.10) to the rescaled data of $N = 4096$. This fit assumes the MF value of γ , but weighs the two different regimes of (4.9) depending on the values of the scaling variable; see discussion in the main text. Interestingly, in this case, the curves do not seem to follow the predicted power-law due to strong finite size effects, but the very good collapsed obtained verifies that the exponent predicted by MF is the correct one. See text for details.

i.e. $h \in [10^{-5}, 10^{-2}]$, in contrast to the analogous results in HS and SS systems (cf. Fig. 4.7). This simply reproduces what I pointed out in the previous paragraph that a seemingly smaller value of γ (greater slope) would better fit the cdf, specially for the smallest systems. Nevertheless, it is remarkable that also in these configuration the linear behaviour of the smallest values of h remains unaltered, for at least two orders of magnitude. In fact, for this model the agreement with Eq. (4.7) of such secondary regime is better than in other cases. This is likely caused by the MF power-law being realized only partially, leading to a quick crossover to the extremal statistics behaviour. Note also that when the N dependence in the prefactor of the scaling function is included the plateau of all curves has a similar value.

I now come back to the marked size effects found in this model, evinced by the size dependence of $\text{cdf}(f_e)$ and the (apparently) different exponent of the gaps distribution. Both of these features contrast with the results of the previous section and the other models studied in Ref. [2]. For concreteness, I will restrict the analysis to the gaps cdf's, where finite size effects are most pronounced. We should first consider that, somehow surprisingly, the individual gaps distributions in the MK model, which by construction should closely resemble the MF value, do not display the right gap exponent. Indeed, from the right panel of Fig. 4.14 we observe that the scaling variable using the MF value of γ is the correct one (curves do collapse when plotted as a function of $\tilde{h} = N^{\frac{1}{1-\gamma}} h$), but the slope of the different curves in the




Figure 4.15. Same distributions of Fig. 4.14 but rescaled by $h^{1-\gamma}$ (upper panel) and $N^{\frac{\gamma}{1-\gamma}} h^{-1}$ (bottom one) to examine more carefully the different scalings present. Here, the near plateau formed after rescaling the cdf with a linear factor is much clearer, and by including the N dependence of the prefactor of Eq. (4.7) the different curves tend to roughly the same value as $h \rightarrow 0$.

range covered in our simulations ($10^{-3} < \tilde{h} < 10^3$) differs from the MF prediction. An important concern is thus whether this deviation is due to finite-size corrections or whether it indicates a failure of the MK model. In order to resolve the matter, we used the expected form of the scaling function, Eq. (4.9), to construct a fitting function, $F(\tilde{h})$, that *assumes* the correct behaviour of the scaling function for large values of \tilde{h} . Explicitly:

$$F(\tilde{h}) = \left[(a\tilde{h})^d + (b\tilde{h}^{1-\gamma})^d \right]^{1/d} \quad (4.10)$$

This fitting function hence only depends on three parameters and fulfils the condition that $F(\tilde{h}) \propto \tilde{h}$ for $\tilde{h} \ll 1$, while the MF form, $\tilde{h}^{1-\gamma}$, is recovered for large values of the scaling variable. Fitting $F(\tilde{h})$ to the largest system size results in the dotted, cyan line in the bottom panel of Fig. 4.14, which clearly interpolates nicely between both regimes. (The values of the parameters obtained from a least squares fit read: $a = 0.179008$, $b = 0.223543$, $d = -1.25569$.) Therefore, the hypothesis that gaps in larger MK systems would eventually follow the MF power-law is very plausible or, at least, cannot be rejected. The convergence of the slope of the scaling function to the predicted value is nevertheless extremely slow in comparison with both SS and HS (and also considering the cases studied in Ref. [2]). This implies that very large values of \tilde{h} are needed to measure the right slope. Quantitatively,

in Fig. 4.16 I report the difference of γ and our estimation of *local* γ_{MK} from the local slope estimated around $\tilde{h} = 1$ and $\tilde{h} = 100$ for the various system sizes. The value of γ_{MK} was determined using a maximum likelihood estimation^[227] (empty markers, dashed lines) and a least squares fit (filled markers, solid line). Notice that both methods yield very similar results, although maximum likelihood values are systematically smaller than the ones of least squares. This is likely due to the fact that by truncating the interval used for the fit, the simple⁵ maximum likelihood method is no longer exact and the resulting exponent is underestimated. In any case, around $\tilde{h} = 1$, the slope clearly differs from the MF prediction, but even when $\tilde{h} \sim 10^2$ very large system sizes are needed for it to approach the theoretical exponent. This deviation results in an *apparent* size dependence of the global exponent, *i.e.* $\gamma_{MK} = \gamma_{MK}(N)$, that is substantially more pronounced than for other models at similar N .




Figure 4.16. Difference of the gaps exponent in the MK model and its MF value. The values of γ_{MK} were determined using a least squares fit (solid lines, filled markers) or a maximum likelihood estimation (dashed lines, empty markers) at two different values of the scaling variable, $\tilde{h} = 1$ (orange) and $\tilde{h} = 100$ (blue). Note that even in this latter case the MF value is not fully recovered, indicating that larger systems, possibly by orders of magnitude, would be needed to observe the theoretical value of γ .

It is likely that this discrepancy is caused, at least in part, by the MK model being a fully connected one. In contrast with their sparse counterparts, fully connected models indeed require much larger system sizes for thermodynamic power-law scalings to be visible^[231–233]. In the MK configurations in particular this feature can be physically understood by recalling that the introduction of random shifts results in neighbours of a given particle (very likely) not being neighbours themselves. A par-

⁵There is also a generalization for truncated power-laws as explained, for instance, in [230]. However, assessing the goodness of fit is much more involved in this scenario and, for the case we are interested in, the results do not improve significantly.

ticle can thus have many more contacts than normally allowed in Euclidean space, as reported in Fig. 4.9. Data in that figure shows that it is not uncommon ($\sim 1\%$) for particles at jamming to have as many as 12 contacts (the $d = 3$ kissing number) or more. In general, particles are thus surrounded by many more particles—in actual and near contact—than usual hard spheres. Additionally, jamming densities in this model are much higher than in standard HS systems. Using the method described in Sec. 4.2—*i.e.* our MD+iLP algorithm with planted configurations^[110] to speed up the growing process—produces packings with densities $\phi_{J,MK} \gtrsim 3.1$ (cf. $\phi_{J,3d} \simeq 0.64$). Considering that $\phi \sim \sigma^{1/d}$, such values in the packing fraction imply that our MK configurations are made out of particles nearly twice as big as those of standard HS. When these two effects are combined, the outcome is that particles in MK packings are surrounded by an abundant cluster of relatively large neighbours. Therefore, with the effective size of the system being thus drastically reduced, the finite-size corrections are correspondingly more pronounced. We thus conclude that gaps in the MK model will probably follow the MF power-law scaling, as expected, but only at system sizes orders of magnitude larger than those considered here. In practice, the finite size effects are so important in the distribution of gaps in the MK model that its MF nature is, perhaps paradoxically, a strong limitation to study its MF behaviour in the thermodynamic limit. These same features are probably also responsible for the more pronounced size effects found in the distribution of extended contact forces.

4.5.2.3 Distribution of localized forces

Lastly, I consider the cdf of localized forces whose cdf is presented in Fig. 4.17, where no clear signatures of finite N effects are visible, although some dispersion can be observed in the extremal values. Putting together these results with the analogous ones in standard configurations of spheres (Sec. 4.4.2.3) and considering that the same absence of size effects were observed in the other two models of Ref. [2], we conclude that the expected behaviour is fully confirmed. Recall that this finding was anticipated previously in this chapter on the basis that the set $\{f_\ell\}$ corresponds to contact forces acting on bucklers, for which opening a contact mostly results in localized displacement field as derived in Refs. [150, 167] and Sec. 1.4.5. In other words, given that opening any of the contacts associated with a buckler only has a non-negligible effect over few particle layers away from its origin (see Fig. 1.20), it is reasonable to assume that their properties should be insensitive to N , or to any border or periodic effects, as we have here verified.

On the other hand, a surprising finding is that f_ℓ is distributed *uniformly*, *i.e.* $\theta_{\ell,MK} = 0$, at variance with previous studies reporting $\theta_\ell \approx 0.17$. A careful analysis suggests that this unexpected distribution is in tune with the spatial properties of MK packings. To explain why, first note that even though buckling forces follow a different pdf, selecting particles with z_ℓ contacts is still a valid selection criterion. (If their contribution had not been isolated, the remaining forces would not follow the MF power-law scaling given in Eq. (1.68b), as it does in Figs. 4.12-4.13, whereas if both kinds of forces are considered together, their joint pdf scales with an exponent ≈ 1.1 , which differs from the analogous quantity for conventional spheres^[150].) Second, analysing the distribution of dot products between contact vectors as in




Figure 4.17. Cumulative distribution function of $\{f_\ell\}$ in MK packings of different sizes (as indicated by the legend). Note that, as anticipated, no signature of finite size effects are present. However, an unexpected finding is that f_ℓ is distributed linearly (pink dash-dotted line), in contradiction the behaviour expected from previous studies (red, dashed line), *i.e.* with an exponent $1 + \theta_\ell$.

