

Hard-sphere jamming through the lens of linear optimizationClaudia Artiago¹, Rafael Díaz Hernández Rojas^{2,*}, Giorgio Parisi^{2,3} and Federico Ricci-Tersenghi^{2,3}¹*Department of Physics, KTH Royal Institute of Technology, Stockholm 106 91, Sweden*²*Dipartimento di Fisica, Sapienza Università di Roma, Piazzale Aldo Moro 5, 00185 Rome, Italy*³*INFN, Sezione di Roma1, and CNR-Nanotec, unità di Roma, Piazzale Aldo Moro 5, 00185 Rome, Italy*

(Received 7 April 2022; accepted 21 September 2022; published 15 November 2022)

The jamming transition is ubiquitous. It is present in granular matter, foams, colloids, structural glasses, and many other systems. Yet, it defines a critical point whose properties still need to be fully understood. Recently, a major breakthrough came about when the replica formalism was extended to build a mean-field theory that provides an exact description of the jamming transition of spherical particles in the infinite-dimensional limit. While such theory explains the jamming critical behavior of both soft and hard spheres, investigating the transition in finite-dimensional systems poses very difficult and different problems, in particular from the numerical point of view. Soft particles are modeled by continuous potentials; thus, their jamming point can be reached through efficient energy minimization algorithms. In contrast, the latter methods are inapplicable to hard-sphere (HS) systems since the interaction energy among the particles is always zero by construction. To overcome these difficulties, here we recast the jamming of hard spheres as a constrained optimization problem and introduce the CALiPPSO algorithm, capable of readily producing jammed HS packings without including any effective potential. This algorithm brings a HS configuration of arbitrary dimensions to its jamming point by solving a chain of linear optimization problems. We show that there is a strict correspondence between the force balance conditions of jammed packings and the properties of the optimal solutions of CALiPPSO, whence we prove analytically that our packings are always isostatic and in mechanical equilibrium. Furthermore, using extensive numerical simulations, we show that our algorithm is able to probe the complex structure of the free-energy landscape, finding qualitative agreement with mean-field predictions. We also characterize the algorithmic complexity of CALiPPSO and provide an open-source implementation of it.

DOI: [10.1103/PhysRevE.106.055310](https://doi.org/10.1103/PhysRevE.106.055310)**I. INTRODUCTION**

Jamming is a pervasive phenomenon: it is present in systems with diverse time and length scales, such as structural glasses, grains, emulsions, foams, and colloids [1–4]. Such ubiquity has been partially understood by recognizing that the jamming point defines a critical point common to all these systems [1,2,5]. Despite being an out-of-equilibrium transition that brings a system to form a mechanically rigid packing, disordered jammed states can be identified as minima of a (properly defined) free-energy landscape (FEL) [6]. In the case of hard-sphere (HS) systems, which are a minimal model for athermal and granular matter [7,8], jamming is reached at infinite pressure, $p = \infty$, and the jammed packings are identified by their packing fraction ϕ_J . In soft-sphere (SS) systems, the jamming point is still identified by ϕ_J , but in the limit $p \rightarrow 0$, at least in the zero temperature limit [5,9].

Despite its physical relevance, a comprehensive theory of jamming is still far from being formulated. Recently, a mean-field theory has provided an exact description of the jamming transition in infinite-dimensional sphere systems [6,10–13]. Such theory established that jamming occurs within the so-called Gardner phase [11–14], where all states become

marginally stable. This feature implies an abundance of soft modes in jammed or nearly jammed packings. Another consequence is that critical jammed packings are isostatic, meaning that the number of mechanical constraints (i.e., contacts between particles) precisely matches the number of degrees of freedom. Additionally, the mean-field theory predicted that near jamming, the FEL of HS configurations is a very rough and hierarchically organized hypersurface.

Another important theoretical step came about with the realization that the jamming transition of hard spheres in the high dimensional limit can be thought of as the satisfiability/unsatisfiability threshold of continuous constraint satisfaction problems (CSPs), where the constraints are induced by the requirement that spheres do not overlap. From this point of view, jamming criticality defines a universality class encompassing the physical systems mentioned above, as well as CSPs [15,16], neural networks [17,18], and inference problems [19]. In particular, in recent years, the perceptron model has gained a prominent role among the CSPs; it has been employed to investigate special instances of the sphere packing problem [20,21], and even to analyze their quantum regime [22,23].

From this perspective, it is quite remarkable that the same jamming criticality predicted by mean-field theory (i.e., as $d \rightarrow \infty$) has been observed in finite dimensional systems, even down to $d = 2$ [6,24–29] which is now believed to be the

*rafael.diazhernandezrojas@uniroma1.it

upper critical dimension. Explaining the observed robustness of the jamming phenomenology in HS systems in different dimensions constitutes a challenging open problem, which has been tackled both numerically [5,6,24–36] and experimentally [37–43].

A related question is whether finite-dimensional HS and SS systems exhibit the same critical scalings when approaching the jamming transition, i.e., as $\varphi \rightarrow \varphi_J^-$ and $\varphi \rightarrow \varphi_J^+$, respectively. While it is generally believed that HS and SS systems share similar critical behaviors, it has been observed that some critical exponents are markedly different. For instance, near jamming the pressure in hard spheres and harmonic soft spheres scales as $p \sim |\varphi - \varphi_J|^{\pm 1}$ for $\varphi \rightarrow \varphi_J^\pm$; nonetheless, both critical behaviors can be captured by a single scaling function [13]. Importantly, it has also been shown that the distributions of forces and interparticle gaps at the critical point are insensitive to the direction from which jamming is reached. The latter feature has been verified numerically only recently [29], thanks to the algorithm for jamming in hard spheres presented in this article.

In summary, while HS and SS systems allow a similar theoretical treatment, from the numerical point of view they represent very different types of systems. So far, computer simulations have amply favored SS systems thanks to the efficiency of energy minimization techniques available for interacting systems. In particular, the powerful FIRE algorithm [44] has been successfully employed for studying jammed packings made of several thousands of soft particles, and there is now abundant numerical evidence that the jamming criticality of SS systems agrees with the mean-field predictions [24,29,35,45–48].

In contrast, due to the singular interaction potential of HS systems, studies analyzing their critical behavior at jamming are much more scarce. By definition, the interaction energy in HS configurations is either zero (if spheres do not overlap) or infinite (whenever two or more spheres overlap). This makes energy minimization strategies inapplicable. To partially overcome this problem, *nearly* jammed HS configurations have been produced using the Lubachevsky–Stillinger (LS) compression protocol [49,50]. By properly tuning the compression rate, this protocol can avoid crystallization (whenever $d \geq 3$) and produce highly compressed HS glasses [51,52]. However, since this algorithm relies on simulating the dynamics of HS configurations through elastic collisions, particles do not remain in contact, nor does limit $p = \infty$ is reached. To solve this problem, other molecular dynamics (MD) algorithms have been devised. For instance, the authors of Ref. [53] considered an overdamped dynamics, in which once two particles collide they are constrained to stay in contact. By cleverly incorporating geometrical information of the configuration into the dynamical equations, they keep a well defined contact network at all times and are able to generate jammed HS packings, coming from this particular dynamical scenario. However, MD approaches that purely rely on elastic collisions—the most common case—lack in general the required precision to resolve the full network of contacts that determines a jammed state (as we exemplify in Appendix D for the LS protocol).

Alternatively, jammed packings of pseudohard spheres have been constructed by introducing an effective interaction,

and subsequently minimizing the corresponding potential energy. The most common choice has been a logarithmic potential, $v(h) \propto -\log h$, where h is the (dimensionless) distance between spheres [27,35,48,54–57]. This potential clearly accounts for the nonoverlapping condition, and ample evidence exists that it correctly accounts for the interactions between particles in terms of their *average* position. Nevertheless, it is also important to investigate how to describe HS systems in terms of their *instantaneous* positions. Modeling HS systems by adding an extraneous potential presents two major drawbacks. First, any effective potential introduces fictitious interactions for $h > 0$, which become more important as the system approaches its jamming point. Second, the associated energy-minimization algorithms identify jammed packings as equilibrium states at zero temperature, while in true HS jammed states correspond to entropy extrema [58]. To certify that the logarithmic, or any other potential capture the physics of HS systems, one should carefully study such systems close to their jamming point. The algorithm we propose here is an ideal candidate for such a task: without relying on any effective potentials, it is able to reach the jamming point making use only of the geometrical constraints of HS configurations and on the particle’s instantaneous positions. This approach relies on a direct mapping of jamming of hard spheres into a constrained optimization problem.

The purpose of this article is to present in full detail the *CALiPPSO algorithm*, that is able to generate jammed packings of hard spheres in arbitrary dimensions and polydispersity without introducing any interaction potential between particles, thus overcoming the issues mentioned above. It relies on the linear approximation of the original, nonconvex optimization problem, and it gradually reaches jamming by iterating over a series of linear programming instances. Hence its name, which stands for *Chain of Approximate Linear Programming for Packing Spherical Objects*. In addition, in this article we show that combining CALiPPSO with the LS compression protocol provides a powerful tool for exploring the physics of jamming in finite-dimensional HS systems. We provide our own implementation as a Julia [59] package at [60].

The CALiPPSO algorithm has been successfully employed before for studying the properties of $3d$ packings near the jamming point [32,61]. Moreover, in combination with the LS protocol, it was used to confirm the jamming criticality of spheres and other mean-field-like models [29]. We refer to those references for extensive discussions on these issues.

The plan of the present article is as follows. In Sec. II, we introduce the CALiPPSO algorithm, and analytically show that the configurations that it produces are well-defined jammed states. Specifically: (i) CALiPPSO jammed packings are solutions for the exact optimization problem; (ii) such solutions satisfy the mechanical equilibrium condition, previously derived for soft spheres [24,27]; (iii) they are also always isostatic. Importantly, points (ii) and (iii) imply that the Hessian obtained for hard spheres matches that of soft spheres which has been found to reproduce the marginal stability condition expected from the mean-field solution. In Sec. III, we demonstrate, through extensive numerical simulations, that CALiPPSO can be readily coupled with the LS compression protocol to produce a robust and fast algorithm for jamming.

We characterize the behavior of the two algorithms combined, and we provide evidence that the LS+CALiPPSO route to jamming allows us to study the hierarchical structure of the FEL of hard spheres with very high accuracy, reproducing the results previously obtained with other methods. In Sec. IV, we investigate the time complexity of our algorithm and conclude that its running time scales with the cube of the system size. Finally, in Sec. V, we briefly summarize our results and discuss why CALiPPSO should be preferred to other jamming algorithms when an accurate identification of the network of contacts is needed. We finish by providing some examples of other problems that could be tackled using CALiPPSO or similar methods.

Before ending the Introduction, we acknowledge that other works, Refs. [62–68], have previously used linear programming methods to produce jammed packings of hard spheres. Our optimization problem is simpler than the ones considered in those studies since it contains fewer optimization variables; thanks to this, we are able to obtain an exact proof of the properties (i)–(iii) mentioned above. To the best of our knowledge, none of the previous works has carried out a similar analysis. Furthermore, the question of whether linear programming algorithms allow probing the FEL of hard spheres has never been put to test, as we do here. Additionally, CALiPPSO, if initialized with sufficiently highly compressed configurations, produces packings that should coincide with the $p \rightarrow \infty$ extrapolation of LS. In contrast, this is not always the case within the other methods [62]. However, it should be noted that some of the previous works [62,66–68] demonstrated that their algorithms are capable of producing both ordered and disordered packings, within a rather broad range of densities; this feature is not realized by CALiPPSO.

II. JAMMED PACKINGS OBTAINED USING CALiPPSO

A. Jamming as an optimization problem

Let us consider N spheres of diameters $\bar{\sigma} = \{\sigma_i\}_{i=1}^N$ inside a cubic box of volume $V = L^d$. We denote the dN -dimensional vector of their centers as $\bar{\mathbf{r}} = \{\mathbf{r}_i\}_{i=1}^N$, where \mathbf{r}_i is the d -dimensional vector identifying the position of the i th particle. The system's packing fraction is given by $\varphi = \frac{1}{V} \sum_{i=1}^N v_d(\sigma_i/2)$, where $v_d(R) = \frac{\pi^{d/2} R^d}{\Gamma(1+d/2)}$ is the volume of a hypersphere of radius R in d dimensions. For later use, we define $\sigma_{ij} = \frac{\sigma_i + \sigma_j}{2}$ as the sum of two particles radii.

For hard particles of arbitrary shape, a jammed state must fulfill an excluded volume constraint, as well as a set of mechanical constraints, related to force and torque balance, absence of attractive forces, etc. [7]. Restricting to frictionless hard spheres, as we do here, the excluded volume constraint reads $|\mathbf{r}_i - \mathbf{r}_j| \geq \sigma_{ij}$, $\forall 1 \leq i < j \leq N$. This simply states that spheres cannot overlap. Moreover, among the mechanical requirements, only the force balance condition is relevant, because once it holds the other mechanical constraints are automatically satisfied. In summary, for frictionless hard spheres the excluded volume constraint and the force balance condition for each particle are necessary and sufficient conditions for having jammed configurations [7,24].

Our aim is to bring an initial HS configuration, with initial packing fraction $\varphi < \varphi_J$, to its jamming point at φ_J . We wish

to reach the jamming point by increasing the packing fraction, $\varphi \rightarrow \varphi_J^-$, until the system becomes mechanically rigid, and using a procedure that does not allow any overlap among particles at any time. All of this can be recast as a constrained optimization problem (OP), as we show next. Without loss of generality, we will assume that the system volume is fixed; φ is then a monotonically increasing function of $\bar{\sigma}$, and maximizing it is equivalent to finding the largest factor by which the particle's diameter can be inflated. Naturally, we look for the optimal value of such inflation factor allowing for particles to be rearranged. Thus, letting $\bar{\mathbf{s}} = \{\mathbf{s}_i\}_{i=1}^N$ be the possible particles displacements, $\sqrt{\Gamma}$ the inflation factor, and using $*$ to denote the optimal value of a quantity, the OP we consider is: find $(\bar{\mathbf{s}}^*, \Gamma^*)$ such that if we transform our HS system according to $(\bar{\mathbf{r}} \leftarrow \bar{\mathbf{r}} + \bar{\mathbf{s}}^*, \sigma_i \leftarrow \sqrt{\Gamma^*} \sigma_i)$, then the packing fraction reaches a (possibly local) maximum value. $\bar{\mathbf{s}}$ and Γ must fulfill the nonoverlapping constraint between spheres; thus, this is a *constrained* OP. The reason for using the square root of Γ will become apparent when we will write the linearized version of this OP [see Eq. (2)]. Clearly, by taking an inflation factor equal for all the particles we preserve the degree of polydispersity of the system.

Hence, in a system with periodic boundary conditions and in absence of external forces, finding a jammed HS configuration is equivalent to solving the following constrained OP:

$$\max \Gamma \quad (1a)$$

$$G_{ij}(\bar{\mathbf{s}}, \Gamma) := \Gamma \sigma_{ij}^2 - |\mathbf{r}_i + \mathbf{s}_i - (\mathbf{r}_j + \mathbf{s}_j)|^2 \leq 0, \quad (1b)$$

$$\forall 1 \leq i < j \leq N,$$

where the OP's variables are $(\bar{\mathbf{s}}, \Gamma)$, while the particles' position and size, $(\bar{\mathbf{r}}, \bar{\sigma})$, play the role of constant parameters. Notice that the constraints in Eq. (1b) are symmetric upon exchanging i and j , making the case $j < i$ redundant. The results presented here can be easily generalized to systems with closed boundaries. In such a case, in the OP (1) one should also require that $\frac{\Gamma \sigma_i^2}{4} \leq (r_{i,\mu} + s_{i,\mu})^2$ and $\frac{\Gamma \sigma_i^2}{4} \leq (L - (r_{i,\mu} + s_{i,\mu}))^2$ for all $i = 1, \dots, N$ and $\mu = 1, \dots, d$. This would add extra constraints that should be taken into account when counting the effective degrees of freedom and analyzing the stability.

Equation (1) represents an exact, albeit nonconvex, formulation of jamming as an OP. Indeed, as depicted in Fig. 1(a), the set of points $(\Gamma, \bar{\mathbf{s}})$ that satisfy the nonoverlapping constraints (represented as the gray region) is nonconvex. Intuitively, the excluded volume constraint between particles creates many “holes” in the set of possible solutions, termed *feasible region*. Thus, as often occurs in nonconvex OP's, the problem (1) becomes intractable for large system sizes [$dN \sim \mathcal{O}(10^2)$], and one has to resort to approximations.

