Population annealing simulations of a binary hard-sphere mixture

Jared Callaham and Jonathan Machta

Population Annealing Simulations of a Binary Hard Sphere Mixture

Jared Callaham[1, *] and Jonathan Machta[1, 2, †]

[1]*Department of Physics, University of Massachusetts, Amherst, Massachusetts 01003 USA*
[2]*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico 87501 USA*

Population annealing is a sequential Monte Carlo scheme well-suited to simulating equilibrium states of systems with rough energy landscapes. Here we use population annealing to study a binary mixture of hard spheres. Population annealing is a parallel version of simulated annealing with an extra resampling step that ensures that a population of replicas of the system represents the equilibrium ensemble at every packing fraction in an annealing schedule. The algorithm and its equilibration properties are described and results are presented for a glass-forming fluid composed of a 50/50 mixture of hard spheres with diameter ratio of 1.4:1. For this system, we obtain precise results for the equation of state in the glassy regime up to packing fractions $\varphi \approx 0.60$ and study deviations from the BMCSL equation of state. For higher packing fractions, the algorithm falls out of equilibrium and a free volume fit predicts jamming at packing fraction $\varphi \approx 0.667$. We conclude that population annealing is an effective tool for studying equilibrium glassy fluids and the jamming transition.

PACS numbers:

I. INTRODUCTION

One of the grand challenges of computational statistical physics is to understand the nature of glassy systems. Profound open questions remain concerning both configurational glasses and spin glasses. The signature property of glassy systems is an extreme slowing of dynamics in both laboratory and computational experiments. The equilibrium properties of glassy systems are thus very difficult to study. The situation is somewhat better for spin glasses where it is established that there is thermodynamic glass transition and a well-understand mean field theory of the low temperature glass phase, though controversies surround the nature of the low temperature phase in finite dimensions. For configurational glasses it is not even known whether there is an underlying thermodynamic transition or whether the glass transition is entirely a kinetic phenomenon.

Much of the progress for finite-dimensional spin glasses has been made possible due to algorithmic advances, particularly the introduction of replica exchange Monte Carlo, also known as parallel tempering [1–5]. The situation for simulating equilibrium glassy fluids is less well-developed and fundamentally a more difficult problem. Parallel tempering has also been extensively applied to fluid systems [6] in the glassy regime [7, 8]. Other recent algorithmic advances such as event chain Monte Carlo [9, 10] and particle-swap Monte Carlo [11] have also shown promises.

In this paper we introduce population annealing Monte Carlo as a method for studying fluid systems and show that it is an effective tool for studying equilibrium properties up to high densities in the glassy regime. Popula-

tion annealing was first developed for spin glasses [12–14] and has been shown to be an efficient method for large-scale studies of equilibrium states [15, 16] and ground states [17] of spin glasses. Here we use population annealing Monte Carlo to simulate a binary mixture of hard spheres.

In particular, we study a glass-forming 50/50 binary mixture of hard spheres with a diameter ratio of 1.4:1, which has been the subject of previous computational studies, e.g. [8, 18, 19]. Using population annealing we obtain precise estimates of the equation of state up to packing fractions in the glassy regime. We compare our results to previous simulations and the well-known Boublik-Mansoori-Carnahan-Starling-Leland equation of state [20, 21].

The paper is organized as follows. In Sec. II we introduce population annealing and describe its properties. In Sec. III we describe the hard sphere fluid model, the simulation methods and the observables that we study. Results for the equation of state, the approach to jamming and the performance of the algorithm algorithm are presented in Sec. IV. The paper concludes with a discussion.

II. POPULATION ANNEALING MONTE CARLO

A. Overview

Population annealing (PA) is closely related to simulated annealing. In both algorithms, a system is taken through an annealing schedule in one or more thermodynamic control parameters, e.g. temperature or density, from a region where equilibration is easy, e.g. high temperature or low density, to a region where equilibration is difficult. At each step in the annealing schedule, an equilibrating procedure, e.g. the Metropolis algorithm,

---

*Electronic address: `jcallaham@physics.umass.edu`
†Electronic address: `machta@physics.umass.edu`

is applied to the system at the current value of the control parameter(s). The objective of simulated annealing is to find ground states or low-lying states and simulated annealing does not sample from the equilibrium distribution at each step along the annealing schedule. Population annealing, by contrast, does sample from the equilibrium distribution at each step along the annealing schedule. In PA a large population of replicas of the system is annealed and, at each step in the annealing schedule, the population is resampled to maintain equilibrium as described in the next section.