Sec. 4.5.1 reveals that particles with z_ℓ contacts in MK packings have a very similar distribution as those in standard hard sphere packings (cf. Fig. 4.11). Bucklers thus mainly give rise to a localized response thanks to them having three nearly coplanar contacts and one nearly orthogonal force. In order to understand why localized forces are uniformly distributed, I follow Ref. [167], which showed that the two types of contact forces are related to two different types of floppy modes: (i) extended forces are related to floppy modes that can couple strongly to external perturbations, and hence their response is bulk dominated; and (ii) buckling forces are associated to floppy modes of a rapidly decaying displacement field. (See also the discussion of Sec. 1.4.5.) With this in mind, it is important to consider that the value of $\theta_\ell \approx 0.17$ was estimated from the statistics of *displacements* in the latter type of contacts. There is therefore a strong connection between the distribution of forces in bucklers and the particle displacements their floppy modes produce. Now, let us assume that in an MK packing we open a buckling contact, $[ij]$, between particles i and j , in order to describe the associated displacement field. In particular, let us focus on the remaining contacts of, say, particle i . Because of the random shifts, the other particles touching i are (very likely) not constrained by each other nor by the other particles near i . Instead, the displacement of each neighbour of i is limited by its own contacts, which are not neighbours themselves, and are typically far apart. By the same token, the effect on the rest of particles in contact with j is determined by secondary contacts that –with high probability– are distant

from each other and from $[ij]$. As a result, opening a buckling contact produces a small series of uncorrelated displacements. No particular length scale is hence favoured over any other. Because of the close relation between localized forces and displacements just mentioned, it is natural for f_ℓ to be uniformly distributed.

At any rate, I would like to stress that our finding of a different value of θ_ℓ is important not merely as a matter of scrupulous curve fitting, but mainly because it violates the stability bound related to local excitations, Eq. (1.74b). Interestingly, a similar result was obtained for the FCC packings studied in [2], suggesting that more general criteria are needed for assessing the stability of such excitations in other classes of disordered models. I comment further about this point below.

4.6 Conclusions

The results I presented here systematically corroborate the non-trivial power-laws of the distributions of forces and gaps in two disordered systems at jamming, fully supporting the description derived from the exact MF theory. I emphasize that for the archetypical minimal model of monodisperse spheres, such validation included changing the direction of approach to the jamming point. Our results obtained with the MK model, a mandatory reference for comparison given its MF-like behaviour, similarly reinforces the $d \rightarrow \infty$ predictions.

Yet, our main finding is the contrasting system-size dependence of these distributions. In the two models studied here, size effects in $p(f_e)$ are practically non-existent, while $g(h)$ exhibits clear and systematic signatures of finite- N deviations from the expected power-law scaling. The same is true in the rest of the models considered in Ref. [2] and analysed with the same methodology. I emphasize that testing for such size scalings not only rigorously assesses the critical exponents^[175,228], but also provides key insight into the length scale of their correlations. Hence, it can be concluded that the MF exponents for all gaps and extended forces distributions are correct. Moreover, a second and more informative conclusion is that the distribution $p(f_e)$ reaches its thermodynamic limit behaviour at smaller values of N than $g(h)$. Two different correlation lengths, ξ_{f_e} and ξ_h , therefore characterise the relevant length scales of correlations of contact forces and gaps, respectively. For the former, it must hold that $N^{1/d} \gg \xi_{f_e}$, and thus finite size corrections are negligible, while the analogous condition for gaps reads $N^{1/d} \lesssim \xi_h$. In other words, gaps are correlated over significantly larger distances than forces, *i.e.* $\xi_h \gg \xi_{f_e}$. Such disparity in the correlation lengths is an unexpected consequence of our results, considering that both quantities are usually treated on an equal footing from the perspective of the SAT-UNSAT transition in the perceptron^[138,139], constraint satisfaction problems^[141,229], and neural networks^[148] as well as from the point of view of marginal stability in amorphous solids^[177,178]. MK results back this hypothesis if we consider that their very high densities and connectivity reduce the effective system size, as discussed in Sec. 4.5.2. Observing the scaling of Eq. (4.5) for the cdf of f_e is thus a manifestation of the smaller effective volume (for a similar N), which confirms that finite-size effects at jamming do occur for $p(f_e)$, but disappear for relatively small system sizes; see also the discussion leading to Fig. 4.5. The significantly more pronounced N dependence of the distributions of h (Figs. 4.12

and 4.14) thus supports our finding that $\xi_h \gg \xi_{f_e}$.

On the other hand, the linear growth of $\text{cdf}(f_\ell)$ in the MK model is at odds with the stability condition of Eq. (1.74b). This finding is more surprising because in these systems there is no long-range order, as there is in near-crystalline packings where similar violations have been reported^[2,162,163]. At the end of Sec. 4.5.2.3 I used the peculiar geometry of the MK packings to suggest a physical explanation for the uniform distribution of f_ℓ , but this reasoning does not explain why the stability condition between γ and θ_ℓ is apparently violated. Given the drastic difference in the inherent structures of the near-crystals and MK packings, they highlight the need to better understand how the spatial correlation caused by external perturbations are influenced by the type of disorder in jammed packings.

Finally, the most persistent observation was that all cumulative distributions, of both gaps and extended forces, behave in a seemingly linear fashion at very small arguments (specially when jamming is reached from the UC phase), in agreement with the MF predictions, p_0 and g_0 in Eqs. (4.7) and (4.8), respectively. We argued that such a cut-off in the main power-law is due to the extra contact with respect to isostatic configurations (see Sec. 4.3) and has been previously reported for the gaps distributions of disks packings^[165], but we are not aware of analogous findings in any other model or for the $\{f_e\}$ distributions. The universality of this secondary scaling has been previously predicted^[141] for all models that can be mapped to jamming of spherical particles, and it has been shown to occur even for non-spherical particles^[165], provided that their jammed states remain sufficiently close to isostaticity. Such robustness can be understood in part by considering that isostaticity is a global property of a system –where the amount of constraints matches the degrees of freedom– and not to the specific distributions of its microstructural variables. Because we have restricted our analysis to packings having 1SS, that is with exactly $N_c = N_{\text{dof}} + 1$ contacts, the ubiquity of the linear growth in the extremal part of the cdf's supports the hypothesis that left tails of the critical distributions is determined by the 1SS property alone, and not by the inherent structure. However, undersampling effects are very hard to avoid when dealing with extremal statistics, so a more stringent analysis must be carried out to verify that $p_0(x) \sim g_0(x) \sim 1$ when $x \ll 1$, specially for the former pdf. For completeness, I should mention that a previous work on the perceptron^[140] also reported a similar transition to a uniform distribution of contact forces that, unexpectedly, depended on the type of algorithm used to reach the jamming point. This is an important consideration because the MD+iLP and FIRE algorithms, used in the UC and OC phases, respectively, indeed yielded distributions with rather different left tails. Given the prominent role of the smallest forces and gaps in determining the global stability of the packings (Sec. 1.4.5), understanding their algorithmic dependence, if any, is undoubtedly an interesting topic for future research.

Appendix

4.A Equivalence of cumulative distributions

Here I briefly compare the cumulative distributions presented before in the chapter with the estimations using two other methods. Given M independent samples, the first of such methods consists in computing the cdf of each individual sample, and then compare such distributions across the M datasets. Because in general computing the cdf of a set of values involves ranking them, when different cdf's are compared, we can obtain an estimate of the typical value of the first, second, etc. element in such set. Moreover, the sample variance of the cdf thus estimated can also be provides useful information about how big are the fluctuations of the ranked elements.

The third way we tried also involves computing the cdf of each sample individually. But then, such function is evaluated at a fixed set of points, and this is repeated for all the systems. Clearly, the set of points at which all the cdf's are evaluated will not coincide in general with the sample values of any given system, but nonetheless the *empirical* cdf can be evaluated in all of them. The reason is that the empirical cdf of a given system is computed as a stepwise function, whose increments occur at the values sampled on such system, but is constant in the intervals formed by such values. Thus, the value of the cdf is well defined essentially in any point (compatible with its domain). In any case, the advantage of this procedure is that by considering the cumulative distributions of the different samples we can obtain the statistics of the probability mass assigned to a fixed set of points.




Figure 4.18. Cumulative distribution functions of extended (left panel) and localized (right) forces in HS computed using three different methods and for all the system sizes considered in this chapter. The method used in the main text (lines) is clearly compatible with (i) the cdf obtained by ranking the forces in each sample (filled diamonds) and (ii) with evaluating each sample's cdf on the same set of points (empty squares). Curves of different N have been displaced horizontally for clarity. Error bars correspond to the standard deviation in the M_N samples considered for each size; their different meaning in the two latter cases is explained in the text.

Naturally, with enough data available both methods should be compatible with each other, as well as with the estimation of the cdf used in the main text, *i.e.* using a single, large dataset obtained from all the samples. However, the two methods described in this appendix are useful to obtain uncertainties of the cdf and are complimentary in this aspect: the method based on each sample's ranking mainly informs on the variance of the random variable itself, while evaluating different cdf's in the same points provides the fluctuations of the probability mass.




Figure 4.19. Cumulative distribution functions of interparticle gaps in HS computed following three different procedures and for all the system sizes considered in this chapter. Once again, combining all the data in a single set as done in the main text (lines) produces compatible results with (i) the cdf obtained from each sample's ranking (filled diamonds) and (ii) with evaluating each sample's cdf on the same set of points (empty squares). Curves of different N have been displaced horizontally for clarity. Error bars correspond to the standard deviation in the M_N samples considered for each size. Note that the smallest gaps show much broader sample-to-sample fluctuations than the smallest forces.