A simple, yet powerful one is to assume that the starting HS configuration is already close to the jamming point; therefore, any feasible displacement $\bar{\mathbf{s}}$ has a negligible magnitude in comparison with the smallest distance between particles centers. Thus, terms of order $\mathcal{O}(|\mathbf{s}_i|^2)$ can be neglected in each of the constraints G_{ij} . This amounts to the so-called approximation of small displacements [69], which has been successfully used to analyze the mechanical properties of rigid structures. Within this approach, the exact OP becomes a




FIG. 1. Geometry of CALiPPSO for a problem of a single moving particle (1, black dotted) and four fixed ones (2, green; 3, blue; 4, orange; 5, purple). In all panels, the exact nonoverlapping constraints, $\{G_{ij}\}_{j=2}^5$, are drawn with dashed circles, while their linearized version, $\{F_{ij}\}_{j=2}^5$, are identified by solid lines. Panel (a) shows the nonconvex OP of Eq. (1) and the associated jamming LOP, Eq. (2). The size of particle 1 is indicated by the dotted circle, but only the position of its center (black dot) is relevant since the contribution of σ_1 has been included in the constraints. $\{G_{ij}\}_{j=2}^5$ induce “holes” in the set of possible displacements (gray region), making it nonconvex. In contrast, $\{F_{ij}\}_{j=2}^5$ always define a convex set: a polytope (light pink). Note that such polytope is always strictly contained in the original feasible set. Panel (b) depicts the optimal displacement, \mathbf{s}_1^* (black star), of the jamming LOP instance of panel (a). As expected for any LOP, \mathbf{s}_1^* is located at the intersection of the linearized constraints, evaluated at the optimal solution, $\{F_{ij}^* = F_{ij}(\mathbf{s}_1^*, \Gamma^*)\}_{j=2}^5$. Note, however, that \mathbf{s}_1^* does not saturate the analogous exact constraints. In panel (c), we show the new instance of the jamming LOP obtained after updating $\mathbf{r}_1 \leftarrow \mathbf{r}_1 + \mathbf{s}_1^*$ and $\sigma_i \leftarrow \sqrt{\Gamma^*} \sigma_i$, for $i = 1, \dots, 5$. This panel shows that CALiPPSO quickly reduces the size of the feasible region, which means that the linear constraints better approximate the exact ones. In this last panel the new size of particle 1 is not shown. Panels (b) and (c) have been magnified by a factor 9.6 and 4.4 with respect to panel (a), respectively.

linear optimization problem (LOP), so it is guaranteed to be convex [70,71]. The resulting LOP reads

$$\max \Gamma \quad (2a)$$

$$F_{ij}(\bar{\mathbf{s}}, \Gamma) := -2\mathbf{r}_{ij} \cdot \mathbf{s}_{ij} + \Gamma \sigma_{ij}^2 - |\mathbf{r}_{ij}|^2 \leq 0, \quad (2b)$$

$$\forall 1 \leq i < j \leq N,$$

with $\mathbf{r}_{ij} := \mathbf{r}_i - \mathbf{r}_j$, and $\mathbf{s}_{ij} := \mathbf{s}_i - \mathbf{s}_j$. Despite the fact that F_{ij} depends on $\bar{\mathbf{s}}$ only through \mathbf{s}_{ij} , to simplify the notation, we will consider the general case where it depends on the full displacement vector. Notice that by increasing the diameters by a factor $\sqrt{\Gamma}$ we are able to keep both the objective function (Γ) and the constraints strictly linear. In the following, we will refer to Eq. (2) as the jamming LOP. For future use, we introduce the Lagrange multiplier, $\lambda_{ij} \geq 0$, associated to the constraint $F_{ij}(\bar{\mathbf{s}}, \Gamma)$. The Lagrange multipliers play a fundamental role in our algorithm since they determine the set of contact forces at jamming, as we will show below when analyzing the dual of the jamming LOP [Eq. (5)].

The main advantage of using the LOP in Eq. (2) with respect to the exact OP is that the optimal solution of the former, $(\bar{\mathbf{s}}^*, \Gamma^*)$, can be found by means of several linear programming methods. Moreover, it is straightforward to show that any point that satisfies the set of linear constraints $\{F_{ij}(\bar{\mathbf{s}}, \Gamma)\}_{1 \leq i < j \leq N}$ satisfies also the exact ones $\{G_{ij}(\bar{\mathbf{s}}, \Gamma)\}_{1 \leq i < j \leq N}$; indeed, $G_{ij}(\bar{\mathbf{s}}, \Gamma) = F_{ij}(\bar{\mathbf{s}}, \Gamma) - |\mathbf{s}_{ij}|^2$. This feature is illustrated by Fig. 1(a), where it is shown that the feasible region of the LOP (pink) is contained in the analogous set of the original OP (gray).

However, a significant drawback is that the optimal solution of the jamming LOP is, in general, a sub-optimal solution for the exact OP. That is, solving Eq. (2) *once*, does

not necessarily yield a jammed packing because, even if the linear constraints (2b) are saturated, the exact constraints in Eq. (1b) might not be so. More precisely, sub-optimality is a consequence of the fact that an optimal solution of the LOP is always located in the boundary of the polytope defined by the linear constraints [70,71], and that such point, in general, does not minimize the separation between inflated particles. Therefore, some of the exact constraints $\{G_{ij}\}$ remain unsaturated. Geometrically, this means that $\bar{\mathbf{s}}^*$ of the jamming LOP is determined by the intersection of the set of constraints $\{F_{ij}\}$ that are saturated once evaluated with Γ^* . But if the intersection point does not coincide with the point where F_{ij} are tangent to G_{ij} , then $(\bar{\mathbf{s}}^*, \Gamma^*)$ will be suboptimal with respect to the exact OP. This is depicted in Fig. 1(b) for a single particle.

Yet, importantly the configuration obtained after solving the LOP, $(\bar{\mathbf{r}}, \bar{\sigma}) \leftarrow (\bar{\mathbf{r}} + \bar{\mathbf{s}}^*, \sqrt{\Gamma^*} \bar{\sigma})$, will have no overlaps and a larger packing fraction. Consequently, it will be closer to jamming and we can use it to generate a new instance of the LOP (2), that better approximates the exact OP. This is shown in Fig. 1(c).

The above considerations suggest that if $x(t)$ denotes the value of the quantity x in the t th instance of the jamming LOP (2), then we can use the optimal solutions as initialization points, and proceed iteratively as $(\bar{\mathbf{r}}(t+1), \bar{\sigma}(t+1)) = (\bar{\mathbf{r}}(t) + \bar{\mathbf{s}}^*(t), \sqrt{\Gamma^*(t)} \bar{\sigma}(t))$ to reach the jamming point. Jamming is realized when the particles cannot be further inflated nor moved, i.e., $(\bar{\mathbf{s}}^*(n), \Gamma^*(n)) = (\bar{\mathbf{0}}, 1)$; we henceforth refer to this optimal solution as the *convergence condition* of the CALiPPSO algorithm and n as the number of linear optimizations required to reach it.

We can intuitively understand the functioning of CALiPPSO in the following way. During the first few

linear optimizations and the corresponding configuration updates, the space of possible solutions is quickly reduced, as depicted in Figs. 1(a)–1(c) for a single particle. This reduction is mainly caused by the tighter bounds imposed by the linear constraints $F_{ij}(\vec{s}, \Gamma)$, in comparison with the exact nonoverlapping constraints of Eq. (1b) [cf. Fig. 1(a)]. Importantly, the polytope defined by the constraints $F_{ij}(\vec{s}, \Gamma)$ suppresses rearrangements that would allow particles to escape from the cages formed by their neighbors, thus efficiently preventing hopping over (entropic) barriers [cf. Fig. 1(c)]. This feature is particularly relevant to prevent crystallization in monodisperse configurations. In this way, the linear constraints of the CALiPPSO algorithm are responsible for efficiently trapping particles, a geometric metaphor for the enchantments used by Calypso to keep Odysseus captive in her island for several years.

As a final remark, we highlight that the convergence condition of CALiPPSO corresponds also to an optimal solution of the exact OP. That is, if $(\vec{r}^{(J)}, \vec{\sigma}^{(J)})$ defines a jammed packing obtained using CALiPPSO and we use these values to generate another instance of the OP (1), then $(\vec{s}^*, \Gamma^*) = (\vec{0}, 1)$ and $(\vec{r}^{(J)}, \vec{\sigma}^{(J)})$ constitutes a local optimum of the exact OP as well. At least in $d = 2$ and $d = 3$ the global optimum corresponds instead to a crystalline structure. For instance, the densest packing of $3d$ monodisperse spheres corresponds to an FCC structure with $\phi_j^{(\text{FCC})} = \pi/\sqrt{18} \approx 0.74$. Similarly, in $d < 10$, the densest known monodisperse packings have an ordered structure [72,73]. Since our aim is to study amorphous packings, we want to avoid precisely such ordered or partially crystallized solutions. It is therefore convenient that the CALiPPSO algorithm prevents particles from performing large displacements: in such a way only a local optimum, very likely corresponding to a disordered configuration, can be obtained. The results we report below confirm that we can avoid producing ordered monodisperse packings whenever a random initial configuration is chosen for $d \geq 3$. When $d = 2$, instead, (partial) crystallization always ensues in monodisperse packings. This is a consequence of the Euler criterion applied to the planar graph formed by the network of contacts [74,75] (see Appendix A for more details). As we show below, adding bidispersity or polydispersity helps to prevent crystallization in bidimensional systems.

While these arguments show that CALiPPSO is able to generate a maximally dense disordered configuration of hard spheres, a more rigorous analysis is needed to show that such configuration is a mechanically stable packing. This is the issue we address next.

B. CALiPPSO produces well-defined, globally stable jammed states

In this section, we analytically prove that a HS packing obtained once the CALiPPSO convergence condition is reached corresponds to a *valid jammed state*. This means that: (a) such HS configuration satisfies the excluded volume and force balance constraints for each particle (local property); (b) it is a collectively stable packing (global property). The stability property follows from a relation between the number of contacts (N_c) and the number of degrees of freedom in a configuration (N_{dof}), specifically, $N_c \geq N_{\text{dof}}$ [3,7,30,76–78].

Counting N_{dof} requires some care due to the presence of symmetries and unstable particles. Fortunately, CALiPPSO also provides a univocal way to determine N_{dof} ; see Eq. (14) below. Before continuing, we recall some common terminology. When the latter inequality does not hold ($N_c < N_{\text{dof}}$), a packing is said to be *hypostatic*; if the equality is verified ($N_c = N_{\text{dof}}$), the packing is *isostatic*; finally, a packing is *hyperstatic* when the strict inequality is fulfilled ($N_c > N_{\text{dof}}$). Note that the condition (b) does not follow necessarily from (a) since there can be hypostatic packings where force balance holds for each particle and no overlaps are present [3,69].

It is known that critical jamming emerges together with isostaticity. Remarkably, isostaticity is verified in packings produced via CALiPPSO once all the relevant degrees of freedom are considered, as we will show here. For our proof, we will adopt a pedagogical approach, so several logical steps will be explicitly made, and we will emphasize the connection between a geometrical analysis of jamming and the optimization approach developed here. Our proof will go along the following lines: to show (a) we will first notice that the excluded volume constraint is verified by construction. Then, we will use general results of convex optimization theory [71] to show rigorously that optimality of the jamming LOP (2) implies mechanical equilibrium. The property (b) will be derived from the existence of the solution of the force balance equations. These equations are encoded in the constraints of the dual optimization problem associated to the jamming LOP; see Eq. (5) below. Leveraging on the property that such dual problem is itself a LOP, we will show that the solution of the force balance equations is unique. Since this can only happen when $N_c = N_{\text{dof}}$, the packings produced must be isostatic.

Let us begin by pointing out that, since $G_{ij}(\vec{s}, \Gamma) \leq F_{ij}(\vec{s}, \Gamma)$, $\forall 1 \leq i < j \leq N$, it is guaranteed that the solutions of the LOP (2) always satisfy the nonoverlapping constraints of HS systems, Eq. (1b). To continue the proof, we will need few results of linear optimization, Eqs. (3)–(7) below, particularized for the jamming LOP. We first introduce some notation. Let $M := \frac{N(N-1)}{2}$ and $\tilde{\lambda} = \{\lambda_{ij} \geq 0\}_{i < j}$ be the set of M nonnegative dual variables or Lagrange multipliers associated with the constraints $\{F_{ij}\}_{i < j}$. In the following, we will show that we only need to consider $\mathcal{O}(N)$ constraints in order for CALiPPSO to work. For the remaining part of this section, we will use the simplified notation $\vec{x} := (\vec{s}, \Gamma)$, and define the $(dN + 1)$ -dimensional vector $\vec{y} := (\vec{0}, -1)$; the zero vector of such space will be denoted by $\vec{0} = (\vec{0}, 0)$. In terms of this new variables, the objective function (2a) is equal to $\vec{y} \cdot \vec{x}$. We rewrite the jamming LOP (2) in the more conventional form:

$$\min \vec{y} \cdot \vec{x}, \quad (3a)$$

$$\text{s.t. } \mathcal{F}\vec{x} \leq \tilde{\rho}, \quad (3b)$$

where \leq denotes element-wise comparison; \mathcal{F} is a $M \times (dN + 1)$ matrix with entries

$$\mathcal{F}_{ij}^{\mu k} = -2r_{ij,\mu}(\delta_{ik} - \delta_{jk}), \quad (4a)$$

$$\mathcal{F}_{ij}^{dN+1} = \sigma_{ij}^2, \quad (4b)$$

for $1 \leq k \leq N$, $1 \leq \mu \leq d$, and $1 \leq i < j \leq N$; while $\tilde{\rho}$ is an M -dimensional vector with components $\rho_{ij} = |\mathbf{r}_{ij}|^2$, $i < j$. Thus, expression (3b) is nothing else than the full set of

nonoverlapping constraints, emphasizing their affine character. In addition, for later use we introduce the *dual* OP of Eq. (3):

$$\max - \tilde{\rho} \cdot \tilde{\lambda}, \quad (5a)$$

$$\text{s.t. } \mathcal{F}^T \tilde{\lambda} + \tilde{\mathbf{y}} = \tilde{\mathbf{0}}, \quad (5b)$$

$$\tilde{\lambda} \succeq \tilde{\mathbf{0}}. \quad (5c)$$

Notice that the dual OP is also a LOP, but one in which the inequality constraints, Eq. (3b) [or Eq. (2b)] have been replaced by equalities, Eq. (5b). This is a standard property of LOPs [70,71], and will be fundamental to prove the isostaticity of CALiPPSO packings.

Readers familiar with a geometrical interpretation of the mechanical properties of granular packings [69] might recognize that Eq. (5b) is equivalent to the mechanical equilibrium condition, provided that the Lagrange multipliers are identified with the contact forces. We expect that the details of our proof will make clear the correspondence between the optimization and the geometrical approach. In fact, the CALiPPSO algorithm works by exploiting the deep connection between the mechanical equilibrium condition to be fulfilled by jammed packings [captured by the jamming LOP's dual (5)] and the constraints that such condition imposes on the density and configurational degrees of freedom of the system [represented in the primal LOP (3)]. In particular, one can rely on the *strong duality* property of LOPs [70,71], which implies that if either the primal or the dual LOP has a finite optimal solution, so does the other. When both LOPs can be solved to optimality, their optimal values are equal (or, in other words, there is no duality gap). In our case, this means that if $\tilde{\mathbf{x}}^*$ is an optimal solution of Eq. (3), then $\tilde{\lambda}^*$ —the optimal solution of the dual (5)—is such that $\tilde{\rho} \cdot \tilde{\lambda}^* = -\tilde{\mathbf{y}} \cdot \tilde{\mathbf{x}}^* = \Gamma^*$. Moreover, as we will show next, the equality constraints (5b) imposed on $\tilde{\lambda}^*$ are equivalent to the force balance requirement for each particle.