B. Population Annealing for Hard Spheres

In this section we describe population annealing for hard sphere fluid systems using packing fraction φ as the control parameter. A population of R independent replicas of the system is initially prepared at some low (or zero) packing fraction, φ_0. In the simulations reported below $R = 10^6$ and $\varphi_0 = 0$. In parallel, each replica is taken through an annealing schedule, which is a sequence of increasing packing fractions $\{\varphi_0, \varphi_1, \ldots, \varphi_K\}$. Each annealing step $(\varphi_i \rightarrow \varphi_{i+1})$ consists of two parts. In the first part, an equilibrating procedure is applied at the initial density, φ_i. In our study, the equilibrating procedure is event chain Monte Carlo (ECMC) [9] (see Sec. III B) but other methods such as molecular dynamics or Metropolis Monte Carlo would also be suitable. The second part of the annealing step is to increase the packing fraction from φ_i to φ_{i+1} holding the relative positions of the particles fixed. Typically some replicas in the population will now suffer overlaps between spheres and therefore have disallowed configurations. These disallowed replicas are removed from the population. Suppose that ϵR replicas are culled from the population. The population size is restored to R by randomly choosing ϵR replicas from among the surviving $(1 - \epsilon)R$ replicas and making one copy of each of them. For an observable $\mathcal{O}$, the PA estimator, $\tilde{\mathcal{O}}$ for the ensemble average of an observable, $\langle \mathcal{O} \rangle$ at each packing fraction is obtained from an average over the population at that packing fraction. The population can be thought of as an approximate, finite realization of the statistical ensemble for hard spheres. Note that it is possible that there are *no* allowed configurations at the new packing fraction. In this case, the algorithm must be terminated without producing results for higher packing fractions.

In the limit of large R, if the original population is an unbiased sample from the equilibrium hard sphere ensemble at packing fraction φ_i then the new population will also be an unbiased sample from the equilibrium hard sphere ensemble at packing fraction φ_{i+1}. However, the new population is now correlated due to the copying of replicas and, for finite R, the new population may be biased due to configurations that are important at the higher density but too rare to be represented in the population at the lower density. The equilibrating procedure

at the new density partially corrects these problems. An analysis of the error in PA due to finite population size is given in Sec. II E.

Population annealing may be implemented with a fixed annealing schedule however we have found it more convenient to use an adaptive annealing schedule in which the fraction culled in each step is fixed. In our simulations the culling fraction is $\epsilon = 0.05$. In this implementation, the annealing schedule is a random list is a random list of packing fractions. However, in practice spacing between successive values of φ is very small so it is straightforward to interpolate to obtain observables at any packing fraction. For the adaptive annealing schedule, in principle, the algorithm will never terminate but it may jam in the sense of taking smaller and smaller steps converging to a maximal or jammed density.

Population annealing for hard spheres in the NVT ensemble is particularly simple because each allowed hard sphere configuration has the same probability in the equilibrium ensemble so the resampling step requires no reweighting or Boltzmann factors. It would also be possible to use the NPT ensemble (with fixed temperature), in which case a Boltzmann factor $\exp[-\beta(p_{i+1} - p_i)V_r]$ would be required when deciding how many copies to make of replica r in the step $(p_i \rightarrow p_{i+1})$ where the volume of replica r at pressure p_i is V_r.

