





Large-scale frictionless jamming with power-law particle size distributionsJoseph M. Monti ^{1,*}, Joel T. Clemmer ¹, Ishan Srivastava ², Leonardo E. Silbert ³,
Gary S. Grest,¹ and Jeremy B. Lechman¹¹*Sandia National Laboratories, Albuquerque, New Mexico 87185, USA*²*Center for Computational Sciences and Engineering, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*³*School of Math, Science, and Engineering, Central New Mexico Community College, Albuquerque, New Mexico 87106, USA*

(Received 17 May 2022; accepted 9 August 2022; published 2 September 2022)

Due to significant computational expense, discrete element method simulations of jammed packings of size-dispersed spheres with size ratios greater than 1:10 have remained elusive, limiting the correspondence between simulations and real-world granular materials with large size dispersity. Invoking a recently developed neighbor binning algorithm, we generate mechanically stable jammed packings of frictionless spheres with power-law size distributions containing up to nearly 4 000 000 particles with size ratios up to 1:100. By systematically varying the width and exponent of the underlying power laws, we analyze the role of particle size distributions on the structure of jammed packings. The densest packings are obtained for size distributions that balance the relative abundance of large-large and small-small particle contacts. Although the proportion of rattler particles and mean coordination number strongly depend on the size distribution, the mean coordination of nonrattler particles attains the frictionless isostatic value of six in all cases. The size distribution of nonrattler particles that participate in the load-bearing network exhibits no dependence on the width of the total particle size distribution beyond a critical particle size for low-magnitude exponent power laws. This signifies that only particles with sizes greater than the critical particle size contribute to the mechanical stability. However, for high-magnitude exponent power laws, all particle sizes participate in the mechanical stability of the packing.

DOI: [10.1103/PhysRevE.106.034901](https://doi.org/10.1103/PhysRevE.106.034901)**I. INTRODUCTION**

Packings of stiff granular particles with a high degree of size dispersity are of widespread geophysical and industrial relevance, with applications including powder technology and the mechanics of soil and construction materials [1–6]. The distribution of particle sizes can adopt discrete or continuous forms, both of which have been shown for frictionless particles to produce overall packing densities, ϕ , that are greater than the frictionless, monodisperse value $\phi^{\text{mono}} \approx 0.64$ [7]. The simplest discrete form is the bidisperse case, for which Furnas [3] predicted the theoretical limiting value of $\phi^{\text{bi}} \approx 0.87$ for an infinitely large size ratio; recent large-scale numerical simulations of bidisperse packings produced packing densities approaching the Furnas limit [8,9]. To date, most three-dimensional (3D) numerical simulations of continuous, highly disperse systems with size distributions of diverse functional forms have been limited to largest-to-smallest particle size ratios of order 10 or less, and typically only reach packing densities ~ 0.71 or smaller at low confining pressures [8,10–15]. To our knowledge, the main exception is the work of Oquendo-Patiño and Estrada [16–18], who considered power-law-like particle size distributions. These distributions were generated by matching the scaling behavior of the cumulative particle size distributions to early experimental observations by Fuller and Thompson [1]. Oquendo-Patiño and Estrada [16–18] simulated particle size ratios of up to 32

and achieved packing densities close to 0.86 depending on the characteristics of the particle size distribution.

Here, we consider power-law particle size distributions to study packings of highly disperse particles. From a numerical perspective, power laws are one of the simplest continuous distributions, since there are only two parameters governing the distribution: the maximum particle size ratio and the power-law exponent. Power-law distributions have been measured to emerge naturally from various fragmentation mechanisms, including that of sea-ice floes [19] and comminution [4,20–23]. Further, power laws display scale invariance and fractal behavior; the geometric Apollonian packing is one example of a fractal packing with an underlying power-law size distribution [24–27].

One of the challenges associated with simulating power-law distributions is that the tail of the distribution has significant weight and decays slowly, thus requiring the simulation of very large size ratios of particles to accurately sample the distribution. Computational costs of 3D discrete element method (DEM) simulations with broad particle size disparities have been prohibitive until recently due to algorithmic limitations. Conventional neighbor list generation methods, e.g., those available by default in popular molecular dynamics (MD) packages like LAMMPS [28,29] exhibit poor scaling with increasing size ratios and become intractable beyond particle size dispersity of order 10 [30]. Ogarko and Luding [31] recently developed an improved neighbor list generation scheme for highly size-dispersed systems. A similar approach has since been implemented in LAMMPS by Shire *et al.* [32]. This implementation has been expanded upon and was used to study bidisperse packings of both frictionless and frictional

*Corresponding author: jmmonti@sandia.gov

particles with particle size ratios of up to 40 [9]. For our study, we have exerted this simulation capability further to investigate strongly disperse power-law-distributed systems of frictionless particles with unprecedented particle size ratios of up to 100.

Packings of highly disperse particles require careful treatment as relaxation may occur over disparate timescales. Unlike volume-controlled jamming protocols, pressure-controlled jamming protocols are guaranteed to produce mechanically stable packings and yield greater accessibility to the jamming point in the low pressure regime [33,34]. Recently, the isobaric-isoenthalpic (NPH) thermodynamic ensemble, one example of a constant-pressure protocol, was successfully applied to multi-friction-mode monodisperse packings [35] and to frictionless and frictional bidisperse packings [9]. The NPH ensemble implementation in LAMMPS can be leveraged to enforce the condition of zero shear stresses applied to the simulation box concurrently with isotropic compression. For the highly disperse systems considered in this article, permitting the relaxation to zero of the off-diagonal components of the internal stress tensor, \mathbf{P}_{int} , is crucial as this technique produces packings that are also stable with respect to shear deformations.

This article describes numerical simulations performed using a constant-pressure (NPH) compression protocol to generate jammed packings of frictionless, power-law-distributed disperse spherical particles. The interparticle contact model is described in Sec. II A. Section II B elucidates how distributions of particle sizes are generated and characterized. Section II C contains a brief description of the multineighboring scheme used in this article. Further details and benchmark results can be found in the Appendix. The packing protocol is described in Sec. II D. In Sec. III A, the results of packed systems of power-law-distributed particles are characterized. Finally, Sec. III B examines several properties of the resultant force-bearing networks.

II. METHODS

A. Contact model

Spherical particle-based 3D DEM packing simulations were conducted using the GRANULAR package in LAMMPS [28,29,36]. The scope of this article is limited to frictionless, purely repulsive normal contacts, where particles interact via a damped Hookean pair potential penalizing overlap. The normal force \mathbf{F}_n between contacting particles i and j with diameters D_i and D_j and separation $\mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j$ is

$$\mathbf{F}_n = k_n \delta \hat{\mathbf{n}} - M_{\text{eff}} \gamma_n \mathbf{v}_n, \quad (1)$$

where $k_n = 1$ is the Hookean spring constant, $\delta = (D_i + D_j)/2 - |\mathbf{r}_{ij}|$ is the overlap, $M_{\text{eff}} = M_i M_j / (M_i + M_j)$ in terms of the particle masses M_i and M_j , and $\gamma_n = 0.5$ is a damping coefficient, reflecting particle inelasticity. The unit vector connecting the particle centers is $\hat{\mathbf{n}} = \mathbf{r}_{ij}/|\mathbf{r}_{ij}|$, and \mathbf{v}_n is the relative velocity of the two particles projected along $\hat{\mathbf{n}}$. Note that, in principle, the net normal force in Eq. (1) can be attractive, i.e., if the damping component is greater than the Hookean component when particles are moving apart. An extra switching function is employed to set the magnitude of \mathbf{F}_n to zero if this condition occurs during the simulation. We

do not expect this formulation to cause any significant changes for slow compression simulations, but it may be an important consideration for high-rate deformation simulations, for instance.

The material density of individual particles is $\rho = 1$, such that particle masses M_i are proportional to the particle volumes V_i and given by $M_i = \rho V_i = \rho \pi D_i^3/6$. The unit of length is the smallest particle diameter $D_{\text{min}} = 1$, and the unit of pressure is k_n/D_{min} ; all lengths and pressures are given in terms of these quantities. The simulation time step is $\Delta t = 0.02\tau$, where $\tau = \sqrt{M_{\text{min}}/k_n}$ with $M_{\text{min}} \equiv \rho \pi/6$.

B. Particle size distributions

Particle sizes are represented using diameters D and are distributed according to power-law distributions such that the probability of finding a particle with diameter between D and $D + dD$ is $P(D)dD \propto D^{-\beta}dD$, where β is the power-law distribution exponent. Particle sizes are limited to a range $1 \leq D \leq \lambda$, where the parameter λ denotes the maximum size ratio of the distribution. Each system is required to have at least ten particles with diameters larger than 0.95λ , meaning that the total number of particles in each system depends upon both λ and β . In the geophysical literature, distributions of particle sizes are often given in terms of their fractal dimensions d_f , meaning that the number of particles N_D larger than size D satisfies $N_D \sim D^{-d_f}$ [4]. For power-law particle size distributions,

$$N_D \propto \int_D^\infty D'^{-\beta} dD' \propto \frac{D^{1-\beta}}{1-\beta}, \quad (2)$$

for $\beta > 1$, so $d_f = \beta - 1$.