An explicit derivation of the force balance equations for $\tilde{\lambda}$ can be obtained using the Lagrangian, defined as

$$\mathcal{L}(\tilde{\mathbf{x}}, \tilde{\lambda}) = \tilde{\mathbf{y}} \cdot \tilde{\mathbf{x}} + \sum_{1 \leq i < j \leq N} \lambda_{ij} F_{ij}(\tilde{\mathbf{x}}). \quad (6)$$

Because of the strong duality between the LOP (3) and (5), the Karush–Kuhn–Tucker (KKT) conditions [71] imply the following results for any primal and dual optimal points $(\tilde{\mathbf{x}}^*, \tilde{\lambda}^*) := (\tilde{\mathbf{s}}^*, \Gamma^*, \tilde{\lambda}^*)$:

$$\lambda_{ij}^* F_{ij}(\tilde{\mathbf{x}}^*) = 0, \quad (7a)$$

$$\nabla \mathcal{L}(\tilde{\mathbf{x}}^*, \tilde{\lambda}^*) = \tilde{\mathbf{y}} + \sum_{1 \leq i < j \leq N} \lambda_{ij}^* \nabla F_{ij}(\tilde{\mathbf{x}}^*) = \tilde{\mathbf{0}}, \quad (7b)$$

where $\nabla = (\frac{\partial}{\partial \tilde{\mathbf{s}}}, \frac{\partial}{\partial \Gamma})$, and $\frac{\partial}{\partial \tilde{\mathbf{s}}}$ has been used as a shorthand notation of $(\frac{\partial}{\partial s_{1,1}}, \frac{\partial}{\partial s_{1,2}}, \dots, \frac{\partial}{\partial s_{1,d}}, \frac{\partial}{\partial s_{2,1}}, \dots, \frac{\partial}{\partial s_{N,d}})$. Importantly, equality (7a), termed *complementary slackness*, implies that if the constraints are inactive, i.e., $F_{ij}(\tilde{\mathbf{x}}^*) < 0$, then $\lambda_{ij} = 0$. Conversely, a Lagrange multiplier will only be positive, $\lambda_{ij} > 0$, when the associated constraint is active, $F_{ij}(\tilde{\mathbf{s}}^*, \Gamma^*) = 0$.

Such active Lagrange multipliers play a major role in our algorithm because they can be used to obtain the physical contact forces between particles. To emphasize that our

analysis of the mechanical equilibrium of the packing is based on the linear constraints (2b) that are saturated in an optimal solution, we will use $[ij]$, with $i < j$, to indicate that there is an active linear constraint between particles i and j . So, we define $\mathcal{C} = \{[ij]\}$, the set of linear contacts, whose cardinality is $|\mathcal{C}| = N_c \ll M$. The last inequality follows from the fact that not all particles are in contact with each other; therefore, the amount of positive Lagrange multipliers is much smaller than M as long as $d \ll N$, as we assume here. It is useful to consider the N_c -dimensional vector of only positive dual variables, $\underline{\lambda} = \{\lambda_{ij}\}_{[ij] \in \mathcal{C}}$.

Proving the mechanical equilibrium condition is now straightforward: the linearity of the LOP (3) implies that each of the derivatives of Eq. (7b) results in one entry of \mathcal{F} (for fixed $i < j$). Thus, plugging Eq. (4) into Eq. (7b) leads to the following equation for the components associated to the i th particle:

$$\sum_{\substack{j \\ j \neq i}}^{1,N} \lambda_{ij}^* \mathbf{r}_{ij} = \sum_{\substack{j \in \partial i \\ [ij] \in \mathcal{C}}} \lambda_{ij}^* \mathbf{r}_{ij} = \mathbf{0}, \quad (8)$$

where, in the second equality, ∂i is the set of all linear contacts of particle i . That is, because of the complementary slackness condition (7a), the sum over $N - 1$ terms is reduced to one that only contains active dual variables, usually of order d . On the other hand, the component associated to $\frac{\partial}{\partial \Gamma}$ of Eq. (7b) reads

$$\sum_{1 \leq i < j \leq N} \lambda_{ij}^* \sigma_{ij}^2 = \sum_{[ij] \in \mathcal{C}} \lambda_{ij}^* \sigma_{ij}^2 = 1. \quad (9)$$

We note in passing that from the last equation we can estimate how far from the jamming point a given $\tilde{\mathbf{x}}^*$ is. Indeed, given that $\Gamma^* \geq 1$, from the absence of duality gap mentioned above, it follows that $\tilde{\rho} \cdot \tilde{\lambda}^* \geq 1$. Whence, from Eq. (9) and the complementary slackness property, it is straightforward to obtain

$$\sum_{[ij] \in \mathcal{C}} \lambda_{ij}^* (|\mathbf{r}_{ij}|^2 - \sigma_{ij}^2) \geq 0. \quad (10)$$

These are conditions to be fulfilled by any optimal solution $\tilde{\mathbf{x}}^*$ of the jamming LOP. However, when the convergence criterion $\tilde{\mathbf{x}}^* = (\tilde{\mathbf{0}}, 1)$ is reached, we have $F_{[ij]}(\tilde{\mathbf{0}}, 1) = 0$, and equality holds in Eq. (10). This means that $|\mathbf{r}_{[ij]}| = \sigma_{[ij]}$, i.e., linear contacts become physical contacts, and \mathcal{C} then determines the full network of contacts at jamming. Once we know that the norm of $\mathbf{r}_{[ij]}$ can be fixed, we can rescale the corresponding dual variables $\lambda_{ij}^* = f_{ij}/\sigma_{ij}$ and introduce the unit vector $\mathbf{n}_{ij} = \frac{\mathbf{r}_{ij}}{|\mathbf{r}_{ij}|}$ to rewrite Eq. (8) as

$$\sum_{j \in \partial i} f_{ij} \mathbf{n}_{ij} = \mathbf{0}. \quad (11)$$

This is the force balance equation for the i th sphere. We can write analogous equations for the full configuration as

$$\mathcal{S} \underline{\mathbf{f}} = \tilde{\mathbf{0}}, \quad (12)$$

where $\underline{\mathbf{f}} = \{f_{ij}\}_{[ij] \in \mathcal{C}}$ is the vector containing the contact forces magnitudes, and \mathcal{S} is a $dN \times N_c$ matrix, whose entries are given by $\mathcal{S}_{k,\mu}^{[ij]} = (\delta_{ik} - \delta_{jk}) n_{ij,\mu}$. Crucially, this expression

is identical to the one derived in previous works [13,24,27,69], and it determines the force balance condition for jammed packings. Notice that \mathbf{f} represents *contact* forces, which are finite despite the singular potential of hard spheres. The connection between contact forces and dual variables is completely analogous to the one between generalized forces and constraint conditions in Lagrangian mechanics; namely, \mathbf{f}_{ij} can be obtained from the Lagrange multiplier λ_{ij}^* times the derivative of the constraints involved [cf. Eq. (7b) from the KKT conditions]. This completes the proof that, upon convergence, CALiPPSO generates packings in mechanical equilibrium and with physical contacts, even if in the intermediate steps this latter feature is not necessarily true.

The global stability property or, equivalently, the requirement that a jammed state produced with our algorithm is not hypostatic, does not follow from the force balance condition, Eq. (12), alone. To prove that CALiPPSO always produces isostatic packings (property (b) above), we note that once convergence has been reached, it follows that $\mathcal{S}_{k,\mu}^{[ij]} = -\frac{1}{2\sigma_{ij}}(\mathcal{F}^T)_{k,\mu}^{i,j}$ if particles i and j are in contact. This means that Eq. (12) is equivalent to the equality constraints of the dual LOP (5b) but with a reduced matrix which only involves particles in contact. Letting \mathcal{R} be such a $(dN + 1) \times N_c$ matrix, Eq. (5b) can be rewritten as

$$\mathcal{R}\underline{\lambda} + \vec{y} = \vec{0}. \quad (13)$$

This expression forms a system of $dN + 1$ equations in N_c unknowns, in which the first dN equations form a homogeneous system, equivalent to Eq. (12) (once rescaled by $\frac{1}{\sigma_{ij}}$). Equation (13) represents the link between the geometric interpretation of HS packings, and the optimization perspective given by the KKT conditions. Indeed, the first dN rows of Eq. (13) are nothing else but the rightmost equality of Eq. (8) for each particle, while the last row corresponds to Eq. (9).

Proving the isostaticity of CALiPPSO packings requires to accurately count the degrees of freedom. We first notice that, because of our assumption of periodic boundaries, there are d uniform translations that leave relative distances and displacements among particles invariant; thus, the number of degrees of freedom is decreased by d . This is nicely reflected by the fact that out of the dN homogeneous equations in Eq. (13), d are linearly dependent. Another important consideration is the presence of rattlers, which are particles with d or fewer contacts. Due to the low number of contacts, these particles are unstable, and in most cases they do not belong to the backbone of the network of contacts. Therefore, they do not contribute to the rigidity of the configuration. The only exception are monodisperse systems in $2d$, where rattlers are subject to forces; this case is discussed in Appendix A. Rattlers should be excluded when counting the number of contacts and degrees of freedom, and only spheres with at least $d + 1$ contacts—henceforth termed stable particles—should be considered. Of course, identifying whether a particle is stable or not can only be done a posteriori, when \mathcal{C} is constructed after CALiPPSO has converged. If N_s is the number of stable particles, in the absence of any other symmetries or external constraints, then we have

$$N_{\text{dof}} = d(N_s - 1) + 1, \quad (14)$$

where the extra degree of freedom is a consequence of the fact that Γ is also a variable of the jamming LOP. Equivalently, one can think of density as an additional degree of freedom. Thus, when we say that the isostatic condition $N_c = N_{\text{dof}}$ is verified, we mean that all the degrees of freedom of the jamming LOP (i.e., the displacements of stable particles and inflation factor) have been considered.

In contrast, if one considers only configurational degrees of freedom, i.e., $N'_{\text{dof}} = d(N_s - 1)$, then we have $N_c = N'_{\text{dof}} + 1$, which makes the packings produced by CALiPPSO hyperstatic. This amounts to say that requiring CALiPPSO configurations to have a finite bulk modulus (or being rigid), imposes an extra constraint [30,76]. To highlight our optimization approach, we will use *isostatic* when the degrees of freedom are counted as in Eq. (14).

As mentioned above, the d uniform translations reduce the number of linearly independent equations in Eq. (13). Naturally, this feature is also present in \mathcal{S} . Therefore, considering Eq. (12), it is easy to see that to have a consistent system of equations, it must happen that $N_c \geq N'_{\text{dof}}$. However, if $N_c = N'_{\text{dof}}$, then only the homogeneous solution exists. Yet, Eq. (9) prevents such scenario, whence $N_c \geq N'_{\text{dof}} + 1 = N_{\text{dof}}$. This shows that CALiPPSO packings can never be hypostatic, and thus are always collectively stable.

Finally, to show that the packings obtained with our method are isostatic we make use of the fact that there is zero duality gap between the jamming primal and dual LOPs, Eqs. (3) and (5), respectively. The absence of duality gap implies that the solution of the linear system (13) is unique [69]. Given that $\underline{\lambda} > \underline{0}$, and that Eq. (13) is not a homogeneous system (recall $y_{dN+1} = -1$), the uniqueness of the solution implies that the number of independent equations matches the number of unknowns, $N_{\text{dof}} = N_c$. Therefore, the resulting packings are isostatic. This completes the proof that packings produced with the CALiPPSO algorithm are valid jammed states.

Let us note that the result of two paragraphs above, i.e., $N_c \geq N_{\text{dof}}$, only relies on the KKT conditions [Eq. (7)], and thus is more general than the isostaticity property. In some “pathological” cases, it may happen that $N_c > N_{\text{dof}}$. In our experience, hyperstaticity only occurs in $2d$ monodisperse packings, where also large crystalline domains are formed. We comment further on this particular case in Appendix A. Importantly, all the tests we performed in $d \geq 3$ (see Secs. III, IV and Appendix B) have never produced such hyperstatic configurations, even with monodisperse systems. Similarly, in bidimensional, polydisperse systems isostaticity is recovered.

It should now be clear that the dual jamming LOP [in particular Eq. (5b)], together with the complementary slackness property [Eq. (7a)], contain all the requirements to guarantee that CALiPPSO produces valid jammed states. We also showed that an equivalent condition is that any CALiPPSO optimal dual solution $\underline{\lambda}^*$ must fulfill Eq. (13). We emphasize that the matrices \mathcal{S} , \mathcal{F} , and \mathcal{R} are determined entirely by the geometrical features of the configuration. Notably however, we derived them not from geometric considerations, but following results of optimization theory; specifically the KKT conditions, Eq. (7). The CALiPPSO algorithm works by exploiting the correspondence between optimization and

geometric descriptions of jammed packings. This can be considered the heart of our proof: any optimal solution $\bar{\mathbf{x}}^*$ must fulfill the KKT conditions, as written in Eqs. (8) and (9) (due to the convex nature of the jamming LOP), and at convergence these conditions are equivalent to Eq. (13) (a geometric property). Mathematically, we can understand that the CALiPPSO algorithm succeeds at producing isostatic configurations because the dual optimization problem associated to the jamming LOP [see Eq. (5)] is itself a LOP. Given that, once an optimal solution is found the two systems of linear equations associated to their constraints are satisfied simultaneously (because there can be no duality gap), the only possibility is that $N_c = N_{\text{dof}}$.

Before closing this part, we comment on an additional property that can be derived from Eq. (9), for a packing in arbitrary dimensions, d . Notice that for monodisperse systems, or whenever the distribution of diameters is significantly peaked around its mean value $\bar{\sigma}$ (i.e., whenever $\text{Var}[\bar{\sigma}] \ll \bar{\sigma}^2$), it is easy to see that $\varphi \sim N\bar{\sigma}^d$. In this case, the scaling of Eq. (9) with the system size implies that the mean force \bar{f} is such that $N\bar{f} \sim 1/\bar{\sigma}$, whence we obtain $\bar{f} \sim N^{1/d-1}$.

C. Algorithmic implementation of CALiPPSO

The CALiPPSO algorithm is rather simple; it consists of a single loop that iterates over successive LOP instances to reach the jamming point.

Its performance can be easily enhanced by using neighbor lists to reduce the number of constraints of the LOP. Indeed, when two particles are far apart, their associated constraint in Eq. (2b) becomes irrelevant. Hence, for sufficiently distant particles, the evaluation of the inequality (2b) can be omitted without affecting the optimization procedure. By implementing the neighbor-list approach, instead of including the $N(N-1)/2$ possible constraints, we only consider $M' \sim c_d N$ of them, where c_d is a prefactor that depends on the dimensionality and should be, at most, of the order of the kissing number (i.e., the maximum number of nonoverlapping spheres such that each of them touch a common sphere) in the corresponding dimension. Clearly, this reduces the size of the constraints matrix \mathcal{F} of Eq. (4) to be $M' \times (dN+1)$.

The neighbor list is constructed utilizing a cutoff distance, $\ell(\varphi)$, which (possibly) depends on the system's packing fraction, as explained below. For each particle i we define its list of neighbors, $\tilde{\partial}i := \{j \mid |\mathbf{r}_{ij}| \leq \ell(\varphi)\}$, where the distance between particles i and j is computed following the nearest image convention [79]. Therefore, the constraint $F_{ij}(\bar{\mathbf{s}}, \Gamma)$ is included only if $j \in \tilde{\partial}i$ (and $i < j$). That is, for a given packing fraction, the full set of constraints becomes $\tilde{\mathbf{F}}(\bar{\mathbf{s}}, \Gamma) = \{F_{ij}(\bar{\mathbf{s}}, \Gamma) \mid i < j, j \in \tilde{\partial}i\}$.

It is useful to consider a cutoff distance dependent on φ since at low packing fraction, when a single update can result in particles being displaced over large distances, $\ell(\varphi)$ should be kept large enough to ensure that no overlaps occur even after such large displacements. In contrast, for $\varphi \lesssim \varphi_J$, when each linear optimization iteration generates only very small rearrangements, $\ell(\varphi)$ can be set to a small value, keeping track only of the nearest neighbors for each particle. Notice that we have assumed that ℓ is the same for each particle, but it is straightforward to generalize our algorithm to

the case where each particle has a different cutoff distance, $\ell_i(\varphi)$. This situation could be useful, for instance, with highly polydisperse packings, in which the smallest particles might perform larger displacements and therefore more neighbors need to be taken into account to avoid overlaps. In such case, the list of neighbors for the i th particle could be defined as $\tilde{\partial}i := \{j \mid |\mathbf{r}_{ij}| \leq \max[\ell_i(\varphi), \ell_j(\varphi)]\}$.