In Figs. 4.18 and 4.19 I compare the cumulative distributions of forces and gaps, respectively, computed using each the three methods in all our HS systems. The agreement between them is excellent for all system sizes and the three of structural variables considered here. Additionally, the error bars of each data series confirms the interpretation given above for the corresponding method.

Bibliography

- ¹R. Díaz Hernández Rojas, G. Parisi, and F. Ricci-Tersenghi, “Inferring the particle-wise dynamics of amorphous solids from the local structure at the jamming point”, *Soft Matter* **17**, 1056–1083 (2021).
- ²P. Charbonneau, E. I. Corwin, R. C. Dennis, R. Díaz Hernández Rojas, H. Ikeda, G. Parisi, and F. Ricci-Tersenghi, “Finite-size effects in the microscopic critical properties of jammed configurations: A comprehensive study of the effects of different types of disorder”, *Physical Review E* **104**, 014102 (2021).
- ³C. Artiago, P. Baldan, and G. Parisi, “Exploratory study of the glassy landscape near jamming”, *Physical Review E* **101**, 052605 (2020).
- ⁴M. van Hecke, “Jamming of soft particles: geometry, mechanics, scaling and isotaticity”, *Journal of Physics: Condensed Matter* **22**, 033101 (2010).
- ⁵L. Berthier and G. Biroli, “Theoretical perspective on the glass transition and amorphous materials”, *Reviews of Modern Physics* **83**, 587–645 (2011).
- ⁶S. Torquato and F. H. Stillinger, “Jammed hard-particle packings: From Kepler to Bernal and beyond”, *Reviews of Modern Physics* **82**, 2633–2672 (2010).
- ⁷A. Donev, “Jammed packings of hard particles”, PhD thesis (Princeton University, 2006).
- ⁸P. W. Anderson, “Through the glass lightly”, *Science* **267**, 1615–1616 (1995).
- ⁹A. Cavagna, “Supercooled liquids for pedestrians”, *Physics Reports* **476**, 51–124 (2009).
- ¹⁰P. G. Debenedetti and F. H. Stillinger, “Supercooled liquids and the glass transition”, *Nature* **410**, 259–267 (2001).
- ¹¹G. Biroli and J. P. Garrahan, “Perspective: the glass transition”, *The Journal of Chemical Physics* **138**, 12A301 (2013).
- ¹²K. Binder and W. Kob, *Glassy Materials and Disordered Solids: An Introduction to Their Statistical Mechanics* (World Scientific, 2011), 562 pp.
- ¹³W. Götze, *Complex Dynamics of Glass-Forming Liquids: A Mode-Coupling Theory* (Oxford University Press, 2008), 655 pp.
- ¹⁴Springer, ed., *Slow relaxations and nonequilibrium dynamics in condensed matter*, Session LXXVII (Springer, Berlin, 2003).
- ¹⁵L. Leuzzi and T. M. Nieuwenhuizen, *Thermodynamics of the Glassy State* (CRC Press, Oct. 29, 2007), 365 pp.

- ¹⁶G. Parisi, P. Urbani, and F. Zamponi, *Theory of simple glasses: exact solutions in infinite dimensions* (Cambridge University Press, New York, 2019), 323 pp.
- ¹⁷L. Berthier, G. Biroli, J.-P. Bouchaud, L. Cipelletti, and W. van Saarloos, *Dynamical heterogeneities in glasses, colloids, and granular media*, Vol. 150 (Oxford University Press, 2011).
- ¹⁸L. M. C. Janssen, “Mode-Coupling Theory of the Glass Transition: A Primer”, [Frontiers in Physics](#) **6**, 10.3389/fphy.2018.00097 (2018).
- ¹⁹G. Biroli, “A crash course on ageing”, [Journal of Statistical Mechanics: Theory and Experiment](#) **2005**, P05014 (2005).
- ²⁰D. R. Reichman and P. Charbonneau, “Mode-coupling theory”, [Journal of Statistical Mechanics: Theory and Experiment](#) **2005**, P05013 (2005).
- ²¹F. Sciortino, “Potential energy landscape description of supercooled liquids and glasses”, [Journal of Statistical Mechanics: Theory and Experiment](#) **2005**, P05015 (2005).
- ²²T. Castellani and A. Cavagna, “Spin-glass theory for pedestrians”, [Journal of Statistical Mechanics: Theory and Experiment](#) **2005**, P05012 (2005).
- ²³L. Berthier and M. D. Ediger, “Facets of glass physics”, [Physics Today](#) **69**, 40 (2015).
- ²⁴C. A. Angell, P. H. Poole, and J. Shao, “Glass-forming liquids, anomalous liquids, and polyamorphism in liquids and biopolymers”, [Il Nuovo Cimento D](#) **16**, 993–1025 (1994).
- ²⁵L.-M. Martinez and C. A. Angell, “A thermodynamic connection to the fragility of glass-forming liquids”, [Nature](#) **410**, 663–667 (2001).
- ²⁶R. Richert and C. A. Angell, “Dynamics of glass-forming liquids. V. On the link between molecular dynamics and configurational entropy”, [The Journal of Chemical Physics](#) **108**, 9016–9026 (1998).
- ²⁷W. Kauzmann, “The Nature of the Glassy State and the Behavior of Liquids at Low Temperatures.”, [Chemical Reviews](#) **43**, 219–256 (1948).
- ²⁸L. Berthier, M. Ozawa, and C. Scalliet, “Configurational entropy of glass-forming liquids”, [The Journal of Chemical Physics](#) **150**, 160902 (2019).
- ²⁹W. Kob, “Supercooled liquids, the glass transition, and computer simulations”, in *Slow relaxations and nonequilibrium dynamics in condensed matter. lecture notes for Les Houches Session LXXVII*, edited by J.-L. Barrat, M. Feigelman, and J. Kurchan (Springer-Verlag, 2003).
- ³⁰L. Berthier and W. Kob, “The Monte Carlo dynamics of a binary Lennard-Jones glass-forming mixture”, [Journal of Physics: Condensed Matter](#) **19**, 205130 (2007).
- ³¹M. Goldstein, “Viscous Liquids and the Glass Transition: A Potential Energy Barrier Picture”, [The Journal of Chemical Physics](#) **51**, 3728–3739 (1969).
- ³²A. S. Keys, L. O. Hedges, J. P. Garrahan, S. C. Glotzer, and D. Chandler, “Excitations are localized and relaxation is hierarchical in glass-forming liquids”, [Physical Review X](#) **1**, 021013 (2011).

- ³³L. Berthier, G. Biroli, J.-P. Bouchaud, and R. L. Jack, “Overview of different characterisations of dynamic heterogeneity”, in *Dynamical heterogeneities in glasses, colloids, and granular media*, edited by L. Berthier, G. Biroli, J.-P. Bouchaud, L. Cipelletti, and W. van Saarloos (Oxford University Press, 2011).
- ³⁴T. R. Kirkpatrick and D. Thirumalai, “Comparison between dynamical theories and metastable states in regular and glassy mean-field spin models with underlying first-order-like phase transitions”, *Physical Review A* **37**, 4439–4448 (1988).
- ³⁵S. Franz and G. Parisi, “Recipes for Metastable States in Spin Glasses”, *Journal de Physique I* **5**, 1401–1415 (1995).
- ³⁶S. Franz and G. Parisi, “Phase Diagram of Coupled Glassy Systems: A Mean-Field Study”, *Physical Review Letters* **79**, 2486–2489 (1997).
- ³⁷S. Franz and G. Parisi, “On non-linear susceptibility in supercooled liquids”, *Journal of Physics: Condensed Matter* **12**, 6335–6342 (2000).
- ³⁸C. Donati, S. Franz, S. C. Glotzer, and G. Parisi, “Theory of non-linear susceptibility and correlation length in glasses and liquids”, *Journal of Non-Crystalline Solids* **307-310**, 215–224 (2002).
- ³⁹P. Harrowell, *The Length Scales of Dynamic Heterogeneity: Results from Molecular Dynamics Simulations*, to update, (Sept. 29, 2010) <http://arxiv.org/abs/1009.5886>.
- ⁴⁰C. Brito and M. Wyart, “Heterogeneous dynamics, marginal stability and soft modes in hard sphere glasses”, *Journal of Statistical Mechanics: Theory and Experiment* **2007**, L08003–L08003 (2007).
- ⁴¹R. Candelier, A. Widmer-Cooper, J. K. Kummerfeld, O. Dauchot, G. Biroli, P. Harrowell, and D. R. Reichman, “Spatiotemporal Hierarchy of Relaxation Events, Dynamical Heterogeneities, and Structural Reorganization in a Supercooled Liquid”, *Physical Review Letters* **105**, 135702 (2010).
- ⁴²L. Berthier and W. Kob, “Static point-to-set correlations in glass-forming liquids”, *Physical Review E* **85**, 011102 (2012).
- ⁴³G. M. Hocky, T. E. Markland, and D. R. Reichman, “Growing Point-to-Set Length Scale Correlates with Growing Relaxation Times in Model Supercooled Liquids”, *Physical Review Letters* **108**, 225506 (2012).
- ⁴⁴P. Charbonneau, E. Dyer, J. Lee, and S. Yaida, “Linking dynamical heterogeneity to static amorphous order”, *Journal of Statistical Mechanics: Theory and Experiment* **2016**, 074004 (2016).
- ⁴⁵S. Karmakar, C. Dasgupta, and S. Sastry, “Length scales in glass-forming liquids and related systems: a review”, *Reports on Progress in Physics* **79**, 016601 (2016).
- ⁴⁶J. Zylberg, E. Lerner, Y. Bar-Sinai, and E. Bouchbinder, “Local thermal energy as a structural indicator in glasses”, *Proceedings of the National Academy of Sciences* **114**, 7289–7294 (2017).
- ⁴⁷H. Tong and H. Tanaka, “Revealing Hidden Structural Order Controlling Both Fast and Slow Glassy Dynamics in Supercooled Liquids”, *Physical Review X* **8**, 011041 (2018).