C. Configurational entropy estimator

One desirable feature of PA is that it gives direct access to thermodynamic potentials. For the hard sphere version of the algorithm described here, we have direct access to the configurational entropy $S_c(\varphi)$ as a function of packing fraction. From the basic definition of entropy, we have $S_c = \log \Omega$ where Ω is the statistical weight or dimensionless volume in configuration space of accessible configurations and the units are chosen so that Boltzmann's constant is unity. The factor $1 - \epsilon$ is an estimator of the ratio of the statistical weight before and after the culling, Ω_{i+1}/Ω_i, during the annealing step $(\varphi_i \rightarrow \varphi_{i+1})$. Thus we have the following expression for the change in the estimator $\tilde{S}_c$ in one annealing step,

$$\tilde{S}_c(\varphi_{i+1}) - \tilde{S}_c(\varphi_i) \approx \log(1 - \epsilon). \tag{1}$$

Summing this estimate over the annealing schedule up to step k, we obtain the entropy at packing fraction φ_k in terms of the entropy at the initial packing fraction and the number of annealing steps taken,

$$\tilde{S}_c(\varphi_k) = \tilde{S}_c(\varphi_0) - k \log(1 - \epsilon). \tag{2}$$

The equation of state can be obtained by differentiating with respect to φ. These estimates become exact in the $R \rightarrow \infty$ limit. Errors are discussed in Sec. II E.

D. Weighted Averaging

A very useful feature of PA is the ability to improve both statistical and systematic errors by combining many independent runs using weighted averaging. If the population is perfectly equilibrated and there are no systematic errors then the most efficient way to reduce statistical errors in an observable is to perform ordinary, unweighted averaging over the values of the observable in each run. However, for finite R, each run is not completely equilibrated and there are systematic errors. Unweighted averaging suppresses statistical errors but does not suppress systematic errors. However, using weighted averaging we can also suppress systematic errors. Suppose we have measured both the configurational entropy estimator, $\tilde{S}_c$ and the estimator of an observable $\tilde{\mathcal{O}}$ at a given value of φ in each of M independent runs, all with the same population size. The weighted average for the observable, $\overline{\mathcal{O}}$ is

$$\overline{\mathcal{O}} = \frac{\sum_{m=1}^{M} \tilde{\mathcal{O}}^{(m)} \exp[\tilde{S}_c^{(m)}]}{\sum_{m=1}^{M} \exp[\tilde{S}_c^{(m)}]}, \tag{3}$$

where m indexes the independent runs. The observable $\mathcal{O}$ must be a quantity that can be measured in a single replica at a single packing fraction. Examples include pressure and various correlation functions but not overlaps between replicas or the entropy itself.

We make the following claim: Weighted averaging, as defined in Eq. (3), yields an exact, unbiased result for fixed population size R, and fixed annealing schedule $\{\varphi_0, \varphi_1, \ldots, \varphi_K\}$, in the limit of infinitely many runs, $M \to \infty$. The same conclusion holds for the adaptive annealing schedule used here and is discussed below.

The validity of weighted averaging is trivial for $R = 1$, (simulated annealing). For a given φ_i, a run either survives to that packing fraction or is terminated at a lower packing fraction. Surviving runs all have the same weight since there is no resampling, while terminated runs have zero weight. Each surviving run is an unbiased sample taken from the hard sphere ensemble. Since the weighting factor is either 0 or 1, the weighted average is a simple average over the surviving runs. The $R = 1$ version of the algorithm lacks resampling and is therefore highly inefficient at high densities.

The validity of weighted averaging for arbitrary R can be established using an inductive argument. The populations of the runs at the initial packing fraction φ_0 are assumed to be equilibrated so that unweighted averaging is appropriate at φ_0. Correspondingly, the weight factors are the same and the claim is trivially valid for φ_0. Now suppose that the weighted averaging claim holds at packing fraction φ_i with weights $w_i^{(m)}$. In the first part of the resampling step from φ_i to φ_{i+1}, the random fraction $\epsilon_i^{(m)}$ of replicas with hard sphere overlaps are culled from the population. Consider the situation before copying is done to restore the population size to R. Each surviving replica in run m should still have the weight $w_i^{(m)}$ if one

averages over all replicas in all runs. However, what we actually do is to first average the observable within each run and then average over runs. Therefore weight of each run must be adjusted to reflect its new population size,