The CALiPPSO algorithm is reported as pseudocode in Algorithm 1. Before further comments, we anticipate that more details on how to initialize the algorithm will be addressed in Sec. III, while the analysis of the convergence time of CALiPPSO is postponed to Sec. IV.

As a first remark, we highlight that CALiPPSO has no free parameters, except for $\ell(\varphi)$. Working with monodisperse configurations, we found that, if the initial packing fraction is not too small, setting $\ell/\sigma \in [3, 4]$ produces good results, while when the system is very close to jamming (i.e., during the last linear optimizations), further reducing the cutoff distance to $\ell = 1.4\sigma$ suffices. However, since φ_J is not known *a priori* for a given configuration, we choose ℓ according to the optimal inflation factor from the previous linear optimization, Γ_0^* . For instance, in our tests we used $\ell = 3.5\sigma$ when $\sqrt{\Gamma^*} - 1 \geq 10^{-5}$, and $\ell = 1.4\sigma$ otherwise. We verified that the results are insensitive to specific value of $\ell \in [1.4, 5]\sigma$.

Second, if the initial packing fraction is not close to φ_J , say $\varphi/\varphi_J < 0.5$ imposing bounds on the particles' displacements, $|\mathbf{s}_i| < s_{\text{bound}}$ for $i = 1, \dots, N$, is convenient. In this way, one can avoid both large rearrangements (that could lead to crystallization in $3d$ monodisperse systems), and the need to make $\ell(\varphi)$ too large. We found that the naive bound, $s_{\text{bound}} = \frac{1}{2\sqrt{d}}[\ell(\varphi) - \sqrt{\Gamma_0^*}\sigma]$, was enough to avoid overlaps in all the cases we tested. However, if the degree of polydispersity is very broad, tighter bounds might be needed. Additionally, bounding $|\mathbf{s}_i|$ might be useful also at later stages of the chain of linear optimizations to effectively reduce the feasible region of the jamming LOP and speed up the optimization. This should be done with some care because, if these constraints become active, they would play the role of external forces. Therefore, they must be taken into account when assessing the mechanical equilibrium of the configuration. If ignored, active displacement bounds might cause a packing to be non-isostatic [in the sense defined above i.e., $N_c = N_{\text{dof}}$ with N_{dof} as in Eq. (14)], given that N_c only considers contact forces. A practical solution to guarantee isostaticity and that only real contacts are included, is to perform the last linear optimization without bounds on any $|\mathbf{s}_i|$.

In the pseudocode of Algorithm 1 we have implicitly assumed, without loss of generality, that $\tilde{\lambda}^*$ is obtained *simultaneously* when solving the jamming LOP (line 8). This is certainly the case when the jamming LOP (2) is solved using interior-point methods [70,71] as we do here (see below). However, if the optimal solution is obtained using, e.g., the primal simplex method [70], the Lagrange multipliers would be computed *after* such solution is found. Conversely, if the dual simplex method was employed [i.e., if the LOP (5) is solved instead of the primal, original jamming LOP], $(\bar{\mathbf{s}}^*, \Gamma^*)$ would be obtained from $\tilde{\lambda}^*$. As already remarked, these arguments rely on the strong duality theorem: once the primal

Algorithm 1 CALiPPSO algorithm for jamming hard spheres.

Require: A HS configuration $(\bar{\mathbf{r}}, \bar{\sigma})$, without overlaps; tolerance for convergence criterion $(\text{tol}_s, \text{tol}_\Gamma)$.

- 1: **procedure** CALiPPSO $(\bar{\mathbf{r}}, \bar{\sigma}; \text{tol}_s, \text{tol}_\Gamma)$
- 2: Compute initial density, φ , and cutoff $\ell(\varphi)$
- 3: **repeat**
- 4: **for** $l = 1, \dots, N$ **do** ▷ Construct neighbor lists
- 5: $\tilde{\delta}l \leftarrow \{j \mid |\mathbf{r}_{lj}| \leq \ell(\varphi)\}$
- 6: **end for**
- 7: $\tilde{\mathbf{F}} := \{F_{ij}(\tilde{\mathbf{s}}, \Gamma) \mid i < j, j \in \tilde{\delta}i\}$ ▷ Define set of relevant constraints, F_{ij} , from Eq. (2b).
- 8: Solve the jamming LOP (2), with constraints $\tilde{\mathbf{F}}$.
- 9: $\tilde{\mathbf{s}} \leftarrow \tilde{\mathbf{s}}^*$
- 10: $\Gamma \leftarrow \Gamma^*$
- 11: Store $\mathcal{C} := \{[ij] \mid \lambda_{ij} > 0\}$ ▷ Define contacts indices from active dual variables
- 12: Store active dual variables, $\underline{\lambda}$.
- 13: $(\bar{\mathbf{r}}, \bar{\sigma}) \leftarrow (\bar{\mathbf{r}} + \tilde{\mathbf{s}}, \sqrt{\Gamma} \bar{\sigma})$ ▷ Update the configuration
- 14: Recompute φ and $\ell(\varphi)$
- 15: **until** $\max_i |s_i| < \text{tol}_s$ **and** $\sqrt{\Gamma} - 1 < \text{tol}_\Gamma$
- 16: $(\bar{\mathbf{r}}_J, \bar{\sigma}_J) \leftarrow (\bar{\mathbf{r}}, \bar{\sigma})$ ▷ Define jammed configuration
- 17: **for** $[ij] \in \mathcal{C}$ **do** ▷ Construct network of contacts
- 18: $\mathbf{n}_{ij} = \frac{\mathbf{r}_{ij}}{\sigma_{ij}}$ ▷ Store contact vectors
- 19: $f_{ij} = \frac{\lambda_{ij}}{\sigma_{ij}}$ ▷ Store forces magnitudes
- 20: **end for**
- 21: **return** $\bar{\mathbf{r}}_J, \bar{\sigma}_J, \{\mathbf{n}_{ij}\}, \{f_{ij}\}$ ▷ Output
- 22: **end procedure**

or dual optimal solution is available, the other one can be accessed straightforwardly.

When the CALiPPSO convergence criterion has been met within a given tolerance $(\text{tol}_s, \text{tol}_\Gamma)$, we directly obtain the jamming packing fraction φ_J , the particles' position and size $(\bar{\mathbf{r}}^{(J)}, \bar{\sigma}^{(J)})$, as well as the set of contact forces $\mathbf{f} = \{f_{ij}\}$. From their knowledge, it is possible to investigate all the properties of the jamming transition of hard spheres, such as the jamming critical exponents [6,11,12,29], and the structure the free-energy landscape (FEL) [26,32] (see also Sec. III A).

In Ref. [60], we provide our own implementation of the CALiPPSO algorithm, written in the Julia programming language [59] and making use of the JuMP [80] modeling package. To solve all the jamming LOP instances, we used the Gurobi Solver [81]. We tested our code also with other free, open-source optimizers such as HiGHS [82] and GLPK [83]. We found that even though all these solvers have a relatively low precision (about 10^{-9}), CALiPPSO is able to produce valid jammed packings systematically. This is a remarkable feature that contrasts with other algorithms, such as FIRE [44], which requires quad-precision computations to avoid over-shooting the jamming critical point [24,46].

The precision of the solver determines the tolerance for satisfying the constraints and, consequently, the precision with which the dual variables are computed. Using Gurobi with the highest overall accuracy available (10^{-9}), overlaps larger than 10^{-8} never occur, and the force balance condition is satisfied with higher precision by several orders of magnitude (about 10^{-13}). Moreover, using such solver, active and nonactive dual variables are distinguished with double precision in most of the systems we examined. This implies that true contacts are identified with high precision and mechanical equilibrium is also guaranteed within a reasonably small tolerance.

Using our implementation [60] of Algorithm 1 we are able to produce packings as the ones shown in Fig. 2. In the top panel, we illustrate the network of contacts in a $2d$ bidisperse configuration, while in the bottom panel spheres are colored according to the number of their contacts (visualization done using Ref. [84]). These packings were crunched from an initial random configuration with $\varphi = 0.4$ (in $2d$) and $\varphi = 0.2$ (in $3d$). We mention that, compared with their corresponding jamming densities, these values amount to $\varphi/\varphi_J < 0.5$ and $\varphi/\varphi_J < 0.3$, respectively. This means that, even if the justification to transform the exact jamming OP (1) into the LOP (2) was based on the assumption that $\varphi \lesssim \varphi_J$, so all $\{|s_i|\}$ are small, the CALiPPSO algorithm is sufficiently robust to produce valid packings even if initialized relatively far from jamming. We tested our algorithm in $d = 2-5$ dimensions (see Appendix B), and we verified that in all cases the jammed packings thus produced satisfy the mechanical equilibrium condition and are isostatic. This shows convincingly that our algorithm generates typical jammed configurations, and can be employed to carefully explore the jamming transition of HS systems.

III. THE LS+CALiPPSO ROUTE TO JAMMING: PROBING THE FREE ENERGY LANDSCAPE

In this section, we present results obtained when combining CALiPPSO with the Lubachevsky–Stillinger (LS) compression protocol [49]. Specifically, the idea is to use configurations compressed with LS as initial conditions of the CALiPPSO algorithm. This scheme will allow us to improve the performance of CALiPPSO and to study in detail the influence of the initial condition on the jammed packings it produces. We argue that the jammed states we obtained following such combined approach, referred as LS+CALiPPSO,




FIG. 2. Upper panel: jammed packing in a two-dimensional, bidisperse system (with diameter ratio 1 : 1.4). Rattlers are coloured in orange and the full network of contacts is shown, with the magnitude of contact forces represented by the thickness of the lines. Bottom panel: monodisperse packing in three dimensions. The contact network is not shown to avoid cluttering but particles are coloured according to the number of their contacts: lighter particles have more contacts, while the darkest ones are rattlers (zero contacts). Both packings were obtained initializing CALiPPSO with low density configurations: $\varphi_0 = 0.4$ in $2d$ ($\varphi_J = 0.839$) and $\varphi_0 = 0.2$ in $3d$ ($\varphi_J = 0.635$).

reflect the hierarchical structure of the free energy landscape (FEL) [6]. Our numerical results suggest that, if CALiPPSO is initialized from a configuration at very high pressure (as specified below), then the packings it produces likely coincide with the ones that would be obtained extrapolating the LS

compression to the infinite pressure limit (see Sec. III B). Hence, using LS+CALiPPSO we can generate jammed packings [5,77,78], reproducing many of the properties observed previously in the literature [6,25,32].

As we mentioned in the Introduction, the LS protocol allows us to compress hard spheres to very high pressures. However, it is not able to strictly reach the jamming condition, $1/p = 0$. Nevertheless, the LS compression protocol is an excellent tool to *approach* the jamming point: Besides being fast, it has been amply verified that it closely reproduces the (phenomenological) equation of state [85] of HS glasses [10,50,86,87]:

$$p = \frac{d}{1 - \varphi/\varphi_J}, \quad (15)$$

with $p = \beta PV/N$ the reduced pressure of the system.

The LS compression protocol increases the particle diameters with a uniform growth rate, $\dot{\sigma}(t) = \kappa$. This compression is performed simultaneously to the dynamical evolution of the configuration. HS dynamics can be efficiently simulated using event-driven MD in arbitrary dimensions [6,25,50,86], or even in some mean-field models [88]. We used the implementation of Ref. [50], so we limit our analysis to finite-dimensional, monodisperse systems. Before proceeding, let us note that depending on κ the LS protocol can produce monodisperse configurations that possess some degree of crystallization [50,51,67]. As stated above, we want to produce only disordered packings; thus, we took care in choosing κ to avoid any ordering in our configurations, as explained in the following.

In this work, we will focus on three-dimensional systems because of their special role as a minimal model displaying jamming phenomenology [89]. Nevertheless, our approach is valid in any dimensionality $d > 1$, (in $d = 1$ the resultant packings are inevitably ordered chains). Monodisperse systems in $d \leq 3$ are prone to crystallize if slow compression rates are used. To avoid the formation of ordered domains in three-dimensional systems, we initially perform a fast compression with $\kappa^{(0)} = 5 \times 10^{-3}$ (see Appendix C for more details); instead, there is no need to include this initial fast compression in higher dimensions.

Our LS+CALiPPSO protocol to reach the jamming point of HS systems is composed of the following steps:

(1) We generate a random HS configuration, i.e., we draw from a uniform distribution the spheres' positions, at low packing fraction φ_0 and without overlaps, and initialize the LS compression protocol with it.

(2) To avoid crystallization, we perform a fast compression, with a compression rate $\kappa^{(0)} = 5 \times 10^{-3}$, until $p_{\text{tar}}^{(0)} = 500$. Note that this step can be safely avoided in $d \geq 4$ (see, e.g., Fig. 8).

(3) We initialize a new LS compression with the HS configuration obtained at $p_{\text{tar}}^{(0)} = 500$, and further compress it until a target pressure $p_{\text{tar}} \gg 1$ (see below for detailed values) is reached. This second compression is performed with a smaller compression rate κ .

(4) We use the HS configuration at p_{tar} from the previous step to initialize the CALiPPSO algorithm. Following the implementation of Algorithm 1, CALiPPSO is executed until the convergence condition $(\Gamma^*, \vec{s}^*) = (1, \vec{0})$ is verified, within a given tolerance. In this way, we obtain an HS configuration

at the jamming point. All the results reported below have been obtained at fixed tolerance: $\text{tol}_s = 10^{-9}$, $\text{tol}_r = 10^{-12}$.

The two protocols involved, LS and CALiPPSO, compress HS configurations at very different speeds. In particular, the LS compression is a finite-time protocol. Thus, while it is unable to produce the states given by a quasistatic construction, such as state-following or adiabatic compression [12,90], it is relatively slow when compared with the CALiPPSO instantaneous inflation of spheres. The CALiPPSO algorithm instead is, for all purposes, a quenched compression, or crunching. Consequently, LS+CALiPPSO is an out-of-equilibrium procedure that brings an HS configuration to its jamming point without following the thermodynamic equation of state (see also Sec. III B).

A. The role of the initial condition and the FEL structure

The most relevant parameters of the LS+CALiPPSO protocol are the LS compression rate, κ (see step 3 above), and the target pressure from which CALiPPSO is initialized, p_{tar} (see step 4 above). In addition, we can change the number of particles in the system, N . Notice that, due to memory effects in hard spheres [48], in principle the value φ_{tar} from which CALiPPSO is initialized, i.e., the packing fraction such that $p(\varphi_{\text{tar}}) = p_{\text{tar}}$, should be considered as a parameter of the protocol. However, for simplicity we will assume that, in the glassy regime, the exact value of φ is uniquely determined by p , and the influence of φ_{tar} is considered to be implicitly captured by p_{tar} . Therefore, in the following, we will investigate only the influence of κ , p_{tar} , and N on the final jammed states.

For fixed κ and N , we construct a *sample* by compressing via LS the same initial HS configuration to different values of p_{tar} in the range $[10^3, 10^{11}]$. In such a way, all the jammed states of a sample belong to the same glassy state. This feature can be exploited to explore the FEL's structure at jamming. Since our compression protocol stops as soon as $p \geq p_{\text{tar}}$, the precise value of p obtained for a target p_{tar} can slightly vary. For a given p_{tar} , there is a small sample variability (usually less than 1%) in the actual value of the pressure.

Let us begin by analyzing the dependence of φ_J on the LS+CALiPPSO parameters. In particular, let us consider its sample average, $\langle\varphi_J\rangle$. All the averages reported in the following are computed over 20 samples, unless otherwise stated. We consider $p_{\text{tar}} \in [10^3, 10^{11}]$, and several values of κ and N . Figure 3(a) shows the dependence of $\langle\varphi_J\rangle$ on p_{tar} , for various κ , at $N = 1024$. Figure 3(b) illustrates analogous results for different system sizes, at fixed $\kappa = 10^{-5}$. In agreement with previous studies [3,5,10], we find that $\langle\varphi_J\rangle$ is always close to 0.64. We see that changing the target pressure of the LS compression protocol has a very small effect in the pressure range we considered. In contrast, slowing down the compression or increasing the system size increases φ_J , albeit always within a narrow interval. These features are not surprising. Lowering κ , the particles have more time to rearrange; therefore, the system can reach deeper minima of the FEL, yielding higher values of φ_J [32]. Instead, border or periodic effects are reduced in larger systems, so the particles are less constrained, and higher values of φ_J can be achieved. Nevertheless, the data suggest a rather quick convergence to the thermodynamic limit value.