- ⁴⁸H. Tong, H. Hu, P. Tan, N. Xu, and H. Tanaka, “Revealing Inherent Structural Characteristics of Jammed Particulate Packings”, *Physical Review Letters* **122**, 215502 (2019).
- ⁴⁹H. Tong and H. Tanaka, “Structural order as a genuine control parameter of dynamics in simple glass formers”, *Nature Communications* **10**, 1–10 (2019).
- ⁵⁰C. P. Royall and S. R. Williams, “The role of local structure in dynamical arrest”, *Physics Reports, The Role of Local Structure in Dynamical Arrest* **560**, 1–75 (2015).
- ⁵¹A. Ikeda, L. Berthier, and G. Biroli, “Dynamic criticality at the jamming transition”, *The Journal of Chemical Physics* **138**, 12A507 (2013).
- ⁵²S. Henkes, C. Brito, and O. Dauchot, “Extracting vibrational modes from fluctuations: a pedagogical discussion”, *Soft Matter* **8**, 6092–6109 (2012).
- ⁵³S. Franz, G. Parisi, P. Urbani, and F. Zamponi, “Universal spectrum of normal modes in low-temperature glasses”, *Proceedings of the National Academy of Sciences* **112**, 14539–14544 (2015).
- ⁵⁴J. Rottler, S. S. Schoenholz, and A. J. Liu, “Predicting plasticity with soft vibrational modes: From dislocations to glasses”, *Physical Review E* **89**, 042304 (2014).
- ⁵⁵L. E. Silbert, A. J. Liu, and S. R. Nagel, “Normal modes in model jammed systems in three dimensions”, *Physical Review E* **79**, 021308 (2009).
- ⁵⁶M. Wyart, L. E. Silbert, S. R. Nagel, and T. A. Witten, “Effects of compression on the vibrational modes of marginally jammed solids”, *Physical Review E* **72**, 051306 (2005).
- ⁵⁷M. Wyart, S. R. Nagel, and T. A. Witten, “Geometric origin of excess low-frequency vibrational modes in weakly connected amorphous solids”, *EPL (Europhysics Letters)* **72**, 486 (2005).
- ⁵⁸P. Charbonneau, E. I. Corwin, G. Parisi, A. Poncet, and F. Zamponi, “Universal Non-Debye Scaling in the Density of States of Amorphous Solids”, *Physical Review Letters* **117**, 045503 (2016).
- ⁵⁹F. Arceri and E. I. Corwin, “Vibrational Properties of Hard and Soft Spheres Are Unified at Jamming”, *Physical Review Letters* **124**, 238002 (2020).
- ⁶⁰Q. Wu, T. Bertrand, M. D. Shattuck, and C. S. O’Hern, “Response of jammed packings to thermal fluctuations”, *Physical Review E* **96**, 062902 (2017).
- ⁶¹S. Franz, T. Maimbourg, G. Parisi, and A. Scardicchio, “Impact of jamming criticality on low-temperature anomalies in structural glasses”, *Proceedings of the National Academy of Sciences* **116**, 13768–13773 (2019).
- ⁶²A. Widmer-Cooper, P. Harrowell, and H. Fynewever, “How Reproducible Are Dynamic Heterogeneities in a Supercooled Liquid?”, *Physical Review Letters* **93**, 135701 (2004).
- ⁶³A. Widmer-Cooper and P. Harrowell, “On the relationship between structure and dynamics in a supercooled liquid”, *Journal of Physics: Condensed Matter* **17**, S4025–S4034 (2005).

- ⁶⁴A. Widmer-Cooper and P. Harrowell, “On the study of collective dynamics in supercooled liquids through the statistics of the isoconfigurational ensemble”, *The Journal of Chemical Physics* **126**, 154503 (2007).
- ⁶⁵R. L. Jack, A. J. Dunleavy, and C. P. Royall, “Information-Theoretic Measurements of Coupling between Structure and Dynamics in Glass Formers”, *Physical Review Letters* **113**, 095703 (2014).
- ⁶⁶A. Widmer-Cooper, H. Perry, P. Harrowell, and D. R. Reichman, “Irreversible reorganization in a supercooled liquid originates from localized soft modes”, *Nature Physics* **4**, 711–715 (2008).
- ⁶⁷A. Widmer-Cooper and P. Harrowell, “Predicting the Long-Time Dynamic Heterogeneity in a Supercooled Liquid on the Basis of Short-Time Heterogeneities”, *Physical Review Letters* **96**, 185701 (2006).
- ⁶⁸L. Berthier and R. L. Jack, “Structure and dynamics of glass formers: Predictability at large length scales”, *Physical Review E* **76**, 041509 (2007).
- ⁶⁹G. M. Hocky, D. Coslovich, A. Ikeda, and D. R. Reichman, “Correlation of Local Order with Particle Mobility in Supercooled Liquids Is Highly System Dependent”, *Physical Review Letters* **113**, 157801 (2014).
- ⁷⁰E. D. Cubuk, S. S. Schoenholz, J. M. Rieser, B. D. Malone, J. Rottler, D. J. Durian, E. Kaxiras, and A. J. Liu, “Identifying Structural Flow Defects in Disordered Solids Using Machine-Learning Methods”, *Physical Review Letters* **114**, 108001 (2015).
- ⁷¹S. S. Schoenholz, “Combining Machine Learning and Physics to Understand Glassy Systems”, *Journal of Physics: Conference Series* **1036**, 012021 (2018).
- ⁷²E. D. Cubuk, S. S. Schoenholz, E. Kaxiras, and A. J. Liu, “Structural Properties of Defects in Glassy Liquids”, *The Journal of Physical Chemistry B* **120**, 6139–6146 (2016).
- ⁷³S. S. Schoenholz, E. D. Cubuk, E. Kaxiras, and A. J. Liu, “Relationship between local structure and relaxation in out-of-equilibrium glassy systems”, *Proceedings of the National Academy of Sciences* **114**, 263–267 (2017).
- ⁷⁴S. S. Schoenholz, E. D. Cubuk, D. M. Sussman, E. Kaxiras, and A. J. Liu, “A structural approach to relaxation in glassy liquids”, *Nature Physics* **12**, 469–471 (2016).
- ⁷⁵V. Bapst, T. Keck, A. Grabska-Barwińska, C. Donner, E. D. Cubuk, S. S. Schoenholz, A. Obika, A. W. R. Nelson, T. Back, D. Hassabis, and P. Kohli, “Unveiling the predictive power of static structure in glassy systems”, *Nature Physics* **16**, 448–454 (2020).
- ⁷⁶T. Castellani and A. Cavagna, “Spin-glass theory for pedestrians”, *Journal of Statistical Mechanics: Theory and Experiment* **2005**, P05012 (2005).
- ⁷⁷F. Zamponi, *Mean field theory of spin glasses*, (Aug. 28, 2010) <http://arxiv.org/abs/1008.4844> (visited on 02/12/2018).
- ⁷⁸M. Mezard, G. Parisi, and M. Virasoro, *Spin Glass Theory and Beyond: An Introduction to the Replica Method and Its Applications*, Vol. 9, World Scientific Lecture Notes in Physics (WORLD SCIENTIFIC, Nov. 1986).

- ⁷⁹C. P. Royall, F. Turci, S. Tatsumi, J. Russo, and J. Robinson, “The race to the bottom: approaching the ideal glass?”, *Journal of Physics: Condensed Matter* **30**, 363001 (2018).
- ⁸⁰F. Ritort and P. Sollich, “Glassy dynamics of kinetically constrained models”, *Advances in Physics* **52**, 219–342 (2003).
- ⁸¹J.-P. Hansen and I. R. McDonald, *Theory of Simple Liquids* (Elsevier, 2013).
- ⁸²A. Santos, S. B. Yuste, and M. López de Haro, “Structural and thermodynamic properties of hard-sphere fluids”, *The Journal of Chemical Physics* **153**, 120901 (2020).
- ⁸³T. Hales, “A proof of the Kepler conjecture”, *Annals of Mathematics* **162**, 1065–1185 (2005).
- ⁸⁴G. Parisi and F. Zamponi, “Mean-field theory of hard sphere glasses and jamming”, *Reviews of Modern Physics* **82**, 789–845 (2010).
- ⁸⁵D. Frenkel, “Order through entropy”, *Nature Materials* **14**, 9–12 (2015).
- ⁸⁶B. J. Alder and T. E. Wainwright, “Phase Transition for a Hard Sphere System”, *The Journal of Chemical Physics* **27**, 1208–1209 (1957).
- ⁸⁷W. W. Wood and J. D. Jacobson, “Preliminary Results from a Recalculation of the Monte Carlo Equation of State of Hard Spheres”, *The Journal of Chemical Physics* **27**, 1207–1208 (1957).
- ⁸⁸B. J. Alder and T. E. Wainwright, “Phase Transition in Elastic Disks”, *Physical Review* **127**, 359–361 (1962).
- ⁸⁹G. Battimelli and G. Ciccotti, “Berni Alder and the pioneering times of molecular simulation”, *The European Physical Journal H* **43**, 303–335 (2018).
- ⁹⁰M. Mareschal, “Early years of Computational Statistical Mechanics”, *The European Physical Journal H* **43**, 293–302 (2018).
- ⁹¹W. G. Hoover and C. G. Hoover, “From hard spheres and cubes to nonequilibrium maps with thirty-some years of thermostatted molecular dynamics”, *The Journal of Chemical Physics* **153**, 070901 (2020).
- ⁹²N. F. Carnahan and K. E. Starling, “Equation of State for Nonattracting Rigid Spheres”, *The Journal of Chemical Physics* **51**, 635–636 (1969).
- ⁹³M. Robles, M. López de Haro, and A. Santos, “Note: Equation of state and the freezing point in the hard-sphere model”, *The Journal of Chemical Physics* **140**, 136101 (2014).
- ⁹⁴M. Skoge, A. Donev, F. H. Stillinger, and S. Torquato, “Packing hyperspheres in high-dimensional Euclidean spaces”, *Physical Review E* **74**, 041127 (2006).
- ⁹⁵M. Isobe and W. Krauth, “Hard-sphere melting and crystallization with event-chain Monte Carlo”, *The Journal of Chemical Physics* **143**, 084509 (2015).
- ⁹⁶A. M. e. A. Malijevský J. Kolafa (auth.), *Theory and simulation of hard-sphere fluids and related systems*, 1st ed., Vol. 753, Lecture Notes in Physics (Springer-Verlag Berlin Heidelberg, 2008).