$$w_{i+1}^{(m)} \propto (1 - \epsilon_i^{(m)}) w_i^{(m)}, \tag{4}$$

where the constant of proportionality is set by the requirement that the weights are normalized. Comparing to Eq. (1), we see that if $w_i^{(m)} \propto \exp[\tilde{S}_c^{(m)}(\varphi_i)]$ then $w_{i+1}^{(m)} \propto \exp[\tilde{S}_c^{(m)}(\varphi_{i+1})]$ and the inductive hypothesis is verified. The remainder of the annealing step $(\varphi_i \to \varphi_{i+1})$ consists of randomly copying replicas to restore the population size to R and then applying the equilibrating procedure to each replica. The copying step yields additional fluctuations but does not change the expectation of the observable for each run. The equilibrating procedure may change the expectation of the observable for a given run if the individual runs were initially out of equilibrium. Nonetheless, the correctness of weighted averaging is preserved because the equilibrating procedure and weighted averaging both converge to the same hard sphere distribution.

Weighted averaging for the entropy itself differs from other observables because the entropy is obtained from a sum over all packing fractions. Nonetheless, the final result for the weighted average of the entropy $\overline{S_c}$ takes a simple form,

$$\overline{S_c} = \log \frac{1}{M} \sum_{m=1}^{M} \exp[\tilde{S}_c^{(m)}]. \tag{5}$$

This result is an example of the Jarzynski equality [22]. A derivation for PA in the the context of free energy rather than the entropy can be found in Ref. [13].

Finally, in the adaptive step algorithm used here, the annealing schedule is a random list so each run visits a distinct set of packing fractions whereas it is necessary to carry out weighted averages at fixed values of φ. To accomplish this we interpolate both the observables and the entropy between packing fractions and then use the weighted averaging formulas, Eqns. (3) or (5).

E. Equilibration and Systematic Errors

How do we know whether PA is yielding equilibrium values of observables? As discussed above, the resampling step for finite population size introduces systematic errors because the distribution at the new packing fraction is not fully sampled. The exactness of weighted averaging gives a way to quantify the deviations from equilibrium. The analysis of systematic errors is discussed in detail in Ref. [13] and is briefly reviewed here.

The expected value of an observable from a single run at population size R of PA is the *unweighted* average over runs while the the exact value of the observable is

the *weighted* average (both averages taken in the limit of infinitely many runs). Thus the systematic error in an observable, $\Delta\mathcal{O}$ is given by $\Delta\mathcal{O} = \langle\tilde{\mathcal{O}}\rangle - \overline{\mathcal{O}}$ where $\langle\tilde{\mathcal{O}}\rangle$ represents an unweighted average over independent runs of the algorithm. The difference between weighted and unweighted averages depends on the variance of the weighting factor in Eq. (3). For the case that the joint distribution of $\tilde{\mathcal{O}}$ and $\tilde{S}_c$ is a bivariate Gaussian, it was shown in Ref. [14] that systematic errors in measuring $\mathcal{O}$ are given by the covariance of $\tilde{\mathcal{O}}$ and $\tilde{S}_c$,

$$\Delta\mathcal{O} = \text{cov}(\tilde{\mathcal{O}}, \tilde{S}_c) = \text{var}(\tilde{S}_c)\left[\frac{\text{cov}(\tilde{\mathcal{O}}, \tilde{S}_c)}{\text{var}(\tilde{S}_c)}\right]. \quad (6)$$

The second, trivial identity in Eq. 6 is useful because the ratio in the square brackets goes to a constant that depends on $\mathcal{O}$ as $R \to \infty$ while $\text{var}(\tilde{S}_c)$ diminishes as $1/R$ and is independent of the observable. The expression on the far right in Eq. (6) motivates defining an equilibration population size ρ_f,

$$\rho_f = \lim_{R\to\infty} R \, \text{var}(\tilde{S}_c). \quad (7)$$

Systematic errors in all observables are proportional to ρ_f/R and PA simulations are well-equilibrated when $R \gg \rho_f$. In population annealing ρ_f plays the same role as the exponential autocorrelation time in Markov chain Monte Carlo methods.