A central quantity of interest is the cumulative volume fraction (CVF), which gives the fraction of particle volume (and mass, since ρ is constant) contained in particles smaller than a given size. The CVF is easily obtained for power-law distributions with $\beta < 4$ as

$$\text{CVF} = \frac{\int_1^D D'^{3-\beta} dD'}{\int_1^\lambda D'^{3-\beta} dD'} = \frac{D^{4-\beta} - 1}{\lambda^{4-\beta} - 1}. \quad (3)$$

The CVF exponent, α , is defined using Eq. (3) as $\alpha \equiv 4 - \beta$, and gives the scaling of the CVF in the limits $D^\alpha \gg 1$ and $\lambda^\alpha \gg 1$ as $\text{CVF} \sim (D/\lambda)^\alpha$. Note that α plays a similar role as the grain-size distribution (GSD) exponent η in Refs. [16–18], i.e., $\text{GSD} \sim [(D-1)/(\lambda-1)]^\eta$, but cannot be compared directly (except in the specific case $\alpha = \eta = 1.0$, corresponding to $\beta = 3$) as the underlying particle size distributions in Refs. [16–18] are not power laws characterized by a single exponent.

This article mainly considers exponents in the range $3.0 \leq \beta \leq 3.8$, or $0.2 \leq \alpha \leq 1.0$, emblematic of soil comminution [4,20–23]. Several of these CVFs are shown in Fig. 1(a) plotted against the reduced particle diameter $(D-1)/(\lambda-1)$. A simple interpretation of Fig. 1(a) is that more than half the particle volume (mass) is contained in particles with diameters smaller than the arithmetic mean diameter $(\lambda+1)/2$, or reduced particle diameter $1/2$, for $\alpha < 1.0$ —the concave down curves in Fig. 1(a)—and more than half the particle volume is contained in particles larger than $(\lambda+1)/2$ for




FIG. 1. (a) Cumulative volume fractions for select values of α and $\lambda = 50$. Points are computed from the particle distributions used in the simulations and lines are the analytic curves obtained from Eq. (3). The dashed black line corresponds to $\alpha = 1.5$ and is included as an example case for $\alpha > 1.0$, for which the CVF is concave up. (b) Example packings obtained at applied pressure $p_a = 10^{-6}$, shaded by particle diameter for the indicated α . Both simulation boxes are triclinic but have small tilt factors. The overall number of particles for $\alpha = 0.2$ is $\sim 16\times$ larger than for $\alpha = 1.0$, and the packed volume fractions are close to $\phi = 0.82$ and $\phi = 0.80$ for $\alpha = 0.2$ and $\alpha = 1.0$, respectively.

$\alpha > 1.0$. As demonstrated in Sec. III A, the range of values of α shown in Fig. 1(a) brackets the densest obtainable packing for $\lambda \gg 1$. As $\alpha \rightarrow 0$ ($\beta \rightarrow 4$), the preponderance of particles have diameters close to $D = 1$ and the total count of particles rises sharply. For example, for $\lambda = 50$, the $\alpha = 1.0$ system has 233 653 particles while for $\alpha = 0.2$ there are 3 739 236 particles. The corresponding packings obtained at low applied pressure (see Sec. II D) for these two systems are shown in Fig. 1(b), rendered in OVITO [37].

C. Efficient multineighboring scheme

To identify potentially interacting atoms (particles) in MD and DEM packages, the most computationally efficient basic algorithm builds a neighbor list with all pairs of nearby particles using a spatial grid with a length scale set by the largest interaction cutoff. This becomes impractical as the size disparity ratio λ increases, as the same bin size is used for all particle pairs. An alternative approach was implemented into LAMMPS by in 't Veld *et al.* [30], which uses the smallest cutoff to set the bin size and to adjust how many bins are searched based on particle types [30]. In LAMMPS, particle types are a discrete categorization used to set interaction parameters such as cutoffs for MD or friction coefficients for DEM. While this method allows simulations to reach larger λ , it also becomes exceedingly expensive as λ increases beyond ~ 10 . To overcome this limitation, an improved algorithm was recently proposed by Ogarko and Luding [31] and initially modified for LAMMPS by Stratford *et al.* [38] and Shire *et al.* [32]. This approach further tailored the neighbor list construction based on a particle's type to ensure that the computational cost of building a neighbor list does not grow faster with λ than the force calculation.

For this article, the implementation by Stratford *et al.* [38] was expanded upon by fully integrating it with the LAMMPS codebase and releasing it in the public LAMMPS distribution [39]. The method is generalized to support DEM by removing

the use of particle types, since these are typically intended to describe material properties and not necessarily particle sizes. Neighbor list construction can be tuned by predefining a set of diameter intervals irrespective of particle types, streamlining optimization of simulations. The crux of the technique is that each particle searches for neighbors with diameters that fall in its own diameter interval and in larger diameter intervals. This approach takes advantage of the inherent asymmetry in the computational effort required to generate lists of neighbors centering on small particles as opposed to using large particles as the point of reference. Previous work demonstrated that this method can be used to model jamming of frictionless and frictional bidisperse packings up to $\lambda = 40$ [9]. Here, this methodology is applied to study packings of frictionless particles with a power-law distribution with λ as large as 100, although larger λ are feasible. See the Appendix for arguments regarding the computational complexity of the algorithm and benchmark results.

D. Constant-pressure packing protocol

Packings are created via a constant-pressure protocol using the NPH ensemble implemented in LAMMPS [9,35]. The symmetric applied pressure tensor, \mathbf{P}_a , has the form $P_{a,xx} = P_{a,yy} = P_{a,zz} = p_a$ and all off-diagonal components are zero. Here, p_a is set to 10^{-6} to work in the limit of small particle overlaps; for context, in systems of monodisperse particles the typical fractional overlap is $\delta/D \sim p_a$ (in units of D/k_n). The simulation box is fully periodic and initially cubic. Under the constraint of no overlaps, i.e., there are no inter-particle forces at time $t = 0$, particles are randomly placed throughout the simulation box at low volume density. The overall particle volume fraction is defined as $\phi \equiv (\sum_i V_i)/V$, where V is the instantaneous simulation box volume and the sum runs over all the particles in the system. During the simulation, the applied pressure compresses the simulation box and forces particles into contact; see Fig. 2 for an illustration




FIG. 2. Particle volume fraction ϕ , overall mean coordination $\langle Z \rangle$ (dashed), and nonrattler mean coordination $\langle Z \rangle_{nr}$ (solid), plotted against simulation times for $\lambda = 50$ and $\alpha = 0.7$. The inset magnifies the end of the simulation, at which time ϕ and $\langle Z \rangle_{nr}$ cease evolving.

of the typical variation of ϕ with simulation time. At the end of the simulation, the internal pressure tensor balances the applied pressure, giving $\mathbf{P}_{int} = \mathbf{P}_a$ within numerical tolerance. While the simulation box is triclinic, the box tilt factors are typically small compared to the characteristic box side length. The rate of compression is slow enough that the simulation box volume monotonically decreases until the system jams, and particle overlaps are much smaller than the particle diameters.

As the number of particles and the distribution of particle volumes vary substantially, we do not identify the final jammed state by using a fixed kinetic energy cutoff. Rather, several criteria are used to determine when to stop jamming simulations. In addition to the numerical equivalence of the final internal and applied pressure tensors, other quantities are also considered, including the evolution of ϕ and the mean number of contacts per particle, $\langle Z \rangle$, where $\langle \cdot \rangle$ refers to the average over all particles. After jamming is achieved, ϕ and $\langle Z \rangle$ do not evolve in time and the total kinetic energy is small: the average kinetic energy per particle is of order 10^{-13} or less. For several of the largest systems— $\lambda = 50$ systems with $\beta \gtrsim 3.6$ ($\alpha \lesssim 0.4$)—a simulation time cutoff of at least $1.2 \times 10^7 \tau$ and up to $\sim 4 \times 10^7 \tau$ is employed out of computational necessity to stop simulations. Most quantities extracted from the simulations, such as ϕ and $\langle Z \rangle$, evolve slowly if at all after such long run times (see Fig. 2).

In frictionless, monodisperse systems under vanishingly small pressure, the isostatic number of contacts per particle is $Z_{iso} = 6$. The number of excess contacts per particle, $\Delta Z = Z - Z_{iso}$, grows systematically with pressure as $\sqrt{p_a}$ [7,35,40]. For highly disperse packings, a large proportion of particles are rattlers, i.e., those particles participating in too few contact pairs to be mechanically stable. Such particles substantially dilute the calculation of $\langle Z \rangle$, but a separate, more informative value, $\langle Z \rangle_{nr}$, can be obtained by excluding rattler




FIG. 3. Variation in the particle volume fraction ϕ of jammed packings plotted against α for the indicated values of the maximum particle size λ . The vertical dashed line corresponds to the CVF exponent $\alpha_A \approx 0.53$ corresponding to the random Apollonian packing.

particles from the calculation. The difference between the two measures is evident by comparing the two red curves in Fig. 2. This issue will be discussed in greater detail in Sec. III.