- ⁹⁷L. Berthier, D. Coslovich, A. Ninarello, and M. Ozawa, “Equilibrium Sampling of Hard Spheres up to the Jamming Density and Beyond”, *Physical Review Letters* **116**, 238002 (2016).
- ⁹⁸Z. W. Salsburg and W. W. Wood, “Equation of State of Classical Hard Spheres at High Density”, *The Journal of Chemical Physics* **37**, 798–804 (1962).
- ⁹⁹L. Berthier, P. Charbonneau, Y. Jin, G. Parisi, B. Seoane, and F. Zamponi, “Growing timescales and lengthscales characterizing vibrations of amorphous solids”, *Proceedings of the National Academy of Sciences* **113**, 8397–8401 (2016).
- ¹⁰⁰T. S. Grigera and G. Parisi, “Fast Monte Carlo algorithm for supercooled soft spheres”, *Physical Review E* **63**, 045102 (2001).
- ¹⁰¹A. Ninarello, L. Berthier, and D. Coslovich, “Models and Algorithms for the Next Generation of Glass Transition Studies”, *Physical Review X* **7**, 021039 (2017).
- ¹⁰²L. Berthier, G. Biroli, P. Charbonneau, E. I. Corwin, S. Franz, and F. Zamponi, “Gardner physics in amorphous solids and beyond”, *The Journal of Chemical Physics* **151**, 010901 (2019).
- ¹⁰³E. Gardner, “Spin glasses with p-spin interactions”, *Nuclear Physics B* **257**, 747–765 (1985).
- ¹⁰⁴P. Charbonneau, J. Kurchan, G. Parisi, P. Urbani, and F. Zamponi, “Glass and jamming transitions: from exact results to finite-dimensional descriptions”, *Annual Review of Condensed Matter Physics* **8**, 265–288 (2017).
- ¹⁰⁵Y. Jin, P. Urbani, F. Zamponi, and H. Yoshino, *A stability-reversibility map unifies elasticity, plasticity, yielding and jamming in hard sphere glasses*, (Mar. 12, 2018) <http://arxiv.org/abs/1803.04597> (visited on 10/09/2018).
- ¹⁰⁶P. Urbani and F. Zamponi, “Shear Yielding and Shear Jamming of Dense Hard Sphere Glasses”, *Physical Review Letters* **118**, 038001 (2017).
- ¹⁰⁷F. Krzakala and L. Zdeborová, “Potts glass on random graphs”, *EPL (Europhysics Letters)* **81**, 57005 (2008).
- ¹⁰⁸C. Scalliet, L. Berthier, and F. Zamponi, “Marginally stable phases in mean-field structural glasses”, *Physical Review E* **99**, 012107 (2019).
- ¹⁰⁹P. Charbonneau, Y. Jin, G. Parisi, and F. Zamponi, “Hopping and the Stokes–Einstein relation breakdown in simple glass formers”, *Proceedings of the National Academy of Sciences* **111**, 15025–15030 (2014).
- ¹¹⁰P. Charbonneau, Y. Jin, G. Parisi, C. Rainone, B. Seoane, and F. Zamponi, “Numerical detection of the Gardner transition in a mean-field glass former”, *Physical Review E* **92**, 012316 (2015).
- ¹¹¹P. Charbonneau, J. Kurchan, G. Parisi, P. Urbani, and F. Zamponi, “Fractal free energy landscapes in structural glasses”, *Nature Communications* **5**, 3725 (2014).
- ¹¹²J. Kurchan, G. Parisi, P. Urbani, and F. Zamponi, “Exact theory of dense amorphous hard spheres in high dimension. II. the high density regime and the gardner transition”, *The Journal of Physical Chemistry B* **117**, 12979–12994 (2013).
- ¹¹³C. L. Hicks, M. J. Wheatley, M. J. Godfrey, and M. A. Moore, “Gardner Transition in Physical Dimensions”, *Physical Review Letters* **120**, 225501 (2018).

- ¹¹⁴Y. Zhang, M. J. Godfrey, and M. A. Moore, “Marginally jammed states of hard disks in a one-dimensional channel”, *Physical Review E* **102**, 042614 (2020).
- ¹¹⁵C. Scalliet, L. Berthier, and F. Zamponi, “Absence of Marginal Stability in a Structural Glass”, *Physical Review Letters* **119**, 205501 (2017).
- ¹¹⁶R. C. Dennis and E. I. Corwin, “Jamming Energy Landscape is Hierarchical and Ultrametric”, *Physical Review Letters* **124**, 078002 (2020).
- ¹¹⁷J. Kurchan, G. Parisi, and F. Zamponi, “Exact theory of dense amorphous hard spheres in high dimension i. the free energy”, *Journal of Statistical Mechanics: Theory and Experiment* **2012**, P10012 (2012).
- ¹¹⁸P. Charbonneau, J. Kurchan, G. Parisi, P. Urbani, and F. Zamponi, “Exact theory of dense amorphous hard spheres in high dimension. III. the full replica symmetry breaking solution”, *Journal of Statistical Mechanics: Theory and Experiment* **2014**, P10009 (2014).
- ¹¹⁹A. Georges and J. S. Yedidia, “How to expand around mean-field theory using high-temperature expansions”, *Journal of Physics A: Mathematical and General* **24**, 2173–2192 (1991).
- ¹²⁰T. Maimbourg, J. Kurchan, and F. Zamponi, “Solution of the Dynamics of Liquids in the Large-Dimensional Limit”, *Physical Review Letters* **116**, 015902 (2016).
- ¹²¹S. F. Edwards and P. W. Anderson, “Theory of spin glasses”, *Journal of Physics F: Metal Physics* **5**, 965–974 (1975).
- ¹²²S. Singh, M. D. Ediger, and J. J. de Pablo, “Ultrastable glasses from in silico vapour deposition”, *Nature Materials* **12**, 139–144 (2013).
- ¹²³C. Rainone and P. Urbani, “Following the evolution of glassy states under external perturbations: the full replica symmetry breaking solution”, *Journal of Statistical Mechanics: Theory and Experiment* **2016**, 053302 (2016).
- ¹²⁴A. Montanari and F. Ricci-Tersenghi, “On the nature of the low-temperature phase in discontinuous mean-field spin glasses”, *The European Physical Journal B - Condensed Matter and Complex Systems* **33**, 339–346 (2003).
- ¹²⁵A. Montanari, G. Parisi, and F. Ricci-Tersenghi, “Instability of one-step replica-symmetry-broken phase in satisfiability problems”, *Journal of Physics A: Mathematical and General* **37**, 2073–2091 (2004).
- ¹²⁶G. Biroli and P. Urbani, “Breakdown of elasticity in amorphous solids”, *Nature Physics* **12**, 1130–1133 (2016).
- ¹²⁷G. Biroli and P. Urbani, “Liu-Nagel phase diagrams in infinite dimension”, *SciPost Physics* **4**, 020 (2018).
- ¹²⁸H. Cohn, *Packing, coding, and ground states*, (Mar. 16, 2016) <http://arxiv.org/abs/1603.05202> (visited on 12/04/2020).
- ¹²⁹J. H. Conway and N. J. A. Sloane, *Sphere Packings, Lattices and Groups* (Springer Science & Business Media, Mar. 9, 2013), 724 pp.
- ¹³⁰J. D. Bernal and J. Mason, “Packing of Spheres: Co-ordination of Randomly Packed Spheres”, *Nature* **188**, 910–911 (1960).