For sufficiently large R, $\tilde{S}_c$ and other observables arise from many independent, additive contributions so non-rigorous "central limit theorem" arguments suggest that the joint distribution of entropy and any observable should be a bi-variate Gaussian. The approach to Gaussianity with R was investigated in Ref. [14]. If R is too small the joint distribution will not be Gaussian and there will be corrections to Eq. (6) containing higher cumulants of the joint distribution. We shall see later that for the fixed population size used here, the joint distribution of pressure and entropy is a bivariate Gaussian for small packing fractions but significant exponential tails appear at high packing fractions. These corrections and their meaning are explored in Sec. IV A

In our simulations we use weighted averaging extensively. What are the systematic errors in weighted averages? The same arguments that shows that $\text{var}(\tilde{S}_c)$ controls systematic errors in a single run demonstrate that the relevant measure of equilibration for a weighted average is $\text{var}(\overline{S_c})$, where $\overline{S_c}$ is the weighted average estimator of the configurational entropy defined in Eq. (5). From $\text{var}(\overline{S_c})$ we define the weighted average equilibration population size $\rho_f^*(R)$ that is a function of the fixed population size R used in the weighted average,

$$\rho_f^*(R) = \lim_{M\to\infty} MR \, \text{var}(\overline{S_c}). \quad (8)$$

Since weighted averaging using M runs of population size R is less efficient than doing a single large run with population size MR, we have that $\rho_f^*(R) \geq \rho_f$ and we also expect that $\lim_{R\to\infty} \rho_f^*(R) \to \rho_f$.

Of course, it is not practical to compute $\text{var}(\overline{S_c})$ from repeated experiments each with M runs. Instead we use a bootstraps procedure to estimate $\text{var}(\overline{S_c})$ by resampling with replacement from our M runs. Each sample of size M is employed to compute $\overline{S_c}$ from Eq. (5) and the variance is compute by resampling many times.

It is known that in the thermodynamic limit and at sufficiently high packing fraction the equilibrium state of the binary mixture studied here is phase separated with separate crystals of the small and large spheres [23]. For a finite number of spheres the dominant configurations in the equilibrium distribution are not known since the entropic cost of the interface between the two crystalline regions may prevent phase separation. It may be that the equilibrium states are dominated either by phase separated configurations or other non-random configurations. We have not seen evidence of either phase separation or other forms of ordering in our PA simulations.

F. Statistical Errors

If an observable $\mathcal{O}$ is averaged over R independent measurements, the statistical error $\delta\mathcal{O}$ in the mean is given by $\sqrt{\text{var}(\mathcal{O})/R}$. The resampling step in PA introduces correlations so that statistical errors are larger than $\sqrt{\text{var}(\mathcal{O})/R}$. As discussed in detail in Ref. [14], we can bound statistical errors by ignoring the de-correlating effects of the equilibrating procedure. If there is no equilibrating procedure, each descendent of an initial replica is the same. We refer to the set of descendents of an initial replica as a *family* and define $\mathfrak{n}_i$ to be the fraction of the population in family i. In the absence of the equilibrating procedure $\mathcal{O}$ takes a single value for every member of the family, call this value $\mathcal{O}_i$. Given this assumption,

$$\tilde{\mathcal{O}} = \sum_i \mathcal{O}_i \mathfrak{n}_i. \quad (9)$$

To proceed we make two additional assumptions: (1) the observable $\mathcal{O}_i$ and family fraction $\mathfrak{n}_i$ are uncorrelated and (2) the distribution of family fraction $\mathfrak{n}_i$ has small fluctuations from run to run for large R. Given these assumptions, the variance of the observable is bounded by,

$$\text{var}(\tilde{\mathcal{O}}) \leq \text{var}(\mathcal{O})\sum_i \mathfrak{n}_i^2. \quad (10)$$

Assuming that the second moment of the family fraction scales as $1/R$ we define the *mean square family size*, ρ_t,

$$\rho_t = \lim_{R\to\infty} R \sum_i \mathfrak{n}_i^2. \quad (11)$$

The bound on the statistical error in $\delta\tilde{\mathcal{O}}$ becomes

$$\delta\tilde{\mathcal{O}} \leq \sqrt{\text{var}(\mathcal{O})\rho_t/R}. \quad (12)$$

Note that if there is no decimation so that $\mathfrak{n}_i = 1/R$, then $\rho_t = 1$ and Eq. (12) (as an equality) reduces to

the expression for statistical errors for uncorrelated measurements. Since ρ_t provides only a bound on statistical errors, in practice we use bootstrapping to estimate errors.