As a separate test of mechanical stability, we conducted several additional simulations to verify that the packings with the most extreme fractions of rattler particles are still stable after removing rattlers (see Sec. III A). For example, for $\alpha = 1.0$ approximately 83% and 97% of all particles are rattlers for $\lambda = 20$ and 50, respectively. For several high- α values, the simulations were restarted after removing rattlers from the packed configurations and checked for reconvergence of the macroscopic quantities, including ϕ and $\langle Z \rangle_{nr}$, under the same stress state.

A final useful metric, the Cundall parameter C , quantifies the typical unbalanced net per-particle force, per contact, in the system [16],

$$C = \frac{\sum_p |\mathbf{F}_p|}{\sum_c |\mathbf{F}_c|}, \quad (4)$$

where the numerator sums over the magnitude of the net per-particle force and the denominator sums over the magnitude of each contact force. Our results showed that a value of $C \lesssim 10^{-6}$ was indicative of a mechanically stable system.

III. RESULTS

A. Characterization of packings

Using the constant-pressure packing protocol, packing volume fractions ϕ were obtained for different power-law particle size distributions. Results comparing the variation of ϕ with α and λ are shown in Fig. 3. For each λ , ϕ is lowest at the endpoints of the range of α considered and peaks near the center of the range. The peak shifts slightly to smaller α and becomes sharper with increasing λ . The sharpening trend

with increasing λ is similar to behavior observed in bidisperse packings, for which increasing the particle size ratio changes ϕ at the Furnas peak from smoothly nonmonotonic to cusped [9,41]. System size and packing equilibration constraints prevent us from exploring the entire range of α for $\lambda = 100$ to see if the sharpening trend persists in the power-law-distributed case.

The results in Fig. 3 are consistent with the packing densities obtained by Oquendo-Patiño and Estrada [16,18], who employed the Hertz contact model and volume-controlled isotropic compression. Several other differences between our simulations and those described in Refs. [16–18] warrant mentioning. Based on the maximum fractional overlaps quoted in Ref. [16], the peak pressure reached in the simulations conducted in that work is estimated to be of order $p_a \sim 10^{-4}$, or two orders of magnitude larger than in our simulations. Test simulations we conducted for $\lambda = 20$ comparing the Hooke and Hertz contact models at identical $p_a = 10^{-4}$ using the compression protocol described in Sec. IID showed that ϕ is $\sim 1\%$ larger in Hertzian systems with all other variables kept constant, perhaps because the Hertz model does not penalize incipient particle overlap (the contact stiffness is zero at first contact). Most importantly, as noted earlier, the GSD characterized by the exponent η in Ref. [16] does not correspond to an underlying single-exponent power-law size distribution, except in the case $\alpha = \eta = 1.0$.

Considering individual values of λ , it is possible to compare our CVFs with the GSDs in Ref. [16] to estimate a value of α that compares most favorably with η . For $\lambda = 32$ —the largest λ used in Ref. [16]—a GSD with $\eta = 0.8$ approximates a CVF with $\alpha = 0.75$, $\eta = 0.6$ is similar to $\alpha = 0.5$, and $\eta = 0.4$ bears some resemblance to $\alpha = 0.2$, etc. For $\eta \lesssim 0.3$, the approximated CVFs are logarithmic or have $\alpha < 0$. In general, α values are smaller than the corresponding values of η . Moreover, the GSD tends to exhibit a higher particle volume fraction contained by small particles, while the CVF and GSD exhibit similar scaling behaviors as $D \rightarrow \lambda$. Despite this, the peak ϕ value obtained in this work is within $\sim 5\%$ of that in Ref. [16] and occurs at comparable α and η values, for $\lambda = 32$. Given the role of λ , the discrepancy between these results should reduce as the maximum particle size ratio increases.

Figure 3 also indicates the CVF exponent of the power-law particle size distribution corresponding to the random Apollonian packing, $\alpha_A \approx 0.53$ [18,24–27]. Interestingly, α_A is quite close to the corresponding α used to obtain the peak ϕ for $\lambda = 50$, $\alpha = 0.55$, and as noted above the peak trends toward smaller α with increasing λ (see also discussion in Ref. [18]). The power-law exponent of the Apollonian packing is conjectured to be the lower bound of exponents that result in full coverage obtained via geometric packing protocols [42]. The corresponding upper bound in 3D is $\beta = 4$ ($\alpha = 0$) [42]. Our $\lambda = 50$ data can be compared with an extrapolation of data from Ref. [27], which explored the physical fractal behavior of random Apollonian packings. The densest packing we obtained gave $\phi \approx 0.85$, a packing significantly less dense than the corresponding Apollonian packing $\phi_A \approx 0.93$. Apollonian packings are created using particle insertion methods, which circumvent physical constraints on particle motion and have not been tested for mechanical stability, while in DEM sim-

ulations particles cannot move freely through constrictions (pores) smaller than their diameters. Thus, in jammed configurations, DEM-generated microstructures are expected to contain larger pores, resulting in overall looser packings than traditional Apollonian states.

The presence of the peak at intermediate α in Fig. 3 suggests that obtaining the optimal packing density depends on balancing the abundance of the large-on-large particle contact pairs comprising the majority of the force-bearing backbone with the amount of small particles filling in the gaps between. This supposition was suggested by Furnas [3] and is qualitatively supported by the snapshots shown in Fig. 1(b) for $\alpha = 1.0$ and $\alpha = 0.2$. For the former, it is apparent that large particles regularly contact others of comparable size, but the relative scarcity of small particles available to populate the gaps results in a somewhat porous microstructure. For the latter, contacts between large particles are rare because they are embedded in a sea of small particles. Qualitatively similar behavior is observed in bidisperse packings, with high- (low-) α power-law disperse packings corresponding to low (high) fractions of small particles in the bidisperse case [9].

To quantify this observation, the joint probability distributions for contacting pairs of particles are computed as determined by the particle diameters, $P_c(D_1, D_2)$, as are the associated conditional contact probability distributions, $P_c(D_1, D_2)/P(D = D_1)$, for $\lambda = 50$. The results, ordered such that $D_1 \geq D_2$, are shown in Fig. 4 for several α . From Fig. 4(a), the most probable pair of sizes for contact for $\alpha = 0.9$ occurs when both particles have intermediate diameters $\sim 5 - 7$, close to the geometric mean of the maximum size ratio, $\sqrt{\lambda}$. Conversely, for $\alpha = 0.3$ most contacts exist between pairs of particles with diameters near the smallest value. Of course, since large particles are less abundant by construction, contacts between them make up a negligible fraction of the full set of contacts. Figure 4(b) shows the conditional contact probability distributions, which better account for large particle scarcity. Figure 4(b) shows that contacts including a particle with $D_1 \sim \lambda = 50$ are more common than might be otherwise expected (because these particles have the greatest surface area) and, further, that the diameter of its contact pair partner falls from $D_2 \sim \sqrt{\lambda}$ for high α to $D_2 \sim 1$ for low α . This result underscores the increasing importance of small particles in stabilizing the packing as α decreases. Minh and Cheng [22] pointed out that changes in large particle connectivity from being large-particle-dominated to being small-particle-dominated may reduce the propensity of large particles to fracture. If this hypothesis is correct, then our results indicate that low- α packings should be less susceptible to interparticle fracture than high- α packings.

Figure 4 does not show a strong signature distinguishing results for the densest packing obtained for $\lambda = 50$ ($\alpha = 0.55$) from results for other α . From Fig. 3, the densest packing is not obtained for contact probability distributions that are large-large (or large-intermediate) pair dominated (e.g., $\alpha = 0.9$ in Fig. 4) or large-small pair dominated (e.g., $\alpha = 0.3$ in Fig. 4), but rather for $\alpha = 0.55$, which from Fig. 4 lies somewhere in between these two extremes. This observation aligns with the behavior of bidisperse packings, which reach optimal density through a saturation of large-large and small-small contact pairs [3,9].