- ¹³¹G. D. Scott and D. M. Kilgour, “The density of random close packing of spheres”, [Journal of Physics D: Applied Physics](#) **2**, 863–866 (1969).
- ¹³²A. Donev, S. Torquato, F. H. Stillinger, and R. Connelly, “Comment on “Jamming at zero temperature and zero applied stress: The epitome of disorder””, [Physical Review E](#) **70**, 043301 (2004).
- ¹³³C. S. O’Hern, L. E. Silbert, A. J. Liu, and S. R. Nagel, “Reply to “Comment on ‘Jamming at zero temperature and zero applied stress: The epitome of disorder’””, [Physical Review E](#) **70**, 043302 (2004).
- ¹³⁴A. Baule, F. Morone, H. J. Herrmann, and H. A. Makse, “Edwards statistical mechanics for jammed granular matter”, [Reviews of Modern Physics](#) **90**, 015006 (2018).
- ¹³⁵A. J. Liu and S. R. Nagel, “Nonlinear dynamics: Jamming is not just cool any more”, [Nature](#) **396**, 21–22 (1998).
- ¹³⁶C. S. O’Hern, L. E. Silbert, A. J. Liu, and S. R. Nagel, “Jamming at zero temperature and zero applied stress: The epitome of disorder”, [Physical Review E](#) **68**, 011306 (2003).
- ¹³⁷A. J. Liu and S. R. Nagel, “The Jamming Transition and the Marginally Jammed Solid”, [Annual Review of Condensed Matter Physics](#) **1**, 347–369 (2010).
- ¹³⁸S. Franz and G. Parisi, “The simplest model of jamming”, [Journal of Physics A: Mathematical and Theoretical](#) **49**, 145001 (2016).
- ¹³⁹S. Franz, A. Sclocchi, and P. Urbani, “Critical jammed phase of the linear perceptron”, [Physical Review Letters](#) **123**, 115702 (2019).
- ¹⁴⁰Y. Kallus, “Scaling collapse at the jamming transition”, [Physical Review E](#) **93**, 012902 (2016).
- ¹⁴¹S. Franz, G. Parisi, M. Sevelev, P. Urbani, and F. Zamponi, “Universality of the SAT-UNSAT (jamming) threshold in non-convex continuous constraint satisfaction problems”, [SciPost Physics](#) **2**, 019 (2017).
- ¹⁴²F. Krzakala and J. Kurchan, “Landscape analysis of constraint satisfaction problems”, [Physical Review E](#) **76**, 021122 (2007).
- ¹⁴³R. Mari and J. Kurchan, “Dynamical transition of glasses: From exact to approximate”, [The Journal of Chemical Physics](#) **135**, 124504 (2011).
- ¹⁴⁴D. Hexner, P. Urbani, and F. Zamponi, “Can a Large Packing be Assembled from Smaller Ones?”, [Physical Review Letters](#) **123**, 068003 (2019).
- ¹⁴⁵R. Mari, F. Krzakala, and J. Kurchan, “Jamming versus Glass Transitions”, [Physical Review Letters](#) **103**, 025701 (2009).
- ¹⁴⁶S. Spigler, M. Geiger, S. d’Ascoli, L. Sagun, G. Biroli, and M. Wyart, “A jamming transition from under- to over-parametrization affects generalization in deep learning”, [Journal of Physics A: Mathematical and Theoretical](#) **52**, 474001 (2019).
- ¹⁴⁷M. Geiger, S. Spigler, S. d’Ascoli, L. Sagun, M. Baity-Jesi, G. Biroli, and M. Wyart, “Jamming transition as a paradigm to understand the loss landscape of deep neural networks”, [Physical Review E](#) **100**, 012115 (2019).

- ¹⁴⁸S. Franz, S. Hwang, and P. Urbani, “Jamming in Multilayer Supervised Learning Models”, *Physical Review Letters* **123**, 160602 (2019).
- ¹⁴⁹F. Antenucci, S. Franz, P. Urbani, and L. Zdeborová, “Glassy Nature of the Hard Phase in Inference Problems”, *Physical Review X* **9**, 011020 (2019).
- ¹⁵⁰P. Charbonneau, E. I. Corwin, G. Parisi, and F. Zamponi, “Jamming Criticality Revealed by Removing Localized Buckling Excitations”, *Physical Review Letters* **114**, 125504 (2015).
- ¹⁵¹E. Bitzek, P. Koskinen, F. Gähler, M. Moseler, and P. Gumbsch, “Structural Relaxation Made Simple”, *Physical Review Letters* **97**, 170201 (2006).
- ¹⁵²E. DeGiuli, E. Lerner, and M. Wyart, “Theory of the jamming transition at finite temperature”, *The Journal of Chemical Physics* **142**, 164503 (2015).
- ¹⁵³P. Charbonneau, E. I. Corwin, G. Parisi, and F. Zamponi, “Universal Microstructure and Mechanical Stability of Jammed Packings”, *Physical Review Letters* **109**, 205501 (2012).
- ¹⁵⁴C. P. Goodrich, A. J. Liu, and S. R. Nagel, “Finite-Size Scaling at the Jamming Transition”, *Physical Review Letters* **109**, 095704 (2012).
- ¹⁵⁵C. P. Goodrich, S. Dagois-Bohy, B. P. Tighe, M. van Hecke, A. J. Liu, and S. R. Nagel, “Jamming in finite systems: Stability, anisotropy, fluctuations, and scaling”, *Physical Review E* **90**, 022138 (2014).
- ¹⁵⁶C. P. Goodrich, A. J. Liu, and J. P. Sethna, “Scaling ansatz for the jamming transition”, *Proceedings of the National Academy of Sciences* **113**, 9745–9750 (2016).
- ¹⁵⁷C. P. Goodrich, W. G. Ellenbroek, and A. J. Liu, “Stability of jammed packings I: the rigidity length scale”, *Soft Matter* **9**, 10993–10999 (2013).
- ¹⁵⁸S. S. Schoenholz, C. P. Goodrich, O. Kogan, A. J. Liu, and S. R. Nagel, “Stability of jammed packings II: the transverse length scale”, *Soft Matter* **9**, 11000–11006 (2013).
- ¹⁵⁹D. Hexner, A. J. Liu, and S. R. Nagel, “Two Diverging Length Scales in the Structure of Jammed Packings”, *Physical Review Letters* **121**, 115501 (2018).
- ¹⁶⁰D. B. Liarte, X. Mao, O. Stenull, and T. C. Lubensky, “Jamming as a Multicritical Point”, *Physical Review Letters* **122**, 128006 (2019).
- ¹⁶¹H. Jacquin, L. Berthier, and F. Zamponi, “Microscopic Mean-Field Theory of the Jamming Transition”, *Physical Review Letters* **106**, 135702 (2011).
- ¹⁶²P. Charbonneau, E. I. Corwin, L. Fu, G. Tsekenis, and M. van der Naald, “Glassy, Gardner-like phenomenology in minimally polydisperse crystalline systems”, *Physical Review E* **99**, 020901 (2019).
- ¹⁶³G. Tsekenis, *Jamming Criticality of Near-Crystals*, (June 12, 2020) <http://arxiv.org/abs/2006.07373> (visited on 07/16/2020).
- ¹⁶⁴H. Ikeda, “Jamming and replica symmetry breaking of weakly disordered crystals”, *Physical Review Research* **2**, 033220 (2020).

- ¹⁶⁵H. Ikeda, C. Brito, and M. Wyart, “Infinitesimal asphericity changes the universality of the jamming transition”, [Journal of Statistical Mechanics: Theory and Experiment](#) **2020**, 033302 (2020).
- ¹⁶⁶C. Brito, H. Ikeda, P. Urbani, M. Wyart, and F. Zamponi, “Universality of jamming of nonspherical particles”, [Proceedings of the National Academy of Sciences](#) **115**, 11736–11741 (2018).
- ¹⁶⁷E. Lerner, G. Düring, and M. Wyart, “Low-energy non-linear excitations in sphere packings”, [Soft Matter](#) **9**, 8252–8263 (2013).
- ¹⁶⁸E. Lerner, E. DeGiuli, G. Düring, and M. Wyart, “Breakdown of continuum elasticity in amorphous solids”, [Soft Matter](#) **10**, 5085–5092 (2014).
- ¹⁶⁹E. DeGiuli, E. Lerner, C. Brito, and M. Wyart, “Force distribution affects vibrational properties in hard-sphere glasses”, [Proceedings of the National Academy of Sciences](#) **111**, 17054–17059 (2014).
- ¹⁷⁰S. Franz, A. Sclocchi, and P. Urbani, “Critical energy landscape of linear soft spheres”, [SciPost Physics](#) **9**, 012 (2020).
- ¹⁷¹V. A. Kuzkin, “On angular momentum balance for particle systems with periodic boundary conditions”, [ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift für Angewandte Mathematik und Mechanik](#) **95**, 1290–1295 (2015).
- ¹⁷²A. Donev, S. Torquato, and F. H. Stillinger, “Pair correlation function characteristics of nearly jammed disordered and ordered hard-sphere packings”, [Physical Review E](#) **71**, 011105 (2005).
- ¹⁷³A. B. Hopkins, F. H. Stillinger, and S. Torquato, “Disordered strictly jammed binary sphere packings attain an anomalously large range of densities”, [Physical Review E](#) **88**, 022205 (2013).
- ¹⁷⁴V. F. Hagh, E. I. Corwin, K. Stephenson, and M. F. Thorpe, “A broader view on jamming: from spring networks to circle packings”, [Soft Matter](#) **15**, 3076–3084 (2019).
- ¹⁷⁵M Newman and G Barkema, *Monte Carlo methods in statistical physics* (Oxford University Press: New York, USA, 1999).
- ¹⁷⁶N. Goldenfeld, *Lectures On Phase Transitions And The Renormalization Group* (CRC Press, Mar. 8, 2018), 250 pp.
- ¹⁷⁷M. Wyart, “Marginal Stability Constrains Force and Pair Distributions at Random Close Packing”, [Physical Review Letters](#) **109**, 125502 (2012).
- ¹⁷⁸M. Müller and M. Wyart, “Marginal Stability in Structural, Spin, and Electron Glasses”, [Annual Review of Condensed Matter Physics](#) **6**, 177–200 (2015).
- ¹⁷⁹L. E. Silbert, A. J. Liu, and S. R. Nagel, “Vibrations and Diverging Length Scales Near the Unjamming Transition”, [Physical Review Letters](#) **95**, 098301 (2005).
- ¹⁸⁰E. Lerner, G. Düring, and E. Bouchbinder, “Statistics and Properties of Low-Frequency Vibrational Modes in Structural Glasses”, [Physical Review Letters](#) **117**, 035501 (2016).
- ¹⁸¹H. Ikeda, “Universal non-mean-field scaling in the density of states of amorphous solids”, [Physical Review E](#) **99**, 050901 (2019).