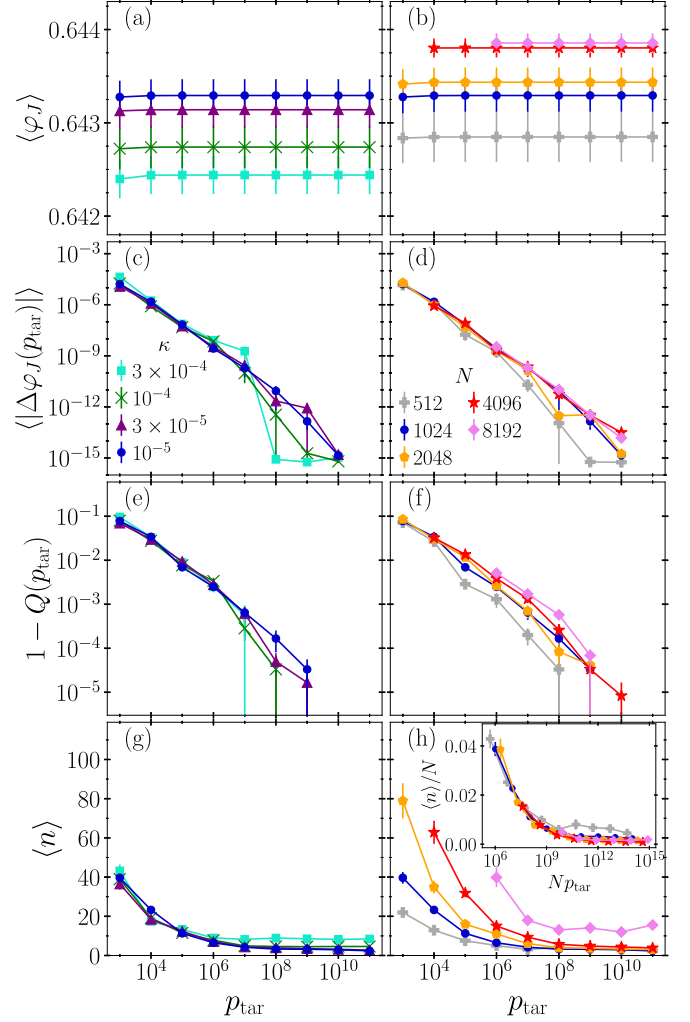


FIG. 3. Dependence of various properties of the jamming configurations obtained via LS+CALiPPSO as a function of p_{tar} , for $N = 1024$ (left) and $\kappa = 10^{-5}$ (right); note that plots on a same row share the scale of their vertical axes. We present results for several values of the compression rate κ (left column) and the system size N (right column). All values reported are the average and standard error over 20 samples. (a, b) Jamming packing fraction, $\langle\varphi_J\rangle$. (c, d) Difference of the jamming packing fraction obtained at p_{tar} from the one obtained at $p_{\text{max}} = 10^{11}$ in the same sample. (e, f) Similarity measure of the contact networks within the same sample, $1 - Q(p_{\text{tar}})$ (defined in the main text). (g, h) Number of linear optimizations, n , needed for CALiPPSO to reach the jamming convergence condition. Inset of panel (h): size scaling collapse showing that n/N is only a function of the thermodynamic pressure $P \sim Np$, as explained in the main text. Panels (c–f) support the presence of a hierarchical structure of the jamming landscape, as explained in the text.

To further investigate the properties of the jamming configurations obtained via the LS+CALiPPSO algorithm, we define the quantity $|\Delta\varphi_J(p_{\text{tar}})| := |\varphi_J(p_{\text{max}}) - \varphi_J(p_{\text{tar}})|$, where $p_{\text{max}} = 10^{11}$ is the largest target pressure considered, and is used as a reference for comparison. $|\Delta\varphi_J|$ quantifies to what extent the packing fraction of a jammed configuration changes if the CALiPPSO crunching begins from a smaller p_{tar} . We seldom find that $\Delta\varphi_J$ is a nonmonotonic function

of p_{tar} , so in Figs. 3(c) and 3(d) we consider its absolute value. We report $\langle |\Delta\varphi_J(p_{\text{tar}})| \rangle$, averaged over 20 samples, for several values of κ at $N = 1024$ in Fig. 3(c), and for different system sizes N at $\kappa = 10^{-5}$ in Fig. 3(d). We observe that in general $\langle |\Delta\varphi_J| \rangle$ is a decreasing function of p_{tar} . Moreover, from Fig. 3(c) we see that $\langle |\Delta\varphi_J(p_{\text{tar}})| \rangle$ is independent of the compression rate only for $p_{\text{tar}} \leq 10^6$. For larger pressures, the effects of κ become relevant. For instance, using the fastest compression ($\kappa = 3 \times 10^{-4}$), $\langle |\Delta\varphi_J(p_{\text{tar}})| \rangle = 0$ within the numerical precision for $p_{\text{tar}} \geq 10^8$. This means that, in each sample, the LS+CALiPPSO algorithm finds the same φ_J . Instead, smaller compression rates yield $\langle |\Delta\varphi_J(p_{\text{tar}})| \rangle \geq 0$ for higher pressure, with $\Delta\varphi_J$ remaining finite within an interval that grows as κ decreases. A cleaner signature of this behavior is discussed below considering the similarity of the network of contacts [see Figs. 3(e) and 3(f)]. On the other hand, Fig. 3(d) shows that size effects are negligible, at least for $N \geq 1024$.

As explained before, employing the LS+CALiPPSO protocol, we can easily extract the network of contacts of a jammed configuration. Notice that this network univocally defines a minimum of the FEL [6]. Thus, a more refined measure of similarity between jammed configurations can be obtained comparing the contact networks rather than the packing fractions of two configurations. In particular, we are interested in quantifying the similarity between two configurations of the same sample at different target pressures. Considering again the configuration produced from p_{max} as reference, we define $Q(p_{\text{tar}}) := |\mathcal{C}(p_{\text{tar}}) \cap \mathcal{C}(p_{\text{max}})|/N_c(p_{\text{max}})$, where $\mathcal{C}(p_{\text{tar}}) \cap \mathcal{C}(p_{\text{max}})$ is the intersection between the contact networks of the two configurations, while $N_c(p_{\text{max}})$ is the number of contacts of the jammed configuration at p_{max} . Therefore, $Q(p_{\text{tar}})$ measures the number of common contacts between the packings obtained initializing CALiPPSO with two configurations in the same sample, one at p_{tar} and the other at p_{max} . $Q(p_{\text{tar}}) = 1$ only when the contact networks exactly coincide. Similar observables have been used elsewhere [6,26,48] to investigate the FEL's structure near and at jamming. Our findings are reported in Figs. 3(e) [respectively, Fig. 3(f)] for fixed N and different κ (respectively, fixed κ and different N). For the values of p_{tar} , κ , and N considered here, we see that the main structure of the contact network is shared by jammed packings in the same sample, i.e., within a given metabasin [6,12]. Yet, a small fraction of the contacts ($\sim 10\%$) are only determined gradually, as the CALiPPSO input configuration goes down in the landscape as p_{tar} increases. Thus, the behavior of $1 - Q(p_{\text{tar}})$ reported in these figures is in accordance with that of $|\Delta\varphi_J(p_{\text{tar}})|$ discussed above. Using $Q(p_{\text{tar}})$ the influence of changing the compression speed, as well as the system size, are distinguished more cleanly. We observe that for a given κ or small N , there is a threshold pressure above which we obtain $Q(p_{\text{tar}}) = 1$ for all samples.

These findings are in agreement with the rough and hierarchical structure of the FEL predicted for hard spheres in infinite dimensions (see Sec. I). Assuming such mean-field picture to be valid also in finite-dimensional systems, as suggested by recent numerical evidence [6,12,26,32], our results indicate that the FEL's structure could only be fully resolved with an infinitesimally slow compression, $\kappa \rightarrow 0$ (assuming that crystallization can be avoided in this limit; see the dis-

cussion at the beginning of Sec. III B). Instead, when κ is finite, there will be a threshold pressure $p^{(\text{th})}$, such that if $p_{\text{tar}} > p^{(\text{th})}$ the jamming packings obtained within the same sample will be inevitably the same: the system is trapped in one minimum. The fact that $p^{(\text{th})}$ increases with N can be understood considering that the number of minima in the FEL increases tremendously with the system size [12,13,91] which suggests that for large systems finding diverse minima is much likelier, even for very high values of p_{tar} . The role of $p^{(\text{th})}$ will be discussed in more detail in the next section.

The dependence of the final packing on p_{tar} is reflected not only in physical quantities, such as φ_J and the network of contacts but also in algorithmic properties. In Figs. 3(g) and 3(h) we consider the number of linear optimizations needed to reach the convergence criterion, n , as a function of p_{tar} . In agreement with the rest of the panels, these curves illustrate that, for sufficiently large p_{tar} , the value of n remains essentially unchanged, despite the fact that the distance to φ_J decreases by several orders of magnitude. Moreover, from Fig. 3(h) we see a clear dependence of n on the system size N . Interestingly, such dependence can be teased out by assuming $n = N\mathcal{G}(Np_{\text{tar}})$, where $\mathcal{G}(x)$ is a scaling function such that $\mathcal{G}(x) \sim \text{constant}$, for $x \rightarrow \infty$. We put to test this scaling in the inset of Fig. 3(h), obtaining a very good collapse of the curves at different N . The small deviations observed in $N = 512$ systems at very large target pressure are likely explained from the fact that, in such cases, n matches the minimal number of iterations, $n_0 = 2$ (see Sec. IV). We can rationalize the scaling variable Np_{tar} by noting that, from the definition of the reduce pressure [92], we have that $Np = \beta PV$, where P is the usual thermodynamic pressure, β is the inverse temperature, and V is the system's volume. Since in all our MD simulations we fix $\beta = 10$ and $V = 1$, $Np_{\text{tar}} \sim P$ provides the natural variable to measure how far from jamming a configuration is, independently of the number of particles. Our results hence indicate that $n = N\mathcal{G}(P)$; for fixed P , this linear relation suggests that CALiPPSO works by blocking few degrees of freedom at each linear optimization (see also the discussion in Sec. IV).

Finally, we compare the properties of the packings obtained via the LS+CALiPPSO protocol to the ones that would be obtained using only the LS compression protocol in the limit $p \rightarrow \infty$. Unfortunately, it is very difficult to access the microscopic details of the jammed packings, such as contacts and interparticle gaps, purely from an extrapolation of the data at finite pressure (see, for instance, Fig. 10 in Appendix D and the corresponding discussion). However, recalling that an HS glass in d dimensions is well described empirically by the equation of state (15), we can use this expression to fit the values of p and φ computed during the MD simulations of the LS protocol, and then estimate the jamming packing fraction $\varphi_J^{(\text{LS})}$ that would correspond to $1/p = 0$.

Fixing $N = 1024$, we estimate $\varphi_J^{(\text{LS})}$ using the LS data in the range $p \in [10^5, 10^{11}]$, and we perform the sample average of such values. Then, we compare it to the sample average of φ_J obtained via the LS+CALiPPSO protocol. The results (mean and standard error) for various compression rates κ are reported in Table I. We only include the LS+CALiPPSO results at $p_{\text{tar}} = 10^3$ and 10^5 because with larger pressures the sample average of φ_J is identical, within the statistical error, to the one at $p_{\text{tar}} = 10^5$. The results in Table I imply that, at

TABLE I. Average values of the jamming density of the LS compression protocol (second column) and LS+CALiPPSO (third and fourth columns). Each row corresponds to a different growth rate (first column). $\varphi_J^{(LS)}$ is the $p \rightarrow \infty$ limit packing fraction within the LS protocol, estimated via a fit as explained in the main text. The values and uncertainties reported are obtained averaging over 20 samples of $N = 1024$ particles.

κ	$\langle \varphi_J^{(LS)} \rangle$	$\langle \varphi_J(p_{\text{tar}} = 10^3) \rangle$	$\langle \varphi_J(p_{\text{tar}} = 10^5) \rangle$
3×10^{-4}	0.64244(20)	0.64240(21)	0.64244(20)
10^{-4}	0.64274(23)	0.64272(22)	0.64274(23)
3×10^{-5}	0.64314(19)	0.64313(19)	0.64314(19)
10^{-5}	0.64329(18)	0.64328(17)	0.64329(18)

fixed κ , the jammed state obtained via the LS+CALiPPSO protocol for sufficiently large p_{tar} (presumably smaller than $p^{(\text{th})}$) coincides with $\varphi_J^{(LS)}$. Similar results have been obtained performing a comparison of φ_J and $\varphi_J^{(LS)}$ within each sample. However, given that for $p_{\text{tar}} \geq 10^5$ the uncertainty in the estimation of $\varphi_J^{(LS)}$ (about 10^{-5}) is much larger than the values of $\varphi_J - \varphi_J^{(LS)}$ (about 10^{-7}) and $|\Delta\varphi(p_{\text{tar}})|$, a more precise comparison for larger p_{tar} is unfeasible.

B. The LS+CALiPPSO phase diagram

In view of the results presented in Sec. III A, we can sketch the path followed by HS configurations brought to jamming via the LS+CALiPPSO algorithm. Such a path is illustrated in Fig. 4 in the $(\varphi, 1/p)$ plane, where it is compared with the path one would obtain considering an infinitesimally slow compression (black solid line), i.e., the hypothetical thermodynamic equation of state of the glass. Let us stress that in practice no algorithm can reach the $\kappa \rightarrow 0$ limit. Moreover, in $3d$, monodisperse systems decreasing κ below a critical compression rate inevitably leads to partial crystallization [52] (accompanied by significantly larger values of φ_J). Since our results have been obtained at finite κ , and we have checked the absence of any sign of crystallization (see Fig. 7), in our discussion we will neglect the possibility of forming a crystal during the LS compression, and use the $\kappa \rightarrow 0$ limit just as an idealized reference.

In Fig. 4, the first half of our protocol, i.e., the LS compression, is identified by the $\kappa > 0$ line which, for small p , is virtually identical to the thermodynamic one. However, for any finite $\kappa > 0$, the glass obtained from the LS protocol inevitably detaches from the path associated with the quasistatic limit. The CALiPPSO crunching is identified by the dashed, red lines, highlighting the fact that no value of p can be assigned during this process. Figure 4 illustrates also the threshold pressure introduced before, such that, if $p_{\text{tar}} < p^{(\text{th})}$, CALiPPSO produces packings with different jamming density and network of contacts. Conversely, when $p_{\text{tar}} > p^{(\text{th})}$, CALiPPSO converges to the same jammed microstate. As discussed above, such a state is presumably identical to the $p \rightarrow \infty$ limit of the LS protocol. From the FEL perspective, this implies that in the regime $p_{\text{tar}} < p^{(\text{th})}$ the hierarchical structure of the landscape affects the final packings realized by the LS+CALiPPSO algorithm, while for $p_{\text{tar}} > p^{(\text{th})}$ such




FIG. 4. Sketch of LS+CALiPPSO route to jamming in the $(\varphi, 1/p)$ plane. We include a schematic comparison of the LS protocol (blue line, $\kappa > 0$) and the quasistatic one (black line, $\kappa \rightarrow 0$) obtained by, e.g., adiabatic compression, that ends in φ_J^* . These two protocols coincide up to a certain pressure, identified by the gray square, but at larger pressure the finite compression rate of LS makes it detach from the idealized thermodynamic line. $\varphi_J^{(LS)}$ is the jamming packing fraction that would be obtained by LS when $1/p_{\text{tar}} = 0$, and is identified with the big blue circle to represent the large uncertainty in its estimation with respect to φ_J from CALiPPSO (see Sec. III A). CALiPPSO crunching is represented by dashed lines, initialized at a given $p_{\text{tar}} < \infty$. If $p_{\text{tar}} \lesssim p^{(\text{th})}$, then CALiPPSO leads to packings of different densities, e.g., $\varphi_J^{(1)} < \varphi_J^{(2)}$ (green lines). Conversely, when $p_{\text{tar}} \gtrsim p^{(\text{th})}$, different initial conditions lead to the same jammed state, $\varphi_J^{(3)}$ (red lines) that very likely coincides with $\varphi_J^{(LS)}$. Note that the interval of different values of φ_J has been drastically magnified, and that the $\kappa \rightarrow 0$ path is used only as a schematic reference since no finite-time numerical algorithm could follow such a path.

structure is not longer detected by the algorithm and it always reaches the same minimum. In addition, our data show that $p^{(\text{th})}$ is not universal: $p^{(\text{th})}$ increases as κ decreases or N increases. Nevertheless, further studies are needed to better characterize such dependence. Finally, we note that the values of φ_J obtained using either CALiPPSO or the $p \rightarrow \infty$ limit of LS are expected to be smaller than the quasistatic jamming density φ_J^* . This discrepancy has already been observed in polydisperse systems [93], and mean-field models [94], although in $3d$ monodisperse HS systems such difference should be very small [25].