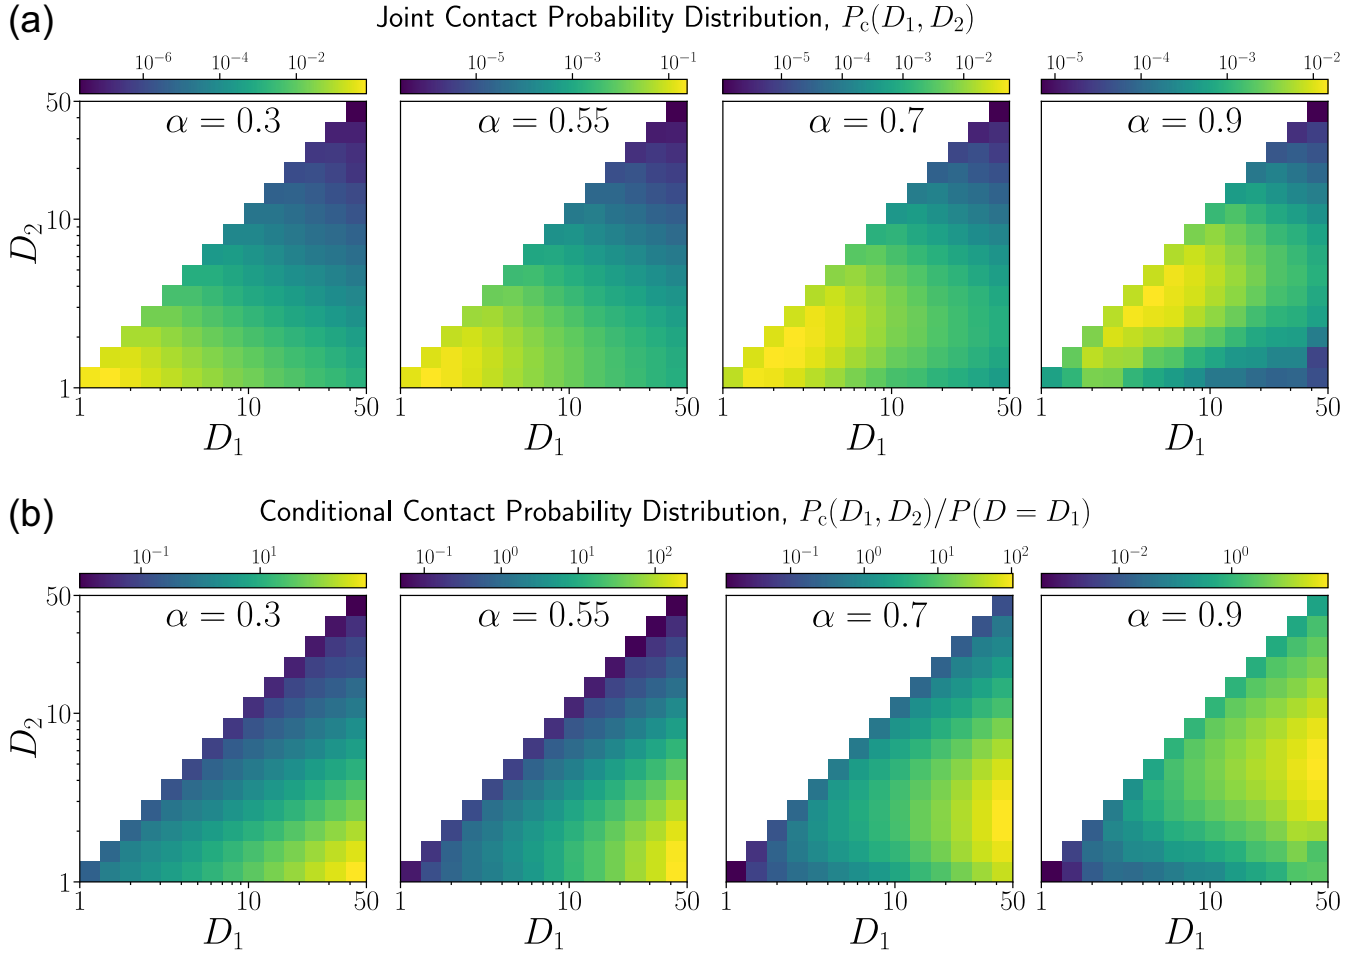


FIG. 4. (a) Joint contact probability distributions for particle diameter pairs with $D_1 \geq D_2$ for the indicated α and $\lambda = 50$. (b) Conditional contact probability distribution for the same set of systems.

In power-law disperse systems, the number of contacts for each particle depends upon its size, but the behavior of the mean number of contacts or coordination $\langle Z \rangle$ is less clear. As noted in Sec. II D, rattler particles strongly influence the calculation of $\langle Z \rangle$. However, the overall value of $\langle Z \rangle$ and a rattlers-excluded value $\langle Z \rangle_{\text{nr}}$ can be computed separately [43]. For the latter, rattler particles are identified by determining particles with fewer than three contacts and removed from the list of contact pairs. Note that each rattler removed decrements the total number of contacts for its contact pair partners, so the removal process is done iteratively [44].

After removing rattlers, the mean coordination is recomputed for the reduced contact list and the smaller set of nonrattler particles. The results of these analyses for all particles and only nonrattler particles are shown in Fig. 5(a) as a function of α and for two separate values of λ . Considering first the overall value of $\langle Z \rangle$, shown using filled symbols, the results are smaller than Z_{iso} for all α and both λ values. The largest $\langle Z \rangle$ is found for $\alpha = 0.2$, while $\langle Z \rangle$ is close to zero for $\alpha = 1.0$ for $\lambda = 50$. After removing rattlers, Fig. 5(a) shows that $\langle Z \rangle_{\text{nr}}$ is approximately equal to Z_{iso} for all α and λ (shown as open symbols). Since the value of p_a used in our simulations is small, the corresponding nonrattler value of $\Delta Z_{\text{nr}} \equiv \langle Z \rangle_{\text{nr}} - Z_{\text{iso}}$ is likewise small but nonzero: typical values found in this article are $\Delta Z_{\text{nr}} \approx 0.01 - 0.03$.

The results in Fig. 5(a) imply that the effect of rattler particles is significant for every system examined. An oppositely related quantity to $\langle Z \rangle$ is the rattler fraction, f_r , the fraction of all particles that are rattlers. Figure 5(b) quantifies how f_r increases with α . In particular, for $\alpha = 1.0$ and $\lambda = 50$ only a few percent of particles are nonrattlers; this system, with rattlers removed, is reproduced in Fig. 5(c). Rattler fractions of comparable magnitude were also observed in Ref. [16] for $\eta = \alpha = 1.0$. Indeed, the trends shown in Fig. 5 closely mirror those of Fig. 5 in Ref. [16] and seem to be fairly universal. Furthermore, the rattler fraction variation with α shown in Fig. 5(b) is reminiscent of the discontinuous jump in small particle rattler behavior observed at the Furnas peak in frictionless, bidisperse packings [3,9]. In Sec. III B, the implications of these results on the distribution of particles that participate in the mechanical stability of the packing will be examined.

While these results show that the mean nonrattler coordination is a constant, the exemplar packings depicted in Fig. 1(b) suggest that the number of contacts per particle, $Z(D)$, depends upon both particle diameter D and the underlying distribution of particle sizes (see Ref. [22] for a similar calculation for frictional particles). The quantity $Z(D)$ is calculated by binning particles by size and computing the mean number of contacts per particle in each bin. Rattler particles

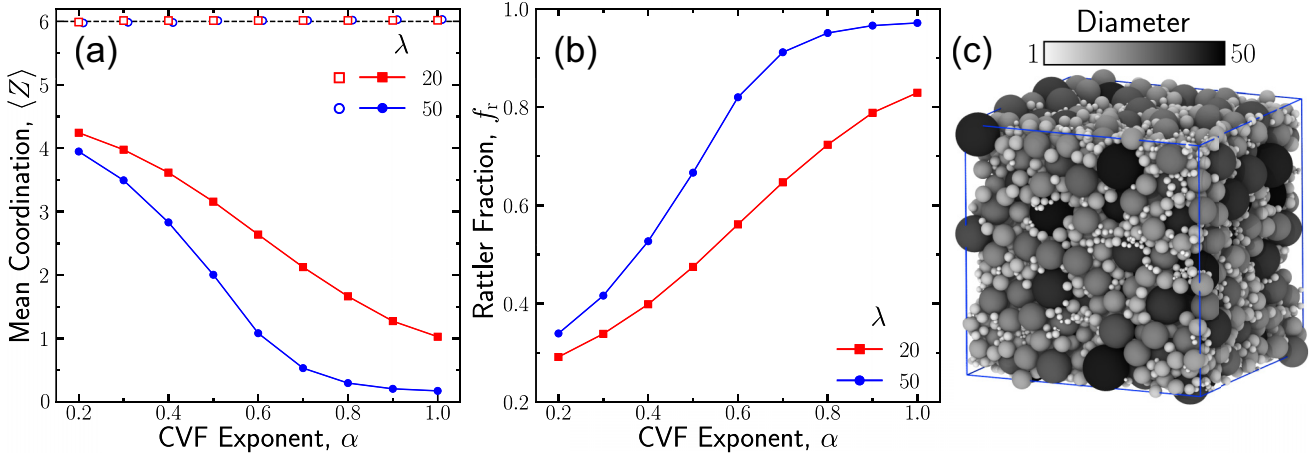


FIG. 5. (a) Mean coordination per particle including rattlers $\langle Z \rangle$ (filled symbols) and without rattlers $\langle Z \rangle_{nr}$ (open symbols) for the indicated values of the maximum particle size, λ . (b) Fraction of rattler particles for the same systems. (c) Reproduced $\alpha = 1.0$ packing from Fig. 1(b) after removing rattlers.