In Ref. [24] it is shown that ρ_f and ρ_t are close to one another. We investigate the relation between these two measures in Sec. IV A. Since measuring ρ_f requires multiple runs while ρ_t can be measured from a single run, ρ_t can serve as a practical measure of equilibration.

It is important to note that in the regime that $R \gg \rho_f$, systematic errors diminish as $1/R$ while statistical errors diminish as $1/\sqrt{R}$ so that, generally, systematic errors are much smaller than statistical errors.

G. Comparison to Parallel Tempering

Population annealing is a sequential Monte Carlo method [25] in contrast to most simulation methods in statistical physics, which are Markov chain Monte Carlo methods. In sequential Monte Carlo, equilibration is approached as the population size R increases while for Markov Chain Monte Carlo, equilibration is approached as the number of Monte Carlo sweeps is increased. Among Markov chain Monte Carlo methods parallel tempering shares the greatest similarity with population annealing. It is comparably efficient to parallel tempering but, as we have seen, it has some interesting advantages–it gives direct access to thermodynamic potentials such as entropy or free energy, it is amenable to massive parallelization and multiple runs can be combined to improve both statistical *and* systematic errors. Parallel tempering explores the disconnected minima in a rough free energy landscapes by annealing to low temperatures multiple times, using correctly designed swap moves to insure that each region of the free energy landscape is visited with the statistical weight determined by the Gibbs distribution. Population annealing explores rough free energy landscapes in a single annealing run using a large number of replicas to populate the disconnected minima in the landscape. Resampling insures that each minimum contains a fraction of the population given by the Gibbs distribution.

For parallel tempering and other Markov chain Monte Carlo methods, systematic and statistical errors can be determined from autocorrelation functions of observables. Systematic errors are related to the exponential autocorrelation time while statistical errors are related to the integrated autocorrelation time. Thus the two characteristic population sizes, ρ_f and ρ_t, play analogous roles for PA to the exponential and integrated autocorrelation times, respectively.

III. MODEL AND METHODS

A. Hard Sphere Model

Hard spheres have long offered a simple and useful model for studying the properties of fluids. Although these systems have been studied for decades, there are still mysteries surrounding some of their behaviors. Monodisperse systems exhibit a distinct entropy-driven first-order phase transition from the disordered liquid to a crystalline solid, but a metastable fluid branch persists beyond the transition and its high-density behavior is still not fully understood. It is as yet unclear to what packing fraction this metastable branch persists [7, 11, 18, 26], and there has also been evidence for [18, 27] and against [8, 11] the existence of a thermodynamic glass transition.

One difficulty in studying the high-density behavior of the metastable fluid is that, since the true thermodynamic state is a solid, full equilibration will lead to crystallization. Size polydispersity is often introduced to avoid crystallization. We study 50/50 mixture of spheres with diameter ratio 1.4:1, which does not easily crystallize [18, 23, 28].

At low and moderate densities the equation of state of a binary mixture of hard spheres is accurately described by the Boublík-Mansoori-Carnahan-Starling-Leland (BMCSL) equation[20, 21]:

$$Z(\varphi) = \frac{(1 + \varphi + \varphi^2) - 3\varphi(y_1 + y_2\varphi) - y_3\varphi^3}{(1 - \varphi)^3} \quad (13)$$

The constants y_1, y_2, and y_3 depend on the choice of polydispersity and for the 50/50 binary mixture with polydispersity ratio 1.4:1 they are given by, $y_1 = 0.0513$, $y_2 = 0.0237$ and $y_3 = 0.9251$.

B. Event Chain Monte Carlo

Population annealing requires an equilibrating procedure performed at each annealing step. We use event chain Monte Carlo (ECMC), which has been shown to be highly efficient for simulating 2D and 3D hard sphere systems [9, 10, 29–32].