IV. TIME COMPLEXITY OF CALiPPSO

In this section, we explore the performance of the CALiPPSO algorithm as a function of the system size N . In particular, we will analyze the time required by the algorithm to converge (τ), the number of linear optimizations (n), and the mean time required per linear optimization in a given system (t_{LOP}), as a function of N . The results presented in this section are obtained for $3d$, monodisperse systems, fixing $p_{\text{tar}} = 10^7$, and within two different LS compression protocols. In the first one, we set $\kappa_f = 10^{-5}$ for all system sizes; in the second one, we introduce a size-dependent growth rate,

κ_N , such that $\dot{\varphi}$ is constant for all values of N . More precisely, given that $\dot{\varphi} \sim N\sigma^2\dot{\sigma} \sim (N\varphi^2)^{1/3}\kappa$, if all configurations must be subject to the same compression rate for a given value of φ , then $\kappa_N \sim N^{-1/3}$ (more generally, in d dimensions we would have $\kappa_N \sim N^{-1/d}$). We consider system sizes up to $N_{\max} = 16384$. For the second compression protocol, we fix $\kappa_N = (N_{\max}/N)^{1/3}\kappa_f$. Notice that, since the smallest system size is $N_{\min} = 256$, the ratio between the inflation rates of the two protocols is $\kappa_{N_{\min}}/\kappa_f = 4$ at most. Comparing the two scenarios is useful to guarantee that the scalings we obtain are intrinsic to CALiPPSO.

Since the performance of the LS protocol has been analyzed before [62], we will not discuss it here. However, in Appendix D we show a comparison of the CALiPPSO and LS times and their dependence on p_{tar} (see Fig. 11).

We test our implementation of the CALiPPSO algorithm in a 6 cores computer, with processor Intel Core i7-8700 at 3.2 GHz. Each jamming LOP is solved using the Gurobi Solver [81] (version 9.1) along with the JuMP package [80] of the Julia programming language [59]. All the jamming LOP instances are solved using the Gurobi's barrier method, with six threads running concurrently, and setting the feasibility and optimality tolerance to their most stringent values, 10^{-9} . The rest of the solver parameters are used with their default values [81].

In Fig. 5, we report the results obtained after averaging over 100 samples for each value of N and compression protocol, along with the corresponding standard error. Notice that in all three panels, the mean values obtained at fixed compression, κ_f (blue crosses), are very similar to the ones obtained with the size-scaled compression rate, κ_N (green squares). Therefore, here we only present the data analysis made within the former protocol. In Fig. 5(a), we illustrate the size dependence of τ , as well as the result of a least-square fit of the form $\tau = \tau_0 + cN^\alpha$ (black line). We chose the offset time, τ_0 , such that residuals of the linear model $\ln(\tau - \tau_0) = \ln c + \alpha \ln N$ were minimized, and obtained $\alpha = 2.84 \pm 0.09$. We can see that data at small N deviate from the large size trend. Performing again the fit, but including only data for $N \geq 1024$, we find that $\alpha = 3.09 \pm 0.13$ (red, dashed line). Similar values are obtained for the compression protocol with κ_N , whence we argue that this is the intrinsic size dependence of τ of CALiPPSO. We can conclude that, for high enough p_{tar} (see Sec. III B), the complexity of CALiPPSO scales approximately as N^3 .

We can also investigate which components of the CALiPPSO algorithm are more sensitive to the system size. In other words, does τ increase with N due to a much larger number of LOP instances required for convergence? Or, instead, is it the time spent for each linear optimization that gives the largest contribution to τ ? To answer this questions, in Fig. 5(b) [respectively, Fig. 5(c)] we plot the behavior of $\langle n \rangle$ (respectively, $\langle t_{\text{LOP}} \rangle$) for different values of N and the two protocols considered. Our results clearly show that the increasing computational cost for jamming larger systems is due to the time required to solve each LOP instance. Indeed, analogous fits to Fig. 5(a) yield that $t_{\text{LOP}} \sim N^{2.34 \pm 0.05}$, while $n \sim N^{1.04 \pm 0.03}$. Hence, the main contribution to the growth of τ with N comes from the increasing difficulty in solving a single instance of the jamming LOP, while the number of steps required for convergence grows only linearly with the

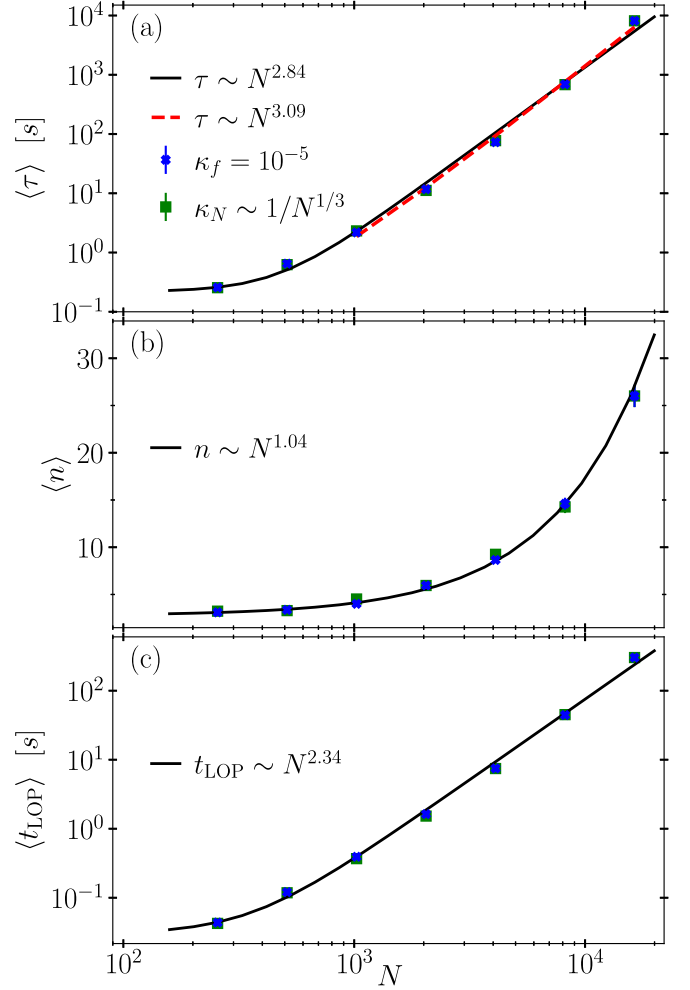


FIG. 5. Complexity of CALiPPSO as a function of system size N , measured by the convergence time, τ (a), number of linear optimizations used to reach the convergence condition, n (b), and the mean time required in solving the jamming LOPs of a system, t_{LOP} (c). We set $d = 3$, $p_{\text{tar}} = 10^7$ and used two LS compression protocols: a fixed growth rate $\kappa_f = 10^{-5}$ (blue crosses) and a size scaled compression κ_N (green squares). Data reported correspond to the average over $M = 100$ samples for each value of N and compression strategy. Solid lines show the size scaling obtained by least-squares fits to the data obtained using κ_f ; the red, dashed line in panel (a) is an analogous fit obtained considering only $N \geq 1024$ (see main text for more details). The fits obtained with the κ_N protocol are virtually identical.

system size. Notice that the scaling exponent obtained for $\langle t_{\text{LOP}} \rangle$ versus N is significantly smaller than the one from the worst-case scenario analysis that would yield a dependence $N^{3.5}$, given the complexity of the interior point method itself [71].

Our data show also that the minimal number of iterations required by CALiPPSO is approximately $n_0 = 2$, a value confirmed by analogous results for very large p_{tar} [cf. Fig. 3(b)]. This finding can be intuitively explained by assuming that, even if the initial condition is very close to jamming and a single linear optimization would make a system reach its final density, \vec{s}^* would be determined by saturating the *linear*

constraints. Therefore, an extra iteration would be needed to make linear contacts precisely match physical ones.

It is worth recalling here that from the scaling presented in Sec. III A, $\langle n \rangle \sim \mathcal{N}\mathcal{G}(P)$, the linear dependence of the number of iterations with N is expected to be valid even for smaller values of p_{tar} . In other words, even if CALiPPSO was initialized from a smaller pressure, the proportionality relation $\langle n \rangle \sim N$ would be valid for a fixed P .

Importantly, our tests have also shown that the convergence condition for Γ is achieved noticeably faster than the analogous condition on \bar{s} , especially for $N > 1024$. That is, the last iterations of CALiPPSO are employed to fine-tune the particles' positions to attain stability (by matching linear with real contacts), and not in increasing the system density. This suggests that our algorithm might be improved by implementing some sort of "relaxation" during the configuration update. For instance, writing the optimal inflation factor of the jamming LOP as $\Gamma^* = 1 + \gamma$, with $\gamma > 0$, we can update the particles' diameter as $\sigma \rightarrow \sqrt{1 + c\gamma}\sigma$, for some $0 < c < 1$, instead of the rule considered so far. By doing so, it is likely that both Γ and \bar{s} would converge at a much similar rate because the slightly larger amount of free volume would allow obtaining larger optimal displacements when needed. This goes beyond the scope of the present work, and such alternatives will be explored in further studies.

Another relevant point emerging from Fig. 5 is that the simple assumption $\langle \tau \rangle \sim \langle n \rangle \cdot \langle t_{\text{LOP}} \rangle$ would give a slightly larger exponent for the scaling of τ with N than the one obtained from the direct fit. There are two mechanisms that might be responsible for this mismatch. The first and most obvious one is the contribution of the other steps described in Algorithm 1. Clearly, besides the linear optimization, CALiPPSO also requires computing and updating the neighbor lists, identifying rattlers, modifying and storing arrays, etc. Because each of these operations is less demanding than the jamming LOP, their combined effect is to reduce the overall exponent. However, a second factor that should be considered is that t_{LOP} and n are not necessarily independent variables, thus $\langle n \rangle \cdot \langle t_{\text{LOP}} \rangle \neq \langle nt_{\text{LOP}} \rangle \sim \langle \tau \rangle$. The fact that the first of these quantities yields an N dependence with a larger exponent implies that n and t_{LOP} are anti-correlated. In other words, it is likely that our algorithm converges by performing a relatively small number of expensive linear optimizations.

As we mentioned in Sec. II, the CALiPPSO algorithm is suitable for systems in any dimensions d , as we evince in Fig. 8 in Appendix B for $d = 4$ and $d = 5$. Here, we anticipate that its performance is affected by the dimensionality of the system because, as d increases, each particle is surrounded by more neighbors. That is, even if at jamming each particle is in contact with an average of $2d$ other spheres, the amount of spheres within $\ell(\varphi)$ that induce a constraint could grow much faster. More precisely, solving each jamming LOP has a polynomial complexity [71,95] on the $\max(N_{\text{dof}}, M')$, following the notation of Sec. II. Now, $N_{\text{dof}} \sim dN$, while $M' = \bar{z}_d N$, where \bar{z}_d denotes the average number of near contacts in a d -dimensional system. Isostaticity and geometrical constraints imply $N_{\text{dof}} \leq M' \leq k_d N$, with k_d the kissing number in d dimensions. Thus, as long as k_d is not much larger than d , we can expect our complexity analysis to hold. Unfortunately, k_d increases exponentially in d , and having tighter

bounds on \bar{z}_d is not trivial. Previous data in $d = 4-6$ [3,50] suggest that the abundance of near contacts increases rapidly. Nevertheless, the constraint matrix \mathcal{F} of the jamming LOP (3) will remain rather sparse as long as $\bar{z}_d \ll N$; this should avoid t_{LOP} from reaching the worst-case complexity mentioned above. Finally, let us notice that \bar{z}_d is also influenced by the CALiPPSO initial configuration parameters (e.g., p_{tar} if the LS compression protocol is used). Nonetheless, our tests have shown that CALiPPSO is an efficient algorithm even for moderately high values of d , as discussed in Appendix B. A more quantitative analysis is beyond the scope of the present work, and is left for future studies.

V. CONCLUSIONS

In this work, we introduced the CALiPPSO algorithm to produce disordered jammed packings of hard spheres with very high accuracy. In contrast with most of the existing algorithms, CALiPPSO does not require introducing any effective potentials between particles. Instead, it is based on formulating the packing problem of hard spheres as a non-convex optimization problem, which is then solved through a series of more tractable linear optimization problems. Section II contains our main results. We showed that, even if the linear problems are only approximations of the original problem, once convergence is attained, CALiPPSO produces hard-sphere (HS) configurations that are optimal also with respect to the original nonconvex problem. Importantly, using results of optimization theory, we analytically proved that CALiPPSO packings are globally stable, in mechanical equilibrium, and isostatic. From the analysis of the complexity of CALiPPSO in $3d$ systems in Sec. IV, we showed that its convergence time likely scales as N^3 , where N is the system size, making CALiPPSO a very efficient algorithm. We provide our own implementation of CALiPPSO in Ref. [60].

Notably, from the isostatic packings produced by the CALiPPSO algorithm, one can easily extract all the relevant information on the microstructure of the jammed HS configurations. Achieving the same precision with techniques based on molecular dynamics (MD) is certainly not as straightforward. When employing MD-based algorithms few gaps might be misclassified as contacts, potentially leading to unstable packings. A detailed discussion on this is provided in Appendix D. CALiPPSO solves these issues thanks to the fact that the complete microstructure of a jammed packing is obtained purely from static quantities. In fact, contact forces are identified from active Lagrange multipliers. Thus, forces are computed independently from interparticle gaps.

To improve the performance of CALiPPSO, in Sec. III we combined it with the Lubachevsky-Stillinger (LS) compression protocol and verified that using these two algorithms together we can readily produce typical jammed packings. By means of extensive numerical simulations, we showed that the LS+CALiPPSO protocol is capable of probing the hierarchical structure of the free-energy landscape with unprecedented accuracy. Studying the landscape structure in finite-dimensional HS systems is relevant for several reasons; for example, it provides a direct test of the recent mean-field theory of glasses and jammed systems [6,12,13]. LS+CALiPPSO represents an optimal candidate to accurately

test whether the landscape of HS configurations possesses the ultrametric structure predicted by the theory, and recently confirmed in soft-sphere packings [26].

Our method opens up numerous research directions towards the characterization of the jamming critical properties in finite-dimensional HS systems. First, exploring how the LS+CALiPPSO protocol navigates the landscape upon reaching jamming could inform on the existence of the Gardner phase [12,14,96] in finite-dimensional models. Moreover, our algorithm can be profitably used to confirm whether the Gardner-like algorithmic transition found in Ref. [48] is a generic feature of packing algorithms. Furthermore, by defining a cost function for hard spheres at jamming, CALiPPSO can be employed to explore the stability of the jammed packings, as well as their spectral properties. Finally, another promising extension would be to adapt this algorithm to tackle constraint satisfaction problems. For instance, one could apply it to the spherical perceptron model [15,16] in finite dimensions and study the properties of the SAT/UNSAT transition when it is approached from the SAT phase. This would provide a numerical validation of the corresponding mean-field results [16,57]. We defer the study of these and other topics to future works.

ACKNOWLEDGMENTS

We thank Edan Lerner for his very valuable help to understand other algorithms dealing with hard-sphere packings, and two anonymous referees whose comments helped us to improve the manuscript. This work was supported by the Simons Foundation [Grant No. 454949 (G.P.)]. C.A. acknowledges financial support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant No. 101001902).