are excluded from this analysis to mitigate transient effects resulting from their short-lived participation in contact pairs. Results for $\lambda = 50$ are shown in Fig. 6. The figure demonstrates the scaling of $Z - 3$, i.e., the mean excess number of contacts over the three contacts required for mechanical stability, plotted against $(D/\lambda)^2$; plotted in this way, the relationship is linear and corresponding linear fits to data for large $(D/\lambda)^2$ were computed (dashed lines). Note that $(D/\lambda)^2 = 0.01$ corresponds to $D = 5$. The relationship $Z - 3 \sim D^2$ exhibits the same D^2 scaling with diameter as the particle surface area and represents a slightly faster scaling than was observed in Ref. [22] for smaller, frictional particles. It is clear from the

linear fits in Fig. 6 that the prefactor decreases with increasing α , an intuitive result given that fewer intermediate and large particles than small particles can be placed in the available solid angle of any central particle [10,45]. From Fig. 4(b), this exclusion of solid angle inherent to large-large contact pairs has strongest significance for high α , resulting in the lowest overall maximum per-particle contact count. Linear fits also worked for both $\lambda = 20$ and $\lambda = 100$ (not shown), though the prefactors generally depended upon λ for $\alpha \lesssim 0.8$.

The results in Fig. 6 are striking given that $\langle Z \rangle_{nr}$ is approximately six, while $Z(D)$ for the largest particles is at least an order of magnitude larger. The inset of Fig. 6 shows the fraction of nonrattler particles f_c participating in exactly c contacts, which is peaked at $c = 4$ and essentially independent of α . This low coordination value is responsible for the deviations away from linear scaling for small $(D/\lambda)^2$. Note that this analysis distinguishes between the discrete contact count c and the bin-averaged quantity $Z(D)$. Similar results for f_c were reported in Ref. [15] for smaller size dispersity and different underlying particle size distributions. The range $3 \leq c \leq 10$ encompasses between 90–95% of all nonrattler particles for each α but only accounts for roughly 80% of $\langle Z \rangle_{nr} = \sum_c c f_c$. The remaining contributions to $\langle Z \rangle_{nr}$ come from the high, but rare, contact participation counts of large particles.




FIG. 6. Mean excess number of contacts per particle than are required for mechanical stability ($Z = 3$) as a function of $(D/\lambda)^2$ for the indicated α and $\lambda = 50$. Rattler particles are omitted from these results. Dashed lines are linear fits to the symbols of matching color. Inset: The fraction of nonrattler particles f_c participating in exactly c contacts.

B. Nonrattler particle distributions

Rattler particles contribute to the overall volume fraction and density of the jammed packing, but have no bearing on its mechanical stability. Large particles are crucial to the force-bearing network, while sizable fractions of the small particles are rattlers. This suggests that the input particle size distribution differs from the distribution of particles in the force-bearing network. This section examines how the shape of the input particle size distribution $P(D)$ dictates the resultant distribution of nonrattler particles, $P_{nr}(D)$, and is motivated by considering the volume fraction contributed solely by nonrattler particles ϕ_{nr} , shown in Fig. 7. This




FIG. 7. Volume fraction ϕ_{nr} contributed exclusively by nonrattler particles for the indicated λ .

measure is akin to the mechanical void ratio in the geophysical literature [46,47]. In contrast to the clear dependence of the overall particle volume fraction ϕ on λ shown earlier in Fig. 3, Fig. 7 shows that ϕ_{nr} is independent of λ (for $\lambda \gtrsim 20$) at high α . While not shown here, results for $\alpha = 1.5$ and $\alpha = 2.0$ also collapsed for $\lambda \geq 8$. In this shallow power-law limit ($\beta \rightarrow 0$, $\alpha \rightarrow 4.0$), λ must become irrelevant. As α grows, ϕ_{nr} gradually approaches the equivalent monodisperse packing value with rattlers removed, ϕ_{nr}^{mono} ; for the constant-pressure protocol and averaged over five realizations each with 10^5 monodisperse particles, we obtained $\phi_{nr}^{\text{mono}} \approx 0.629$, compared to the overall packing density $\phi^{\text{mono}} \approx 0.639$.

Several interesting trends are apparent for $\alpha \rightarrow 0.3$ in Fig. 7. First, ϕ_{nr} levels off at low α , and second, the plateau values steadily increase with λ . The $\lambda \geq 20$ data in Fig. 7 suggest that increasing λ may shift the collapse of ϕ_{nr} to progressively smaller α . Since reducing α corresponds to increasing the relative abundance of small particles compared to large particles, a plateau in ϕ_{nr} implies that there may be diminishing returns to adding *more* small particles as most become rattlers. However, adding smaller and smaller particles, i.e., increasing λ , does lead to denser force-bearing networks in the plateau regime. It is interesting that while $\alpha \sim 0.55 - 0.6$ lead to the densest overall packings, the force-bearing components of such packings are less dense than those for smaller α . In most cases, the volume fraction lost when rattlers are removed, $\phi - \phi_{nr}$, is smaller than 0.1, with the largest shifts occurring for $\alpha = 1.0$. This is an indication that while rattlers may constitute a large fraction of the total number of particles, they typically only account for a small fraction of the total particle volume.

The removal of rattler particles permits the identification of the nonrattler particle size distributions P_{nr} , the distribution governing the force-bearing backbone. Results for P_{nr} obtained using $\lambda = 50$ are shown in Fig. 8. For each α , the tail of the distribution maintains its original power-law character, while the probability of retaining small particles is reduced, with substantial dependence on α . As $\alpha \rightarrow 1.0$, the most prob-




FIG. 8. Distributions of nonrattler particles for the indicated α and $\lambda = 50$. For each curve, the power-law behavior at large D is unchanged from the initial particle size distribution, with exponent $\beta = \alpha - 4$. The limiting cases of such power-law scalings are indicated.

able remaining particle diameters are close to $\sqrt{\lambda}$ [e.g., refer to the $\alpha = 1.0$ packing image in Fig. 5(c)], while for $\alpha \rightarrow 0.2$ it is clear that the smallest nonrattler particles remain the most probable. Indeed, aside from the exponent of the power-law tail, only a small amplitude change for $D \sim 1$ differentiates $\alpha = 0.2$ from $\alpha = 0.55$, the densest overall packing. The results shown in Fig. 8, taken together with the context given by Fig. 5(b) and Fig. 7, imply that high- α packings do not derive mechanical stability from small particles, while particles of all sizes are necessary to stabilize low- α packings.

The collapse of ϕ_{nr} with increasing λ for high α indicates that in such cases the underlying distributions of nonrattler particles should have similarities. Rather than considering the nonrattler particle size distributions themselves, this analysis focuses on the fraction of nonrattler particles $F(x)$ that are larger than x . This quantity has the advantages that it varies monotonically from 1 to 0, and Eq. (2) dictates how it should scale with D away from the endpoints. To compare data for different λ on an equal basis, the particle diameters are normalized by λ such that the scaled diameters fall in the domain $1/\lambda \leq x = D/\lambda \leq 1$. Results for $F(D/\lambda)$ are shown in Fig. 9 for several α and a wide range of λ . For each data set, $F(D/\lambda)$ is unity until the smallest nonrattler particle is encountered, beyond which $F(D/\lambda)$ drops to 0 in a manner that exhibits the expected power-law behavior over narrow ranges of D/λ and accelerates as $D/\lambda \rightarrow 1$. The power-law regime broadens as α increases, consistent with the results depicted in Fig. 8.