- ¹⁸²H. Mizuno, H. Shiba, and A. Ikeda, “Continuum limit of the vibrational properties of amorphous solids”, [Proceedings of the National Academy of Sciences](#) **114**, E9767–E9774 (2017).
- ¹⁸³E. DeGiuli, A. Laversanne-Finot, G. Düring, E. Lerner, and M. Wyart, “Effects of coordination and pressure on sound attenuation, boson peak and elasticity in amorphous solids”, [Soft Matter](#) **10**, 5628–5644 (2014).
- ¹⁸⁴M. L. Manning and A. J. Liu, “Vibrational Modes Identify Soft Spots in a Sheared Disordered Packing”, [Physical Review Letters](#) **107**, 108302 (2011).
- ¹⁸⁵A. Donev, S. Torquato, F. H. Stillinger, and R. Connelly, “A linear programming algorithm to test for jamming in hard-sphere packings”, [Journal of Computational Physics](#) **197**, 139–166 (2004).
- ¹⁸⁶K. Krabbenhoft, A. V. Lyamin, J. Huang, and M. Vicente da Silva, “Granular contact dynamics using mathematical programming methods”, [Computers and Geotechnics](#) **43**, 165–176 (2012).
- ¹⁸⁷S. Torquato and Y. Jiao, “Robust algorithm to generate a diverse class of dense disordered and ordered sphere packings via linear programming”, [Physical Review E](#) **82**, 061302 (2010).
- ¹⁸⁸C. Artiaco, “Studio del landscape al jamming”, M.Sc. thesis (Sapienza Università di Roma, Rome, Italy, 2017).
- ¹⁸⁹Y. Jiao, F. H. Stillinger, and S. Torquato, “Nonuniversality of density and disorder in jammed sphere packings”, [Journal of Applied Physics](#) **109**, 013508 (2011).
- ¹⁹⁰S. P. Boyd and L. Vandenberghe, *Convex optimization* (Cambridge University Press, Cambridge, UK ; New York, 2004), 716 pp.
- ¹⁹¹R. J. Vanderbei, *Linear Programming: Foundations and Extensions*, Vol. 196, International Series in Operations Research & Management Science (Springer US, Boston, MA, 2014).
- ¹⁹²D. G. Luenberger and Y. Ye, *Linear and Nonlinear Programming*, Vol. 228, International Series in Operations Research & Management Science (Springer International Publishing, Cham, 2016).
- ¹⁹³J. Nocedal and S. J. Wright, *Numerical optimization*, 2nd ed, Springer Series in Operations Research (Springer, New York, 2006), 664 pp.
- ¹⁹⁴A. Donev, S. Torquato, and F. H. Stillinger, “Neighbor list collision-driven molecular dynamics simulation for nonspherical hard particles. I. Algorithmic details”, [Journal of Computational Physics](#) **202**, 737–764 (2005).
- ¹⁹⁵B. D. Lubachevsky, “How to simulate billiards and similar systems”, [Journal of Computational Physics](#) **94**, 255–283 (1991).
- ¹⁹⁶B. D. Lubachevsky and F. H. Stillinger, “Geometric properties of random disk packings”, [Journal of Statistical Physics](#) **60**, 561–583 (1990).
- ¹⁹⁷C. Rainone, P. Urbani, H. Yoshino, and F. Zamponi, “Following the Evolution of Hard Sphere Glasses in Infinite Dimensions under External Perturbations: Compression and Shear Strain”, [Physical Review Letters](#) **114**, 015701 (2015).

- ¹⁹⁸C. Rainone, *Metastable Glassy States Under External Perturbations : Monitoring the Effects of Compression and Shear-strain*, Springer Theses (Springer International Publishing, 2017).
- ¹⁹⁹Gurobi Optimization LLC, *Gurobi*, version 8.0, 2018.
- ²⁰⁰P. Rissone, “Forces correlation in jammed hard spheres”, M.Sc. thesis (Sapienza Università di Roma, Rome, Italy, 2017).
- ²⁰¹P. Rissone, E. I. Corwin, and G. Parisi, *Long-range anomalous decay of the correlation in jammed packings*, (Dec. 21, 2020) <http://arxiv.org/abs/2012.10181> (visited on 02/15/2021).
- ²⁰²P. Charbonneau and P. Morse, *Memory formation in jammed hard spheres*, (Sept. 1, 2020) <http://arxiv.org/abs/2009.00492> (visited on 11/11/2020).
- ²⁰³M. N. Bannerman, R. Sargant, and L. Lue, “DynamO: a free $\mathcal{O}(N)$ general event-driven molecular dynamics simulator”, *Journal of Computational Chemistry* **32**, 3329–3338 (2011).
- ²⁰⁴Q. Huangfu and J. A. J. Hall, “Parallelizing the dual revised simplex method”, *Mathematical Programming Computation* **10**, 119–142 (2018).
- ²⁰⁵H. G. E. Hentschel, I. Procaccia, and S. Roy, “Diffusion in agitated frictional granular matter near the jamming transition”, *Physical Review E* **100**, 042902 (2019).
- ²⁰⁶O. Dauchot, G. Marty, and G. Biroli, “Dynamical Heterogeneity Close to the Jamming Transition in a Sheared Granular Material”, *Physical Review Letters* **95**, 265701 (2005).
- ²⁰⁷F. Lechenault, O. Dauchot, G. Biroli, and J. P. Bouchaud, “Critical scaling and heterogeneous superdiffusion across the jamming/rigidity transition of a granular glass”, *EPL (Europhysics Letters)* **83**, 46003 (2008).
- ²⁰⁸C. Brito, O. Dauchot, G. Biroli, and J.-P. Bouchaud, “Elementary excitation modes in a granular glass above jamming”, *Soft Matter* **6**, 3013–3022 (2010).
- ²⁰⁹J. Deseigne, O. Dauchot, and H. Chaté, “Collective Motion of Vibrated Polar Disks”, *Physical Review Letters* **105**, 098001 (2010).
- ²¹⁰J. Deseigne, S. Léonard, O. Dauchot, and H. Chaté, “Vibrated polar disks: spontaneous motion, binary collisions, and collective dynamics”, *Soft Matter* **8**, 5629–5639 (2012).
- ²¹¹C. Coulais, R. P. Behringer, and O. Dauchot, “How the ideal jamming point illuminates the world of granular media”, *Soft Matter* **10**, 1519–1536 (2014).
- ²¹²A. Seguin and O. Dauchot, “Experimental Evidence of the Gardner Phase in a Granular Glass”, *Physical Review Letters* **117**, 228001 (2016).
- ²¹³A. Stukowski, “Visualization and analysis of atomistic simulation data with OVITO—the Open Visualization Tool”, *Modelling and Simulation in Materials Science and Engineering* **18**, 015012 (2009).
- ²¹⁴S. Patinet, D. Vandembroucq, and M. L. Falk, “Connecting Local Yield Stresses with Plastic Activity in Amorphous Solids”, *Physical Review Letters* **117**, 045501 (2016).