APPENDIX A: OVERCONSTRAINED SYSTEMS: THE CASE OF $2d$ MONODISPERSE PACKINGS

As we showed in Sec. II B, our protocol produces configurations that satisfy the stability condition, $N_c \geq N_{\text{dof}}$. Nevertheless, even though the vast majority of packings produced by CALiPPSO are isostatic, in the case of *monodisperse* disks additional topological and geometrical constraints lead to hyperstatic configurations (i.e., the strict inequality is verified). As we explain here, in such overconstrained packings the excess of contacts only occur between pairs of rattlers.

When dealing with monodisperse configurations in $2d$, CALiPPSO constructs a solution that is mathematically valid [i.e., (\vec{s}^*, Γ^*) optimize the jamming LOP, and λ^* fulfills the force balance conditions] but *physically unstable*. The reason is that some rattlers have two parallel force bearing contacts. Hence, such particles could move in the direction perpendicular to the line defined by the contacts, without affecting any mechanical constraint. Moreover, force bearing rattlers are part of the backbone of the system and, if they were removed, mechanical equilibrium would be broken across the system. We observed that such instabilities are invariably accompanied by the formation of large crystalline domains. These features are exemplified in Fig. 6.




FIG. 6. Unstable, jammed packing of $N = 1024$ monodisperse hard disks. The presence of large crystalline domains is apparent. Particles coloured in orange are rattlers upon which no contact force is acting, while green disks are rattlers with finite contact forces. Particles of this latter type cannot be removed without breaking the force balance condition.

This atypical case can be understood by first noting that in $2d$ monodisperse systems partial crystallization is practically unavoidable. This is due to the fact that the Euler criterion (applied to planar graphs), and the requirement that the faces of such a graph are regular polygons add extra constraints to the network of contacts in monodisperse disks [74,75]. Now, in a region with partial crystallization, the average coordination number is usually larger than the isostaticity requirement, $z = 2d - \mathcal{O}(1/N)$. For instance, in a triangular lattice $z = 6$. However, as we have showed in Sec. II B, the CALiPPSO algorithm always generates isostatic packings with respect to the stable particles. To reconcile these two opposing conditions, in $2d$ monodisperse systems CALiPPSO produces packings where some rattlers exert finite forces upon stable particles. From the point of view of linear optimization, the excess of active constraints causes the jamming dual LOP (5) to have degenerate solutions [70].

Formally, such configurations are hyperstatic since there are more contacts than the number of degrees of freedom. Nevertheless, if the contacts exerted by rattlers on stable particles are taken as “external” forces acting individually on such stable particles, isostaticity is recovered. In other words, if N_c exclusively counts contacts between stable-stable and rattler-stable particles, and all rattlers (even the ones with force bearing contacts) are excluded from N_{dof} , then the isostatic condition $N_c = N_{\text{dof}}$ is once again verified. We stress that, even in this uncommon scenario, the force balance condition is always satisfied for both rattlers and stable particles.

From these considerations, we argue that whenever the solution to the jamming LOP leads to finite forces between rattlers and stable particles, the resultant system must be




FIG. 7. Radial distribution function of jammed packings obtained initializing CALiPPSO from different target pressures (displaced vertically for clarity), as indicated in the legend, using $\kappa = 3 \times 10^{-4}$ for the LS compression. Each curve is the average over $M = 20$ configurations of $N = 1024$ spheres. Changing the compression rate does not alter $g(r)$, as illustrated by the black lines corresponding to $\kappa = 10^{-4}$ (dashed), $\kappa = 3 \times 10^{-5}$ (dash-dotted), and $\kappa = 10^{-5}$ (dotted). Inset: We show, for comparison, $g(r)$ for 10 jammed configurations with partial crystallization (gray lines), and their average (cyan curve). Such kind of configurations have been excluded from all the results presented in this study. The peaks at $\sqrt{2}$ and $\sqrt{7}$ (vertical dashed lines) are not present in the amorphous packings, confirming the absence of crystallization. Besides, the ones at $\sqrt{3}$ and 2 are considerably smaller and broadened.

hyperstatic. Such a hyperstaticity in CALiPPSO packings signals the presence of some structural ordering.

APPENDIX B: FURTHER CHARACTERIZATION OF CALiPPSO JAMMED PACKINGS

First, we show that our LS+CALiPPSO protocol produces packings without any crystallization. To do so, we compute the radial distribution function $g(r)$ of the jammed packings with $N = 1024$ monodisperse particles in $3d$. The curves we obtained initializing CALiPPSO from different values of p_{tar} are reported in Fig. 7 and are fully consistent with the analogous results obtained with other methods [50,97]. In particular, the absence of a peak at $r = \sqrt{2}\sigma_J$ indicates that no crystalline order is present in our packings. To exemplify that even a small degree of crystallization yields a noticeable peak at such location, in the inset of the same figure, we plot $g(r)$ of 10 slightly denser configurations, obtained using a smaller $\kappa^{(0)}$ during the LS compression which have partially crystallized. The densest of these 10 extra configurations has a φ_J that is only 5% larger than the average of the packings without crystallization. In addition, we see that the very sharp peaks at $r = \sqrt{3}\sigma_J$ and $r = 2\sigma_J$ broaden considerably in the disordered configurations, while the one at $r = \sqrt{7}\sigma_J$ disappears and only the underlying shoulder remains.

Finally, we characterize the jammed packings obtained via the LS+CALiPPSO protocol in $d > 3$. In particular, we analyze monodisperse configurations of $N = 1024$ particles in $d = 4$ and $d = 5$, for $\kappa = 10^{-5}$. As mentioned at the




FIG. 8. Results with the LS+CALiPPSO protocol in higher dimensions: $d = 4$ (blue circles) and $d = 5$ (green squares). In analogy with Fig. 3, we explore the effect of p_{tar} on the jamming packing fraction [$d = 4$ in panel (a) and $d = 5$ in panel (b)], as well as $\Delta\varphi_J(p_{\text{tar}})$ (c), and the number of linear optimizations, n (d). The values reported correspond to the average over 20 samples together with the corresponding standard error.

beginning of Sec. III, when producing these packings, the fast compression with $\kappa^{(0)}$ (point 2 of our LS+CALiPPSO protocol) has been omitted. In analogy with the analysis of Sec. III A, we compute φ_J , $\Delta\varphi_J(p_{\text{tar}})$, and n when CALiPPSO is initialized from different values of the target pressure. In Fig. 8, we report the average over 20 samples of these quantities, which agree with the behavior found in $d = 3$ systems and the hierarchical FEL structure we described above. We verified that all the obtained configurations at jamming are isostatic and in mechanical equilibrium. These results demonstrate that CALiPPSO can readily produce valid jammed packings in higher dimensions.

APPENDIX C: FURTHER DETAILS ON THE LS COMPRESSION PROTOCOL

In Fig. 9, we exemplify the LS compression part of our LS+CALiPPSO algorithm (i.e., the first three steps described at the beginning of Sec. III). From this Figure it is clear that using the fast compression the liquid (crosses) undergoes a smooth transition to a glass (circles); thus, there are no signatures of the crystalline phase. However, it should be noted that having a finite κ , the LS fails to equilibrate the liquid all the way up to the glass transition density [10,92], $\varphi \simeq 0.58$. Yet, the glass phase is well described by the free-volume equation of state, as shown by the excellent agreement between the numerical data and the plot of Eq. (15) (solid red line). As mentioned above, this is not the true thermodynamic equation of state. The inset of Fig. 9 confirms the divergent behavior of p as $\varphi \rightarrow \varphi_J$. In short, even if the LS is a rather quick compression protocol, it efficiently produces glassy configurations. Naturally, a more complex compression method can be used to attain better thermalized systems before




FIG. 9. Comparison of the pressure obtained during the LS compression (symbols) with the Carnahan–Starling equation of state [92] for the liquid (black, dashed line), and Eq. (15) for the glass phase (solid, red line). The fast compression is represented by the blue crosses, while the circles correspond to the slower compression in the glass phase. The inset shows that the pressure diverges as predicted by Eq. (15) in the main text.

initializing the CALiPPSO crunching part. However, as we argued in Sec. III A, this protocol suffices to probe the most salient features of jammed packings.

APPENDIX D: COMPARING CALiPPSO TO MD BASED ALGORITHMS

In the main text, we mentioned that using only the LS compression it is impossible to extract the full structural information of a jammed state. To support this claim, in Fig. 10 we present the probability distribution function (pdf) of the interparticle gaps, defined as $h_{ij} = \frac{r_{ij}}{\sigma_{ij}} - 1$, obtained at different values of p_{tar} . That is, the gaps’ pdf once the MD simulations have reached p_{tar} , but without using CALiPPSO. For a given value of target pressure, the associated pdf has three different regimes: (i) for very small values of the gaps, $h \lesssim 1/p_{\text{tar}}$, a plateau forms, whose height increases (linearly) with p_{tar} ; (ii) for $1/p_{\text{tar}} \ll h$, the pdf follows the scaling predicted by mean-field theory [6,12,13], $p(h) \sim h^{-\gamma}$, with $\gamma = 0.4163$ (dashed, black line); (iii) for $1/p_{\text{tar}} \lesssim h$ a third, intermediate regime sets in which cannot be identified with any well-known property of the configurations. The first two regimes can be identified, respectively, as the nascent contact-singularity and the critical gap distribution of the radial distribution function characteristic of spheres packings [50,76,92]. We verified that if only the gaps smaller than $1/p_{\text{tar}}$ are counted as contacts, the resulting system is highly hypostatic. Notably, these features are independent of the value of κ employed.

Figure 10 illustrates also the gap distributions obtained at convergence of CALiPPSO (circular markers), using the LS configurations as seeds. Notice that using CALiPPSO, $p(h)$ is rather insensitive to the value of p_{tar} from which a configuration is crunched. Specifically, we also recover the mean-field prediction, but in a consistent range spanning sev-




FIG. 10. Probability distribution function (pdf) of the interparticle gaps, h , after terminating the LS compression at different target pressures and with different compression rates (as indicated by colors and legend in the lower left). We show also the analogous pdf’s obtained with CALiPPSO for the same values of p_{tar} , with configurations compressed using $\kappa = 3 \times 10^{-5}$ (other values yield virtually identical distributions). These curves have been displayed downwards by a small fraction for clarity. For each value of κ and p_{tar} , we report the distributions obtained averaging over 20 configurations of $N = 1024$ particles. Down arrows indicate the intersection of the distributions with $h = 1/p_{\text{tar}}$ and correspond to the beginning of the plateau. Thus, for a given value of p_{tar} , LS “contacts” can be identified as the gaps smaller than $1/p_{\text{tar}}$. We observe the presence of the regime $p(h) \sim h^{-\gamma}$ (black, dashed line) which is in good agreement with the distribution of gaps predicted by the mean-field theory. As we argue in the main text, since intermediate values of h cannot be associated unequivocally to either of these two classes, using exclusively LS-type algorithms we cannot recover the detailed structure of jammed packings.

eral decades, independently of p_{tar} . Moreover, deviations from $p(h) \sim h^{-\gamma}$ observed in the left tail of $p(h)$ can be ascribed to finite-size effects [29]. The plateau proportional to p_{tar} present in the LS distribution, completely disappears after the crunching: it becomes $\delta(h)$, i.e., particles are in contact. Thus, the intermediate regime (iii) for $1/p_{\text{tar}} \lesssim h$ disappears as well because using CALiPPSO all the interparticle distances can be classified either as force bearing contacts, or as true gaps.

To better understand the difference between the results of our approach and highly compressed configurations obtained from MD simulations, let us assume a fixed value of p_{tar} . From the discussion of the previous paragraphs, by comparing $p(h)$ obtained within LS with the corresponding distribution within CALiPPSO, we can conclude that a significant fraction of gaps in the intermediate domain of the LS distribution will form valid contacts once ϕ_J is reached. Yet, for any $\phi < \phi_J$, there is no clear criterion for distinguishing this type of “precontact” gaps from true ones (i.e., those that remain finite at ϕ_J). To overcome this difficulty, previous studies have computed a time-averaged network of contacts. In a nutshell, this technique identifies the contacts of a particle by averaging the collisions it undergoes with its neighbors over a sufficiently large time window. Naturally, the plateau of $h < 1/p_{\text{tar}}$ is a consistent contribution to such average because such very small gaps constitute a rather constant background




FIG. 11. Running time of LS (dotted lines) and CALiPPSO algorithms (dashed), and their sum (solid), as a function of p_{tar} . We explore the effect of using different compression rates with fixed $N = 1024$ (upper panel), as well as changing the system size but a constant $\kappa = 10^{-5}$ (lower one), as indicated in the legends. For the analysis of Sec. IV, we used $p_{\text{tar}} = 10^7$ (highlighted by the vertical line).

of collisions associated with potential contacts. The remaining contacts needed to achieve stability are therefore obtained from the collisions with particles whose typical distance lies in the intermediate regime. Instead, as we have shown here, identifying forces with strictly positive Lagrange multipliers of the final jamming LOP offers a clear-cut distinction between gaps and contacts. This makes CALiPPSO a more precise packing algorithm than dynamical or time-average approaches coming from the near jamming regime. We emphasize that precisely identifying real contacts and gaps is of utmost importance because of the marginal stability of jammed packings and the long-range correlations of their networks of contacts.

To conclude, we compare (i) the time required by LS to reach a given p_{tar} , (ii) the time CALiPPSO takes to successfully jam the configuration from the target pressure, and (iii) their sum, i.e., the total time of our LS+CALiPPSO method. For the first one, it is clear that the larger the desired p_{tar} the longer the LS protocol takes (dotted lines in Fig. 11). Besides, intuition suggests the duration of the CALiPPSO crunching process is reduced the larger p_{tar} gets, leading to a trade-off of the optimal total time of LS+CALiPPSO. However, from the results of Figs. 3(g) and 3(h), we know that if $p_{\text{tar}} > p^{(\text{th})}$ this is not necessarily the case. Indeed, those figures show that $\langle n \rangle$ remains essentially unchanged in this very high-pressure regime. Hence, it is expected that the convergence time of CALiPPSO remains roughly constant, as confirmed in Fig. 11 (dashed lines). Thus, at least for moderately large N , increasing p_{tar} is actually detrimental to the performance of LS+CALiPPSO. Recall, however, that to properly sample the free energy landscape (see Sec. III A) a large target pressure, $p^{(\text{th})} > p_{\text{tar}} \gg 1$ is useful. Therefore, even if choosing a very high value of p_{tar} leads to a longer LS compression, it should be favored to better model the thermodynamic route to the jamming line. This is the reason why we fixed $p_{\text{tar}} = 10^7$ in Sec. IV to explore the CALiPPSO algorithmic complexity.

From a more practical point of view, if many independent configurations have to be compressed simultaneously, then a long LS compression does not necessarily hinder the performance of LS+CALiPPSO, provided $p_{\text{tar}} < p^{(\text{th})}$. That is, given that LS [50] relies on event-driven MD that can be efficiently implemented using a serial algorithm, as many configurations as available threads can be compressed. However, some optimizers, as Gurobi [81], benefit from being executed with parallelized algorithms, which hinders the possibility of concurrently executing CALiPPSO on several systems. Therefore, when selecting p_{tar} to optimize the total time of LS+CALiPPSO these two different behaviors should be considered.

As a final remark, we mention that a relatively large κ can produce a nonmonotonic behavior of time of the full LS+CALiPPSO protocol. In this case, a longer LS compression to reach a larger pressure is convenient up to a certain p_{tar} , that depends on κ . Beyond such pressure, however, further increasing the target pressure leads to a larger time because the CALiPPSO crunching is not substantially accelerated. In other words, the small speed gain of CALiPPSO does not compensate for the longer LS compression.

- [1] A. J. Liu and S. R. Nagel, Nonlinear dynamics: Jamming is not just cool any more, *Nature (London)* **396**, 21 (1998).
- [2] A. J. Liu and S. R. Nagel, The jamming transition and the marginally jammed solid, *Annu. Rev. Condens. Matter Phys.* **1**, 347 (2010).
- [3] S. Torquato and F. H. Stillinger, Jammed hard-particle packings: From Kepler to Bernal and beyond, *Rev. Mod. Phys.* **82**, 2633 (2010).
- [4] M. van Hecke, Jamming of soft particles: Geometry, mechanics, scaling and isostaticity, *J. Phys.: Condens. Matter* **22**, 033101 (2010).
- [5] C. 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, *Phys. Rev. E* **68**, 011306 (2003).
- [6] P. Charbonneau, J. Kurchan, G. Parisi, P. Urbani, and F. Zamponi, Fractal free energy landscapes in structural glasses, *Nat. Commun.* **5**, 3725 (2014).
- [7] A. Baule, F. Morone, H. J. Herrmann, and H. A. Makse, Edwards statistical mechanics for jammed granular matter, *Rev. Mod. Phys.* **90**, 015006 (2018).
- [8] B. P. Tighe, J. H. Snoeijer, T. J. H. Vlugt, and M. van Hecke, The force network ensemble for granular packings, *Soft Matter* **6**, 2908 (2010).