The most striking result shown in Fig. 9 is that $F(D/\lambda)$ is identical for all λ for $\alpha = 1.5$, signifying that the shape of the nonrattler particle size distribution is constant with respect to increases in λ beyond $\lambda = 8$. Similar results were obtained for larger α (not shown). However, as α falls to 1.0 and lower, data for the smallest λ increasingly deviate from the other curves until all data sets are clearly distinct for $\alpha \lesssim 0.7$. In cases where $F(D/\lambda)$ collapses, the constant

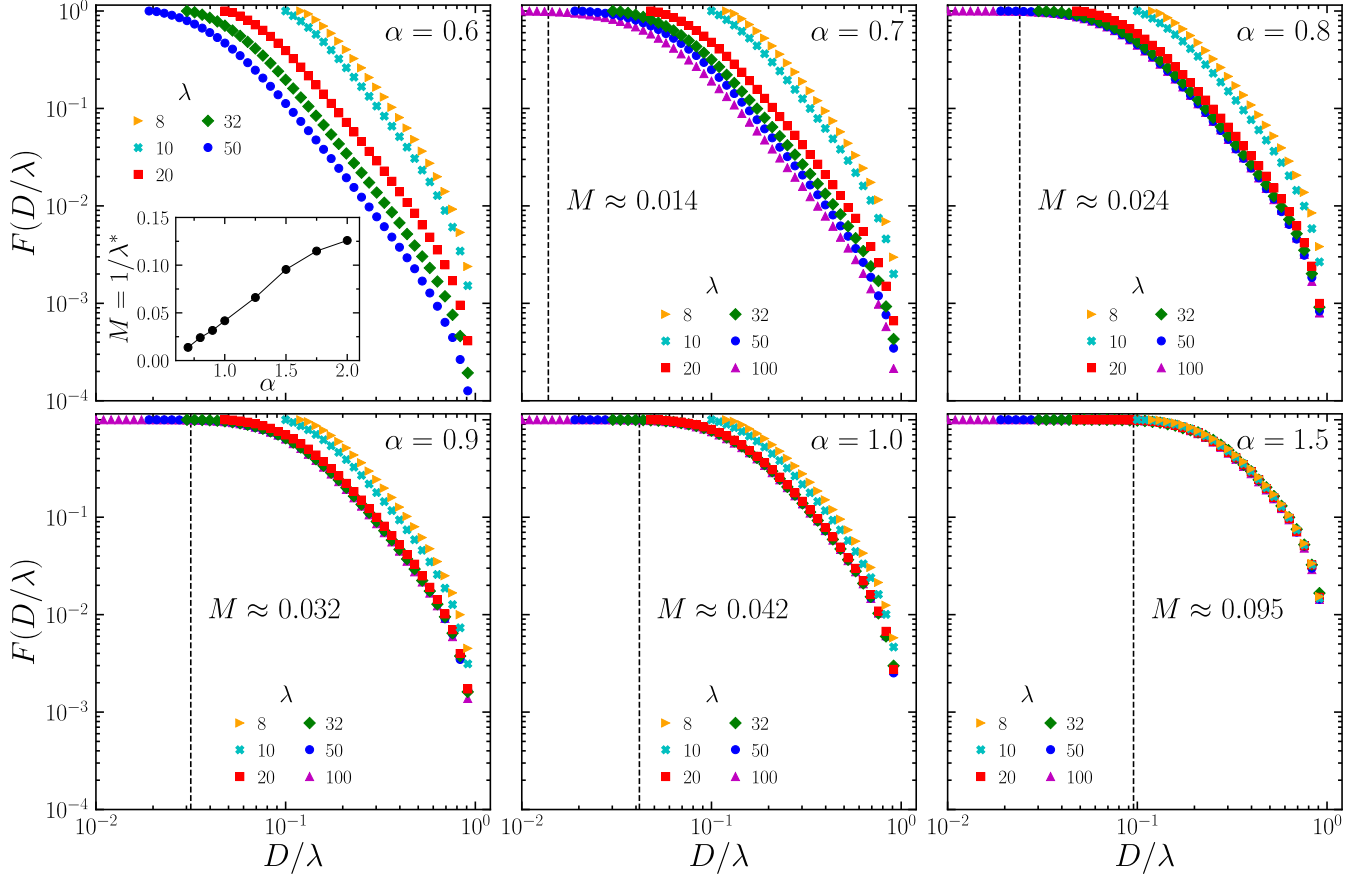


FIG. 9. Fraction of nonrattler particles with diameters larger than D/λ . The dashed black lines and printed labels M mark the value of D/λ for which $F(D/\lambda) = 0.99$ as estimated from the $\lambda = 100$ data. The inset to the upper left panel plots computed values of M for $0.7 \leq \alpha \leq 2.0$ that were obtained for $\lambda = 100$. Note that M was computed for several additional values of α for which we omit plots of $F(D/\lambda)$ data.

value of D/λ determining the onset of $F < 1$ means that there is a reduction of the effective width of the non-rattler particle size distribution, i.e., $\lambda \rightarrow \lambda^*$, with $\lambda^* \leq \lambda$. This results from the removal of rattler particles with diameters smaller than a threshold $D_{\min}^* > D_{\min} = 1$. In what follows, the onset value is referred to as the magnification, $M \equiv 1/\lambda^* = D_{\min}^*/\lambda$, which can be estimated for different α . Specifically, using the $\lambda = 100$ data shown in Fig. 9, M is determined by extracting the onset value of D/λ for which $F < 0.99$. Note that there is ambiguity in the precise determination of M based on the threshold F value—for example, for $\alpha = 1.5$, both the $\lambda = 8$ and $\lambda = 10$ data nearly collapse on top of the larger λ data despite being smaller than the nominal $\lambda^* \approx 10.5$. However, our tests showed that changes in the estimated magnifications are sharper for higher F thresholds and the estimates of M we obtained for our threshold choice are sufficient for the discussion here.

The estimated values of M are given in the corresponding panels of Fig. 9 for $\alpha \geq 0.7$ and indicated with dashed black lines. In addition, an inset plotting $M(\alpha)$ estimated from the $\lambda = 100$ data is shown in the $\alpha = 0.6$ panel, and includes data for several α that are not shown in Fig. 9. In essence, the definition of M permits us to determine a criterion, given by $1/\lambda \lesssim M$, for which the nonrattler particle size distribution is independent of λ . For systems satisfying this criterion, particles with $D < D_{\min}^*$ are almost always rattlers. Moreover,

D_{\min}^* replaces $D_{\min} = 1$ as the unit of length of the system; as an intensive quantity, ϕ_{nr} is also independent of λ when M is constant. Note that since the tail of the nonrattler particle size distribution is unchanged from the original power law, it is still possible to collapse F for each α , provided that a λ -dependent rescaling factor is used. In such cases, however, ϕ_{nr} does not collapse.

To help contextualize these results, the two separate limits of $M = 1$ and $M = 1/\lambda$ can be defined. The former is relevant for high- α systems that approach the limit of monodisperse systems, which is defined by a singular length scale set by the particle diameter. The latter occurs in cases where the smallest particles are necessary for ensuring mechanical stability of the packing. From the trend depicted in the inset of Fig. 9, the $M = 1$ limit is likely slowly approached for $\alpha > 2$.

The results in Fig. 9 show that the force-bearing component of the packing is invariant with respect to changes in λ provided that the α -dependent scale λ^* is exceeded. From this standpoint, no additional benefit is gained by adding successively smaller particles once $\lambda \geq \lambda^*$. However, the overall properties and structure of the packing, including the packing density, still depend on the full particle size distribution. Indeed, the relatively low density of the force-bearing component suggests that the rattler particles play a significant role in determining the final configuration, perhaps by restricting the intermediate configurations that the force-bearing

component can adopt. Intriguing avenues for possible future study include successive, repeated jamming and removal of rattler particles to isolate the limiting particle size distribution, and in designing particle size distributions for use in constructing the densest possible packings that can be obtained via compaction protocols. In addition, DEM simulations of frictional and/or cohesive particles with large size dispersity have not yet been systematically performed but are crucial to connecting simulation results with real-world applications.

IV. CONCLUSION

We performed large-scale 3D DEM simulations to study the packing properties of power-law disperse spherical particles. This article considered a wide range of power-law particle size distributions, varying the range of particle sizes and the exponents characterizing the power laws.

At fixed particle size ratio, the results showed that the densest overall packings were obtained for power-law particle size distributions that achieved mechanical stability while balancing contacts between pairs of large-large and large-intermediate particles with pairs of small-small and large-small particles. Distributions containing too many large particles do not need small particles for mechanical stability, and so most small particles were rattlers. Conversely, distributions with too many small particles produce packings that are dominated by large-small contact pairs, and so do not generate contacts between pairs of particles with diameters in the intermediate size classes. Further, despite the strong dependence of the mean coordination and rattler fraction on the CVF exponent, the mean coordination of nonrattler particles was close to the isostatic value, while the mean number of contacts per nonrattler particle scaled quadratically with particle diameter.

Considering only nonrattler particles, volume fractions of nonrattler particles for input distributions with high α were independent of size dispersity for λ larger than an α -dependent cutoff value λ^* , while for low α the nonrattler volume fraction was insensitive to α . In the former case, the fraction of nonrattler particles with normalized diameters larger than D/λ was independent of λ , provided that $\lambda \geq \lambda^*$. This result signifies a separate effective length scale of the force-bearing network, as determined by α . For the latter case, the results indicated that increasing the *proportion* of small particles has little effect on the force-bearing component of the packing, while adding *smaller* particles tended to improve both the overall and nonrattler packing densities.

The results presented here provide insight into the internal microstructure of large size dispersity particle packings and broaden our understanding of the relationship between features of the overall packing and its force-bearing backbone. In turn, the understanding gleaned from studying these systems may provide a pathway to optimizing the properties, mechanical and otherwise, of designed particle packings.