- ²¹⁵B. Xu, M. L. Falk, J. F. Li, and L. T. Kong, “Predicting Shear Transformation Events in Metallic Glasses”, [Physical Review Letters](#) **120**, 125503 (2018).
- ²¹⁶A. Barbot, M. Lerbinger, A. Hernandez-Garcia, R. García-García, M. L. Falk, D. Vandembroucq, and S. Patinet, “Local yield stress statistics in model amorphous solids”, [Physical Review E](#) **97**, 033001 (2018).
- ²¹⁷Z. Schwartzman-Nowik, E. Lerner, and E. Bouchbinder, “Anisotropic structural predictor in glassy materials”, [Physical Review E](#) **99**, 060601 (2019).
- ²¹⁸C. Coulais, R. P. Behringer, and O. Dauchot, “Dynamics of the contacts reveals Widom lines for jamming”, [EPL \(Europhysics Letters\)](#) **100**, 44005 (2012).
- ²¹⁹X. Ma, Z. S. Davidson, T. Still, R. J. S. Ivancic, S. S. Schoenholz, A. J. Liu, and A. G. Yodh, “Heterogeneous Activation, Local Structure, and Softness in Supercooled Colloidal Liquids”, [Physical Review Letters](#) **122**, 028001 (2019).
- ²²⁰F. Lechenault, R. Candelier, O. Dauchot, J.-P. Bouchaud, and G. Biroli, “Superdiffusion around the rigidity transition: Lévy and the Lilliputians”, [Soft Matter](#) **6**, 3059–3064 (2010).
- ²²¹J. Wang and A. P. Young, “Monte Carlo study of the six-dimensional Ising spin glass”, [Journal of Physics A: Mathematical and General](#) **26**, 1063 (1993).
- ²²²J. J. Ruiz-Lorenzo, “Logarithmic corrections for spin glasses, percolation and Lee-Yang singularities in six dimensions”, [Journal of Physics A: Mathematical and General](#) **31**, 8773–8787 (1998).
- ²²³R. Kenna, “Finite size scaling for $O(N)$ Φ^4 -theory at the upper critical dimension”, [Nuclear Physics B](#) **691**, 292–304 (2004).
- ²²⁴H. Ikeda, “Jamming Below Upper Critical Dimension”, [Physical Review Letters](#) **125**, 038001 (2020).
- ²²⁵P. K. Morse and E. I. Corwin, “Geometric signatures of jamming in the mechanical vacuum”, [Physical Review Letters](#) **112**, 115701 (2014).
- ²²⁶P. K. Morse and E. I. Corwin, “Echoes of the Glass Transition in Athermal Soft Spheres”, [Physical Review Letters](#) **119**, 118003 (2017).
- ²²⁷M. E. J. Newman, “Power laws, Pareto distributions and Zipf’s law”, [Contemporary Physics](#) **46**, 323–351 (2005).
- ²²⁸D. J. Amit and V. Martin-Mayor, *Field Theory, the Renormalization Group, and Critical Phenomena: Graphs to Computers*, 3rd ed. (WORLD SCIENTIFIC, June 2005).
- ²²⁹H. Ikeda, P. Urbani, and F. Zamponi, “Mean field theory of jamming of nonspherical particles”, [Journal of Physics A: Mathematical and Theoretical](#) **52**, 344001 (2019).
- ²³⁰A. Deluca and Á. Corral, “Fitting and goodness-of-fit test of non-truncated and truncated power-law distributions”, [Acta Geophysica](#) **61**, 1351–1394 (2013).
- ²³¹C. Lucibello, F. Morone, G. Parisi, F. Ricci-Tersenghi, and T. Rizzo, “Anomalous finite size corrections in random field models”, [Journal of Statistical Mechanics: Theory and Experiment](#) **2014**, P10025 (2014).

- ²³²C. Lucibello, F. Morone, G. Parisi, F. Ricci-Tersenghi, and T. Rizzo, “Finite-size corrections to disordered Ising models on random regular graphs”, [Physical Review E](#) **90**, 012146 (2014).
- ²³³U. Ferrari, C. Lucibello, F. Morone, G. Parisi, F. Ricci-Tersenghi, and T. Rizzo, “Finite-size corrections to disordered systems on Erdős-Rényi random graphs”, [Physical Review B](#) **88**, 184201 (2013).

List of Figures

1.1	Some examples of systems that exhibit a jammed regime.	2
1.2	Glass – supercooled liquid diagram.	4
1.3	Angell plots obtained experimentally and from simulations.	6
1.4	Structure factor in a LJ mixture at four different temperatures. . . .	7
1.5	Signatures of two-steps relaxations in the MSD and the self-intermediate scattering function.	8
1.6	Sketch of the potential energy landscape of a glass former.	9
1.7	Prediction of typical behaviour of correlation functions according to MCT.	11
1.8	Dynamical heterogeneity in glassy systems.	12
1.9	Criticality of the four point susceptibility.	13
1.10	Different methods for connecting the local structure with dynamical features or plastic rearrangements in a variety of models.	15
1.11	Dynamics-structure correlation in the isoconfigurational ensemble. .	16
1.12	Machine Learning methods attain a high predictability of dynamical features, but use artificial representations of the particles' local environment.	17
1.13	Phase diagram of hard sphere fluids: liquid and solid phase	19
1.14	Liquid equation of state in HS systems, for different compression protocols	27
1.15	Liquid and glass equations of state in HS systems, comparing theory with numerical simulations.	29
1.16	Phase diagram of $d = 3$ HS polydisperse mixture.	31
1.17	Phase diagram of HS systems in the $d \rightarrow \infty$ limit.	40
1.18	MF picture of the fragmentation of the phase space at different densities.	41
1.19	Liu–Nagel phase diagram	45
1.20	Floppy modes of localized and extended excitations.	58
2.1	Exact and linearised non-overlapping constraints of the optimization problem of HS jamming.	66
2.2	Optimal solution and new constraints after the first LP optimization of the linear problem in Fig. 2.1b.	74
2.3	Jammed configuration and network of contact forces of $N = 1024$ bidisperse disks and periodic boundary conditions.	77
2.4	Diagram to determine the collision time of HS in MD simulations. .	78

2.5	Comparison of event-driven MD simulations and the exact liquid EOS in the MK model	81
2.6	Liquid and glass branches in the MK model, both with and without planting.	83
2.7	Probability distribution function of interparticle gaps from MD simulations after reaching different values of $p \gg 1$	85
2.8	Jamming packing fraction as a function of p_{tar} , κ , and N	88
2.9	Similarity in the network of contacts as a function of p_{tar} , κ , and N	89
2.10	Radial distribution function of jammed packings obtained from different p_{tar}	91
2.11	Difference in jamming densities obtained from iLP, φ_J , and fitting data of MD simulations, $\varphi_J^{(MD)}$	93
2.12	Diagram of the MD+iLP route to jamming	95
2.13	Histograms of displacements during iLP crunching	96
2.14	MSD during at different steps of iLP	97
2.15	Execution times of the MD+iLP algorithm, for different values of κ , p_{tar} , and N	100
2.16	Number of LP optimizations required for convergence as a function of p_{tar} and κ	102
2.17	Iterations and convergence time as a function of the system's size	103
2.18	Time per LP iteration and its dependence on N	104
3.1	Jammed configuration of $N = 1024$ spheres used to perform dynamical simulations.	111
3.2	Cumulative distributions of contact forces, gaps, and frequencies.	113
3.3	Distributions of dot products of contact vectors and contact forces.	115
3.4	Relation of the physical time and collision events in MD simulations.	119
3.5	Temporal evolution of MSD for different packing fractions.	120
3.6	Probability distributions of single particle displacements and mobilities.	121
3.7	Comparison of mean and standard deviation of the single-particle displacements.	122
3.8	Scatter plots of $\langle \delta \mathbf{r}_i \rangle$ (from MD simulations) and \mathbf{C}_i	124
3.9	Scatter plot of $Var[\delta \mathbf{r}_i]$ (from MD simulations) and the absolute value of \mathbf{C}_i	126
3.10	Scatter plots of the particles' average mobility (from MD simulations) and three scalar quantities.	127
3.11	Average energy per particle as a function of the number of MC steps using three different potentials.	128
3.12	Scatter plots of $\langle \delta \mathbf{r}_i \rangle$ (from MC dynamics) and $\vec{\mathbf{C}}_i$	129
3.13	Scatter plot of neighbouring particles $Var[\delta \mathbf{r}_i]$ (from MC simulations) and the absolute value of \mathbf{C}_i	130
3.14	Scatter plots of the particles mobility (from MC dynamics) and three scalar quantities.	131
3.15	Predicting preferential directions and mobility from the ranking of structural variables.	135

3.16	Evolution of the Spearman correlation between preferential directions and total contact vector and between mobility and the sum of dot products of contact vectors. Results from MD simulations	136
3.17	Evolution of the Spearman correlation between preferential directions and total contact vector and between mobility and the sum of dot products of contact vectors. Results from MC dynamics.	138
3.18	Projection of the single particle mean displacement along each normal mode.	142
3.19	Projection of the single particle mean displacement along each normal mode, as a function of the participation ratio.	144
3.20	Curvature induced by the Hessian and the comparison with the density of states.	145
3.21	Analogous of Fig. 3.8 but for a second, independent configuration. .	149
3.22	Analogous of Fig. 3.10 but for a second, independent configuration. .	150
3.23	Comparison of the temporal evolution of the statistical correlations between structure and dynamics in two independent configurations. .	150
4.1	Probability density function of the coordination number for different system sizes.	160
4.2	Radial distribution function of HS jammed packings.	160
4.3	Size scaling of the cdf of $\{f_e\}$ in HS and SS configurations.	162
4.4	Rescaled cdf of $\{f_e\}$ in spheres packings to study the two different power-laws.	163
4.5	Cdf of f_e in spheres packings with contribution of extremes removed. .	164
4.6	Size scaling of the cdf of $\{h\}$ in HS and SS configurations.	165
4.7	Rescaled cdf of $\{h\}$ in SS and HS configurations to study the two different power-laws.	166
4.8	Cumulative distribution function of localized forces in spheres packings	168
4.9	pdf of the coordination number in MK packings.	169
4.10	pdf of dot products between contact vectors in MK configurations .	170
4.11	pdf of S_i of bucklers in MK and HS of different size	170
4.12	Size scaling of the cdf of $\{f_e\}$ in MK configurations	171
4.13	Rescaled cdf of $\{f_e\}$ in MK configurations to study the two different power-laws.	172
4.14	Size scaling of the cdf of $\{h\}$ in MK configurations.	173
4.15	Rescaled cdf of $\{h\}$ in MK configurations to study the two different power-laws.	174
4.16	Difference of the gaps exponent in the MK model and its MF value .	175
4.17	Cdf of localized forces in jammed MK configurations	177
4.18	Three different (and compatible) ways of computing the cdf of forces in HS.	180
4.19	Three different (and compatible) ways of computing the cdf of gaps in HS.	181