- [9] E. DeGiuli, E. Lerner, and M. Wyart, Theory of the jamming transition at finite temperature, *J. Chem. Phys.* **142**, 164503 (2015).
- [10] G. Parisi and F. Zamponi, Mean-field theory of hard sphere glasses and jamming, *Rev. Mod. Phys.* **82**, 789 (2010).
- [11] P. 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, *J. Stat. Mech.: Theory Exp.* (2014) P10009.
- [12] P. Charbonneau, J. Kurchan, G. Parisi, P. Urbani, and F. Zamponi, Glass and jamming transitions: From exact results to finite-dimensional descriptions, *Annu. Rev. Condens. Matter Phys.* **8**, 265 (2017).
- [13] G. Parisi, P. Urbani, and F. Zamponi, *Theory of Simple Glasses: Exact Solutions in Infinite Dimensions* (Cambridge University Press, Cambridge, UK, 2020).
- [14] L. Berthier, G. Biroli, P. Charbonneau, E. I. Corwin, S. Franz, and F. Zamponi, Gardner physics in amorphous solids and beyond, *J. Chem. Phys.* **151**, 010901 (2019).
- [15] S. Franz and G. Parisi, The simplest model of jamming, *J. Phys. A: Math. Theor.* **49**, 145001 (2016).
- [16] S. Franz, G. Parisi, M. Sevelev, P. Urbani, and F. Zamponi, Universality of the SAT-UNSAT (jamming) threshold in nonconvex continuous constraint satisfaction problems, *SciPost Phys.* **2**, 019 (2017).
- [17] M. 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, *Phys. Rev. E* **100**, 012115 (2019).
- [18] S. 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, *J. Phys. A: Math. Theor.* **52**, 474001 (2019).
- [19] F. Antenucci, S. Franz, P. Urbani, and L. Zdeborová, Glassy Nature of the Hard Phase in Inference Problems, *Phys. Rev. X* **9**, 011020 (2019).
- [20] S. Franz, G. Parisi, P. Urbani, and F. Zamponi, Universal spectrum of normal modes in low-temperature glasses, *Proc. Natl. Acad. Sci. USA* **112**, 14539 (2015).
- [21] S. Franz, A. Sclocchi, and P. Urbani, Critical Jammed Phase of the Linear Perceptron, *Phys. Rev. Lett.* **123**, 115702 (2019).
- [22] S. Franz, T. Maimbourg, G. Parisi, and A. Scardicchio, Impact of jamming criticality on low-temperature anomalies in structural glasses, *Proc. Natl. Acad. Sci. USA* **116**, 13768 (2019).
- [23] C. Artiaco, F. Balducci, G. Parisi, and A. Scardicchio, Quantum jamming: Critical properties of a quantum mechanical perceptron, *Phys. Rev. A* **103**, L040203 (2021).
- [24] P. Charbonneau, E. I. Corwin, G. Parisi, and F. Zamponi, Jamming Criticality Revealed by Removing Localized Buckling Excitations, *Phys. Rev. Lett.* **114**, 125504 (2015).
- [25] P. Charbonneau, E. I. Corwin, G. Parisi, and F. Zamponi, Universal Microstructure and Mechanical Stability of Jammed Packings, *Phys. Rev. Lett.* **109**, 205501 (2012).
- [26] R. C. Dennis and E. I. Corwin, Jamming Energy Landscape is Hierarchical and Ultrametric, *Phys. Rev. Lett.* **124**, 078002 (2020).
- [27] E. DeGiuli, E. Lerner, C. Brito, and M. Wyart, Force distribution affects vibrational properties in hard-sphere glasses, *Proc. Natl. Acad. Sci. USA* **111**, 17054 (2014).
- [28] E. Lerner, G. Düring, and M. Wyart, Low-energy nonlinear excitations in sphere packings, *Soft Matter* **9**, 8252 (2013).
- [29] P. 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, *Phys. Rev. E* **104**, 014102 (2021).
- [30] C. P. Goodrich, A. J. Liu, and S. R. Nagel, Finite-Size Scaling at the Jamming Transition, *Phys. Rev. Lett.* **109**, 095704 (2012).
- [31] C. P. Goodrich, A. J. Liu, and J. P. Sethna, Scaling ansatz for the jamming transition, *Proc. Natl. Acad. Sci.* **113**, 9745 (2016).
- [32] C. Artiaco, P. Baldan, and G. Parisi, Exploratory study of the glassy landscape near jamming, *Phys. Rev. E* **101**, 052605 (2020).
- [33] D. Hexner, P. Urbani, and F. Zamponi, Can a Large Packing be Assembled from Smaller Ones? *Phys. Rev. Lett.* **123**, 068003 (2019).
- [34] D. Hexner, A. J. Liu, and S. R. Nagel, Two Diverging Length Scales in the Structure of Jammed Packings, *Phys. Rev. Lett.* **121**, 115501 (2018).
- [35] F. Arceri and E. I. Corwin, Vibrational Properties of Hard and Soft Spheres are Unified at Jamming, *Phys. Rev. Lett.* **124**, 238002 (2020).
- [36] V. 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 (2019).
- [37] C. Coulais, R. P. Behringer, and O. Dauchot, How the ideal jamming point illuminates the world of granular media, *Soft Matter* **10**, 1519 (2014).
- [38] O. Dauchot, G. Marty, and G. Biroli, Dynamical Heterogeneity Close to the Jamming Transition in a Sheared Granular Material, *Phys. Rev. Lett.* **95**, 265701 (2005).
- [39] F. Lechenault, O. Dauchot, G. Biroli, and J. P. Bouchaud, Critical scaling and heterogeneous superdiffusion across the jamming/rigidity transition of a granular glass, *Europhys. Lett.* **83**, 46003 (2008).
- [40] A. Seguin and O. Dauchot, Experimental Evidence of the Gardner Phase in a Granular Glass, *Phys. Rev. Lett.* **117**, 228001 (2016).
- [41] Y. Wang, J. Shang, Y. Jin, and J. Zhang, Experimental observations of marginal criticality in granular materials, *Proc. Natl. Acad. Sci. USA* **119**, e2204879119 (2022).
- [42] T. Aste, M. Saadatfar, and T. J. Senden, Geometrical structure of disordered sphere packings, *Phys. Rev. E* **71**, 061302 (2005).
- [43] T. Aste, Volume Fluctuations and Geometrical Constraints in Granular Packs, *Phys. Rev. Lett.* **96**, 018002 (2006).
- [44] E. Bitzek, P. Koskinen, F. Gähler, M. Moseler, and P. Gumbsch, Structural Relaxation Made Simple, *Phys. Rev. Lett.* **97**, 170201 (2006).
- [45] P. K. Morse and E. I. Corwin, Geometric Signatures of Jamming in the Mechanical Vacuum, *Phys. Rev. Lett.* **112**, 115701 (2014).
- [46] P. Charbonneau, E. I. Corwin, G. Parisi, A. Poncet, and F. Zamponi, Universal Non-Debye Scaling in the Density of States of Amorphous Solids, *Phys. Rev. Lett.* **117**, 045503 (2016).
- [47] P. K. Morse and E. I. Corwin, Echoes of the Glass Transition in Athermal Soft Spheres, *Phys. Rev. Lett.* **119**, 118003 (2017).
- [48] P. Charbonneau and P. K. Morse, Memory Formation in Jammed Hard Spheres, *Phys. Rev. Lett.* **126**, 088001 (2021).

- [49] B. D. Lubachevsky and F. H. Stillinger, Geometric properties of random disk packings, *J. Stat. Phys.* **60**, 561 (1990).
- [50] M. Skoge, A. Donev, F. H. Stillinger, and S. Torquato, Packing hyperspheres in high-dimensional Euclidean spaces, *Phys. Rev. E* **74**, 041127 (2006).
- [51] S. Torquato, T. M. Truskett, and P. G. Debenedetti, Is Random Close Packing of Spheres Well Defined? *Phys. Rev. Lett.* **84**, 2064 (2000).
- [52] K. Zhang, W. W. Smith, M. Wang, Y. Liu, J. Schroers, M. D. Shattuck, and C. S. O'Hern, Connection between the packing efficiency of binary hard spheres and the glass-forming ability of bulk metallic glasses, *Phys. Rev. E* **90**, 032311 (2014).
- [53] E. Lerner, G. Düring, and M. Wyart, Simulations of driven overdamped frictionless hard spheres, *Comput. Phys. Commun.* **184**, 628 (2013).
- [54] C. Brito and M. Wyart, On the rigidity of a hard-sphere glass near random close packing, *Europhys. Lett.* **76**, 149 (2006).
- [55] C. Brito and M. Wyart, Geometric interpretation of previtrification in hard sphere liquids, *J. Chem. Phys.* **131**, 024504 (2009).
- [56] S. Henkes, C. Brito, and O. Dauchot, Extracting vibrational modes from fluctuations: A pedagogical discussion, *Soft Matter* **8**, 6092 (2012).
- [57] A. Altieri, S. Franz, and G. Parisi, The jamming transition in high dimension: An analytical study of the TAP equations and the effective thermodynamic potential, *J. Stat. Mech.: Theory Exp.* (2016) 093301.
- [58] D. Frenkel, Order through entropy, *Nat. Mater.* **14**, 9 (2015).
- [59] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, Julia: A fresh approach to numerical computing, *SIAM Rev.* **59**, 65 (2017).
- [60] Our implementation of the CALiPPSO algorithm, <https://github.com/rdhr/CALiPPSO.jl>.
- [61] R. 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 (2021).
- [62] S. Torquato and Y. Jiao, Robust algorithm to generate a diverse class of dense disordered and ordered sphere packings via linear programming, *Phys. Rev. E* **82**, 061302 (2010).
- [63] K. Krabbenhoft, A. V. Lyamin, J. Huang, and M. Vicente da Silva, Granular contact dynamics using mathematical programming methods, *Comput. Geotech.* **43**, 165 (2012).
- [64] A. Donev, S. Torquato, F. H. Stillinger, and R. Connelly, A linear programming algorithm to test for jamming in hard-sphere packings, *J. Comput. Phys.* **197**, 139 (2004).
- [65] A. B. Hopkins, Y. Jiao, F. H. Stillinger, and S. Torquato, Phase Diagram and Structural Diversity of the Densest Binary Sphere Packings, *Phys. Rev. Lett.* **107**, 125501 (2011).
- [66] A. B. Hopkins, F. H. Stillinger, and S. Torquato, Densest binary sphere packings, *Phys. Rev. E* **85**, 021130 (2012).
- [67] A. B. Hopkins, F. H. Stillinger, and S. Torquato, Disordered strictly jammed binary sphere packings attain an anomalously large range of densities, *Phys. Rev. E* **88**, 022205 (2013).
- [68] Y. Jiao, F. H. Stillinger, and S. Torquato, Nonuniversality of density and disorder in jammed sphere packings, *J. Appl. Phys.* **109**, 013508 (2011).
- [69] J.-N. Roux, Geometric origin of mechanical properties of granular materials, *Phys. Rev. E* **61**, 6802 (2000).
- [70] D. G. Luenberger and Y. Ye, *Linear and Nonlinear Programming*, International Series in Operations Research & Management Science, Vol. 228 (Springer International Publishing, Cham, 2016).
- [71] S. P. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, Cambridge, UK, 2004).
- [72] J. H. Conway and N. J. A. Sloane, *Sphere Packings, Lattices and Groups* (Springer Science & Business Media, Cham, 2013).
- [73] P. Charbonneau, P. K. Morse, W. Perkins, and F. Zamponi, Three simple scenarios for high-dimensional sphere packings, *Phys. Rev. E* **104**, 064612 (2021).
- [74] R. Blumenfeld, Disorder Criterion and Explicit Solution for the Disc Random Packing Problem, *Phys. Rev. Lett.* **127**, 118002 (2021).
- [75] E. L. Hinrichsen, J. Feder, and T. Jøssang, Random packing of disks in two dimensions, *Phys. Rev. A* **41**, 4199 (1990).
- [76] A. Donev, S. Torquato, and F. H. Stillinger, Pair correlation function characteristics of nearly jammed disordered and ordered hard-sphere packings, *Phys. Rev. E* **71**, 011105 (2005).
- [77] A. Donev, S. Torquato, F. H. Stillinger, and R. Connelly, Comment on “Jamming at zero temperature and zero applied stress: The epitome of disorder,” *Phys. Rev. E* **70**, 043301 (2004).
- [78] C. 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’,” *Phys. Rev. E* **70**, 043302 (2004).
- [79] A. Donev, S. Torquato, and F. H. Stillinger, Neighbor list collision-driven molecular dynamics simulation for nonspherical hard particles. I. Algorithmic details, *J. Comput. Phys.* **202**, 737 (2005).
- [80] I. Dunning, J. Huchette, and M. Lubin, Jump: A modeling language for mathematical optimization, *SIAM Rev.* **59**, 295 (2017).
- [81] Gurobi Optimization, LLC, Gurobi Optimizer Reference Manual (2021), <http://www.gurobi.com>.
- [82] Q. Huangfu and J. A. J. Hall, Parallelizing the dual revised simplex method, *Math. Program. Comput.* **10**, 119 (2018).
- [83] A. Makhorin, GLPK (GNU linear programming kit) (2008), <https://www.gnu.org/software/glpk/>.
- [84] A. Stukowski, Visualization and analysis of atomistic simulation data with OVITO—The open visualization tool, *Modell. Simul. Mater. Sci. Eng.* **18**, 015012 (2010).
- [85] Note that Eq. (15) has been derived from a free volume analysis [87]. Therefore, it does not correspond to the true thermodynamic equation of state [12,13] (see also Fig. 9 in Appendix C for more details).
- [86] L. Berthier, P. Charbonneau, Y. Jin, G. Parisi, B. Seoane, and F. Zamponi, Growing timescales and lengthscales characterizing vibrations of amorphous solids, *Proc. Natl. Acad. Sci. USA* **113**, 8397 (2016).
- [87] Z. W. Salsburg and W. W. Wood, Equation of state of classical hard spheres at high density, *J. Chem. Phys.* **37**, 798 (1962).
- [88] P. Charbonneau, Y. Jin, G. Parisi, C. Rainone, B. Seoane, and F. Zamponi, Numerical detection of the Gardner transition in a mean-field glass former, *Phys. Rev. E* **92**, 012316 (2015).
- [89] Jamming criticality can also be clearly observed in two-dimensional systems, provided that bidisperse or polydisperse packings are used. In that sense, 3d monodisperse systems are simpler because the random particles' positions are the only source of disorder.

- [90] C. Rainone and P. Urbani, Following the evolution of glassy states under external perturbations: The full replica symmetry breaking solution, *J. Stat. Mech.: Theory Exp.* (2016) 053302.
- [91] F. H. Stillinger, Exponential multiplicity of inherent structures, *Phys. Rev. E* **59**, 48 (1999).
- [92] A. Santos, S. B. Yuste, and M. López de Haro, Structural and thermodynamic properties of hard-sphere fluids, *J. Chem. Phys.* **153**, 120901 (2020).
- [93] L. Berthier, D. Coslovich, A. Ninarello, and M. Ozawa, Equilibrium Sampling of Hard Spheres Up to the Jamming Density and Beyond, *Phys. Rev. Lett.* **116**, 238002 (2016).
- [94] S. Hwang and H. Ikeda, Force balance controls the relaxation time of the gradient descent algorithm in the satisfiable phase, *Phys. Rev. E* **101**, 052308 (2020).
- [95] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed., Springer Series in Operations Research (Springer, New York, 2006).
- [96] J. 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, *J. Phys. Chem. B* **117**, 12979 (2013).
- [97] P. Rissone, E. I. Corwin, and G. Parisi, Long-Range Anomalous Decay of the Correlation in Jammed Packings, *Phys. Rev. Lett.* **127**, 038001 (2021).