ACKNOWLEDGMENTS

The authors thank Andrew P. Santos for helpful discussions and Kevin Stratford, Tom Shire, and Kevin Hanley for

providing an initial implementation of their disperse neighboring technique in LAMMPS and Steve Plimpton and Axel Kholmeyer for their assistance in expanding this implementation and merging it into LAMMPS for public use. I.S. acknowledges support from the U.S. Department of Energy (DOE), Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics Program under Contract No. DE-AC02-05CH11231. This work was performed at the Center for Integrated Nanotechnologies, a U.S. DOE and Office of Basic Energy Sciences user facility. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy National Nuclear Security Administration under Contract No. DE-NA-0003525. The views expressed in the article do not necessarily represent the views of the U.S. DOE or the United States Government.

APPENDIX: MULTINEIGHBORING DETAILS AND PERFORMANCE

To simulate systems with large disparity ratios λ in particle sizes, we leveraged contact detection algorithms optimized for highly disperse systems. In this Appendix, we focus on dense, homogeneous systems and derive the leading order computational costs in λ for three methods. First, we describe a typical approach for monodisperse systems [28,29]. Next, we consider two additional methods designed for highly disperse systems: an older method by Veld *et al.* [30] and a more recent technique by Ogarko and Luding [31] and Krijgsman *et al.* [48] that was extended and ported to LAMMPS by Stratford *et al.* [38] and Shire *et al.* [32]. Our arguments will highlight a fundamental advantage of the newest method. A brief description of each approach is included although further details can be found in their respective references. Finally, we briefly describe changes in our implementation of the newest method in LAMMPS aimed at continuous particle size distributions and provide some benchmark data for the power-law-distributed systems studied in the main text. Before calculating contact forces, particle-based simulations often construct a Verlet neighbor list which contains all pairs of interacting particles [28,29]. In LAMMPS, a link-cell method is used where particles are spatially binned onto a grid with a bin size Δ in a process that takes $O(N_T)$ time, where N_T is the total number of particles. Particles use this binning to efficiently generate a list of potential neighbors consisting of particles within their own bin and in other nearby bins that are within the interaction distance. This set of bins that needs to be searched is known as a stencil. A distance is then only calculated between these candidate neighbors as opposed to all $\sim N_T^2$ pairs of particles in the system. In the default algorithm in LAMMPS, labeled *default*, Δ is approximately half of the maximum interaction distance such that the stencil only includes a small number of adjacent bins [39,49]. This method is very efficient for nearly monodisperse systems.

To illustrate how this algorithm fails at large λ , we consider a d -dimensional, bidisperse packing of N_L large particles and N_S small particles with diameters D_L and D_S , respectively, and a volume fraction of small particles $f_S \equiv N_S D_S^d / (N_S D_S^d +$




FIG. 10. Relative cost of constructing a neighbor list compared to a single force evaluation for 3D binary packings near jamming with $f_s = 0.5$ and the indicated values of λ and N_L using the *default* (green), *multi/old* (red), and *multi* (blue) algorithms. Tests were run on a single processor. The dashed line represents λ^3 scaling. Note that typically there are many force evaluations between neighbor list builds in LAMMPS and this reported value does not represent a typical ratio over many time steps. For perspective at $\lambda = 200$, building a single neighbor list took less than nine seconds using *multi* but over 72 hours using *multi/old*. Builds were impractically long to measure using the *default* method at $\lambda > 100$ or *multi/old* at $\lambda > 200$.

$N_L D_L^d$). As in the main text, one can treat $D_S = D_{\min} = 1$ such that $D_L/D_S = \lambda$. We mainly focus on a fixed volume fraction $0 < f_s < 1$ such that $N_S = N_L \lambda^d f_s / (1 - f_s)$ and $N_T \approx N_S \gg N_L$ at large λ , but other cases are briefly discussed at the end of the Appendix. In the *default* method, the bin size Δ is set by D_L such that each bin contains $N_{\text{particles/bin}} \sim D_L^d/D_S^d = \lambda^d$ particles and the stencil only contains a finite number $N_{\text{bins/stencil}}$ of nearby bins. Therefore, there are $N_T N_{\text{particles/bin}} N_{\text{bins/stencil}}$ candidate neighbors and calculating their distances is $O(N_L \lambda^{2d})$. This cost is greater than binning particles and dominates neighbor list construction for large λ .

To put this scaling in context, one can compare it to the total number of contacts in the system or the computational cost of calculating forces. In a jammed system, each small particle can only have a finite number of contacts independent of λ while each large particle can have up to $\sim D_L^{d-1}/D_S^{d-1} = \lambda^{d-1}$ contacts. Therefore, the total number of contacts in the system scales as $N_S + N_L \lambda^{d-1} \sim N_S \sim N_L \lambda^d$. This implies that the cost to build the neighbor list using the *default* method dominates the total simulation time and simulations become prohibitively expensive with increasing λ . This disparity is seen in Fig. 10 for $d = 3$ and $f_s = 0.5$, where the ratio of time to construct the neighbor list versus the time to calculate forces grows as λ^3 .

To reduce costs, an alternate algorithm was implemented in LAMMPS by in 't Veld *et al.* [30]. We refer to this algorithm by its current name in LAMMPS, *multi/old*. The *multi/old* method adjusts spatial binning based on a particle's type, a

categorization of particles that is used to set interactions parameters including the distance cutoff for non-DEM particles. In this approach, the size of a bin is set by the smallest interaction length D_S such that $N_{\text{particles/bin}}$ is constant and no longer grows with λ . Therefore, a different stencil is needed for each combination of particle types. These stencils extend out to the order of $D_L/D_S = \lambda$ bins for large-large pairs, $(D_L + D_S)/D_S \sim \lambda$ bins for large-small pairs, and a constant number of bins for small-small pairs. While this method reportedly accelerates highly disperse simulations up to a factor of 100 for $\lambda = 20$ [30], it does not address the fundamental scaling with λ due to the search for large-small neighbors. Each small particle searches $N_{\text{bins/stencil}} \sim (D_L/D_S)^d = \lambda^d$ bins within the large-small stencil to find potential large neighbors, where most bins will not contain a large particle, such that $O(N_S \lambda^d)$ or $O(N_L \lambda^{2d})$ operations are still performed. This scaling, with a smaller prefactor than *default*, is seen in Fig. 10. Practically, we find simulations become intractable around λ of order 10.

To address this shortcoming, an additional twist described in Refs. [31,38] uses a hierarchy of binning grids, one for each particle type in the initial LAMMPS implementation by Stratford *et al.* [38]. This method is referred to as *multi*, reflecting its current name in LAMMPS, and includes a separate binning grid for each particle type with a bin size Δ set by the same-type interaction distance. In a binary system, Δ is set by D_S for small particles and D_L for large particles such that $N_{\text{particles/bin}}$ does not depend on λ , similar to *multi/old*. The key difference is that each particle looks for same-type neighbors using its own set of bins while only small particles look for large neighbors using the large bins. Large particles do not search for small neighbors. Therefore, $N_{\text{bins/stencil}}$ is also independent of λ and construction costs are proportional to $N_T N_{\text{particles/bin}} N_{\text{bins/stencil}} \sim N_T \sim N_L \lambda^d$, equivalent to the force calculation. This scaling is demonstrated in Fig. 10 where the time to construct a neighbor list normalized by the time to calculate forces has no significant dependence on λ up to $\lambda = 300$ for *multi*. The only factor preventing simulations at larger λ was the growing number of particles in the system, already reaching $N_T = 27$ million at $\lambda = 300$.

Here, we considered the case of fixed f_s , although one could consider a value of f_s that grows or shrinks with λ . If f_s grew with λ approaching the limit of unity, then $N_T \sim N_S$ would grow faster than λ^d , e.g., as $\lambda^{d+\epsilon}$. For *multi*, $N_{\text{particles/bin}}$ and $N_{\text{bins/stencil}}$ would still both be finite such that neighbor costs would still grow as $N_T \sim \lambda^{d+\epsilon}$. This increase in cost would mirror an increase in cost to calculate forces simply due to having more particles. In contrast for *default*, $N_{\text{particles/bin}}$ would still scale as λ^d such that neighbor list construction would be $O(N_T \lambda^d)$ or $O(\lambda^{2d+\epsilon})$ and still dominate simulation time. In the opposite limit, one could consider a shrinking f_s exemplified by the extreme case of $N_S = 1$ where large particles effectively only have a finite number of contacts with other large particles. Again, $N_{\text{particles/bin}}$ and $N_{\text{bins/stencil}}$ would be finite in *multi* but now only large-large lookups would be relevant such that building a neighbor list would be $O(N_L)$, equivalent to calculating forces. For *default*, the one small particle would be irrelevant and building the neighbor list would resemble the $O(N_L)$ process for a purely monodisperse system, identical to the scaling of *multi* although with reduced overhead. Therefore, in all of these cases, the scaling of *multi*




FIG. 11. Time in seconds to build a single neighbor list per particle as a function of the number of collection intervals for jammed packings of power-law-distributed grains for the indicated λ . Calculations were performed on a single processor. Open symbols: $\beta = 3.0$ ($\alpha = 1.0$); filled symbols: $\beta = 3.6$ ($\alpha = 0.4$).

will always match that of a force evaluation and will either scale better than or equivalent to *default*, although prefactors depend on the specific system.

For continuous particle size distributions, increasing the number of collection intervals with increasing λ generally

improves performance as particles are binned using a bin size Δ closer to their actual diameter. These savings grow until the overhead of having additional collection intervals exceeds the benefit. This is seen in Fig. 11 where the cost of neighbor list construction is plotted as a function of the number of linearly spaced intervals for a variety of jammed systems with different power-law size distributions. The optimal number of bins and their spacing ultimately depends on the specific system. It is worth noting that as the optimal number of collection intervals continues to grow with λ , the computational costs of building a neighbor list could begin growing faster than the force evaluations with λ as more binning grids need to be created and searched. Practically, we have not yet found this to be a limitation, particularly since the optimal number of bins only reaches ~ 10 for the systems considered here. Similar studies on the optimization of power-law-distributed particle sizes were also performed in Ref. [48].

For this article, we expanded the original *multi*-implementation of Stratford *et al.* [38] to fully integrate it with LAMMPS and added the method to the public distribution of LAMMPS. As the particle type in LAMMPS is intended to represent material properties such as moduli or friction coefficients and not necessarily the size of a particle, we also generalized their implementation and provided the option for users to select an arbitrary set of particle size intervals or collections independent of particle types. This approach more closely reflects the original discussion of the method by Ogarko and Luding [31] and helps streamline optimization of neighbor list construction for continuous particle size distributions.

-
- [1] W. B. Fuller and S. E. Thompson, *Trans. Am. Soc. Civ. Eng.* **59**, 67 (1907).
- [2] A. H. M. Andreassen and J. Andersen, *Kolloid-Zeitschrift* **50**, 217 (1930).
- [3] C. C. Furnas, *Ind. Eng. Chem.* **23**, 1052 (1931).
- [4] D. L. Turcotte, *J. Geophys. Res.* **91**, 1921 (1986).
- [5] P. A. Langston, M. S. Nikitidis, U. Tüzün, D. M. Heyes, and N. M. Spyrou, *Powder Technol.* **94**, 59 (1997).
- [6] H. Liu, S. Dong, L. Tang, N. M. Krishnan, G. Sant, and M. Bauchy, *J. Mech. Phys. Solids* **122**, 555 (2019).
- [7] C. S. O’Hern, L. E. Silbert, A. J. Liu, and S. R. Nagel, *Phys. Rev. E* **68**, 011306 (2003).
- [8] R. S. Farr and R. D. Groot, *J. Chem. Phys.* **131**, 244104 (2009).
- [9] I. Srivastava, S. A. Roberts, J. T. Clemmer, L. E. Silbert, J. B. Lechman, and G. S. Grest, *Phys. Rev. Research* **3**, L032042 (2021).
- [10] M. Danisch, Y. Jin, and H. A. Makse, *Phys. Rev. E* **81**, 051303 (2010).
- [11] M. Hermes and M. Dijkstra, *Europhys. Lett.* **89**, 38005 (2010).
- [12] K. W. Desmond and E. R. Weeks, *Phys. Rev. E* **90**, 022204 (2014).
- [13] V. Baranau and U. Tallarek, *Soft Matter* **10**, 3826 (2014).
- [14] D. Cantor, E. Azéma, P. Sornay, and F. Radjai, *Phys. Rev. E* **98**, 052910 (2018).
- [15] P. Mutabaruka, M. Taiebat, Roland J.-M. Pellenq, and F. Radjai, *Phys. Rev. E* **100**, 042906 (2019).
- [16] W. F. Oquendo-Patiño and N. Estrada, *Granular Matter* **22**, 75 (2020).
- [17] W. F. Oquendo-Patiño and N. Estrada, *EPJ Web Conf.* **249**, 02003 (2021).
- [18] W. F. Oquendo-Patiño and N. Estrada, *Phys. Rev. E* **105**, 064901 (2022).
- [19] A. Herman, *Entropy* **15**, 4802 (2013).
- [20] R. R. Filgueira, L. L. Fournier, C. I. Cerisola, P. Gelati, and M. G. García, *Geoderma* **134**, 327 (2006).
- [21] O. Ben-Nun and I. Einav, *Philos. Trans. R. Soc. A* **368**, 231 (2010).
- [22] N. H. Minh and Y. P. Cheng, *Géotechnique* **63**, 44 (2013).
- [23] J. P. de Bono and G. R. McDowell, *Int. J. Solids Struct.* **187**, 133 (2020).
- [24] M. Borkovec, W. De Paris, and R. Peikert, *Fractals* **02**, 521 (1994).
- [25] S. V. Anishchik and N. N. Medvedev, *Phys. Rev. Lett.* **75**, 4314 (1995).
- [26] H. Herrmann, R. Mahmoodi Baram, and M. Wackenhut, *Phys. A* **330**, 77 (2003).
- [27] F. Varrato and G. Foffi, *Mol. Phys.* **109**, 2923 (2011).
- [28] S. Plimpton, *J. Comput. Phys.* **117**, 1 (1995).
- [29] A. P. Thompson, H. M. Aktulga, R. Berger, D. S. Bolintineanu, W. M. Brown, P. S. Crozier, P. J. in ’t Veld, A. Kohlmeyer, S. G. Moore, T. D. Nguyen, R. Shan, M. J. Stevens, J. Tranchida, C. Trott, and S. J. Plimpton, *Comput. Phys. Commun.* **271**, 108171 (2022).

- [30] P. J. in 't Veld, S. J. Plimpton, and G. S. Grest, *Comput. Phys. Commun.* **179**, 320 (2008).
- [31] V. Ogarko and S. Luding, *Comput. Phys. Commun.* **183**, 931 (2012).
- [32] T. Shire, K. J. Hanley, and K. Stratford, *Comp. Part. Mech.* **8**, 653 (2021).
- [33] S. Dagois-Bohy, B. P. Tighe, J. Simon, S. Henkes, and M. van Hecke, *Phys. Rev. Lett.* **109**, 095703 (2012).
- [34] K. C. Smith, I. Srivastava, T. S. Fisher, and M. Alam, *Phys. Rev. E* **89**, 042203 (2014).
- [35] A. P. Santos, D. S. Bolintineanu, G. S. Grest, J. B. Lechman, S. J. Plimpton, I. Srivastava, and L. E. Silbert, *Phys. Rev. E* **102**, 032903 (2020).
- [36] L. E. Silbert, D. Ertaş, G. S. Grest, T. C. Halsey, D. Levine, and S. J. Plimpton, *Phys. Rev. E* **64**, 051302 (2001).
- [37] A. Stukowski, *Modell. Simul. Mater. Sci. Eng.* **18**, 015012 (2010).
- [38] K. Stratford, T. Shire, and K. Hanley, *Implementation of multi-level contact detection in LAMMPS*, Tech. Rep. No. eCSE12-09 (University of Edinburgh, United Kingdom, 2018)
- [39] These capabilities are described in the documentation found at <https://docs.lammps.org/neighbor.html>. An example input script `in.powerlaw` is included with the LAMMPS distribution in the `examples/multi` subdirectory. It is a demo of shearing a 2D packing of particles with power-law-distributed sizes using neighbor list options similar to those leveraged in this paper.
- [40] A. J. Liu and S. R. Nagel, *Annu. Rev. Condens. Matter Phys.* **1**, 347 (2010).
- [41] I. Prasad, C. Santangelo, and G. Grason, *Phys. Rev. E* **96**, 052905 (2017).
- [42] T. Aste, *Phys. Rev. E* **53**, 2571 (1996).
- [43] J.-N. Roux, *Phys. Rev. E* **61**, 6802 (2000).
- [44] A. Donev, S. Torquato, F. H. Stillinger, and R. Connelly, *J. Comput. Phys.* **197**, 139 (2004).
- [45] E. I. Corwin, M. Clusel, A. O. N. Siemens, and J. Brujić, *Soft Matter* **6**, 2949 (2010).
- [46] M. Otsubo, Particle scale analysis of soil stiffness and elastic wave propagation, Ph.D. thesis, Imperial College London, 2016.
- [47] D. Liu, C. O'Sullivan, and J. A. H. Carraro, *J. Geotech. Geoenviron. Eng.* **147**, 04020182 (2021).
- [48] D. Krijgsman, V. Ogarko, and S. Luding, *Comp. Part. Mech.* **1**, 357 (2014).
- [49] Some aspects of neighbor list construction in LAMMPS, such as the addition of a skin to the interaction distance and the existence of full and half stencils, are not described in this Appendix as they do not affect the underlying scaling of costs with λ .