In an ECMC step, a particle is chosen at random and moved in a given direction. The particle moves until it collides with some other particle (the "event"). The original particle remains at the point of collision, while the particle that was struck then moves in the same direction. This process continues until the total displacement length of the chain of particles reaches a predetermined chain length ℓ_c. The chain length is a parameter that can be adjusted, e.g. to minimize correlation times.

There are several variants of ECMC, but the simplest and fastest [9] version has only two directions of motion, $+x$ or $+y$. This method, called "x-y straight event chain," violates detailed balance, but preserves global

balance and thus converges to the correct equilibrium state [9].

C.   Population Annealing Pseudo-Code

The following is pseudocode for our implementation of population annealing with event chain Monte Carlo for a polydisperse hard sphere system. The diameters of the small and large particles are fixed respectively at 1 and 1.4.

1. Initialize each of R replicas at zero packing fraction by choosing N points at random within a unit cube simulation cell with periodic boundary conditions. Each particle is also labeled with a size according the prescription for polydispersity.

2. For each replica, determine $d_{\min}$, the minimum relative center-to-center distance between any two particles as shown in Fig. 1:

$$d_{\min} = \min_{i,j} \left[\frac{2\,|\mathbf{r_i} - \mathbf{r_j}|}{\sigma_i + \sigma_j} \right], \qquad (14)$$

where $\mathbf{r_i}$ is the position of particle i and σ_i is its diameter. The dimensions of this system and particle positions may be rescaled by $1/d_{\min}$ without causing overlaps. Note that for the first annealing step this rescaling increases the size of the system, and for all later steps the size is decreased.

3. Sort the list of $d_{\min}$ values and define the $(\epsilon R)^{\text{th}}$ smallest value of $d_{\min}$ as d^*.

4. Rescale the simulation box size L, packing fraction φ and position $\mathbf{r}$ of each particle according to $L \leftarrow L/d^*$, $\varphi \leftarrow \varphi/(d^*)^3$ and $\mathbf{r} \leftarrow \mathbf{r}/d^*$. (Note that φ is initially zero and needs to be calculated after the first annealing step).

5. After rescaling there are ϵR replicas with invalid configurations [33]. These replicas are eliminated from the population. The same number of replicas are randomly chosen (with replacement) from among the valid replicas and copies of these replicas are added to the population. After this step, the population consists of R valid configurations.

6. Perform one ECMC sweep on each replica. A sweep should displace on average N particles. We choose the chain length ℓ_c to displace an average of $\sqrt{N}$ particles (See Appendix A) and then perform $\sqrt{N}$ ECMC moves.

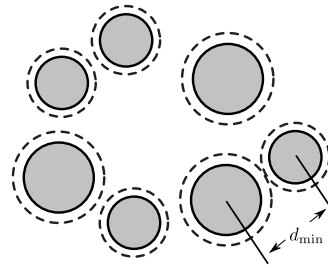7. Repeat steps 2-6 on each replica until a predefined packing fraction is reached.

FIG. 1: The minimum relative center-to-center distance over all particle pairs in a replica is $d_{\min}$.

D.   Observables

1.   Entropy

The relative configurational entropy $\tilde{S}_c$ at each step in the annealing schedule is obtained from Eq. (2). The simulations are initialized at $\varphi_0 = 0$ and the entropy at zero packing fraction is taken to be zero so that all entropy values are negative. We would like to know the weighted average value of the entropy $\bar{S}_c$ as a function of φ. However, our annealing schedule traverses a fixed set of entropies and the packing fraction at each entropy is variable. Values of the entropy at each packing fraction are obtained by interpolation and Eq. (5) is applied to obtain the weighted average entropy.

2.   Equation of State

The dimensionless equation of state, Z is expressed in terms of the pressure P,

$$Z = \frac{\pi\sigma^3}{6} \frac{P}{\varphi}, \qquad (15)$$

where σ is the average cubed sphere diameter, which for a binary mixture with diameters σ_A and σ_B is given by $\sigma^3 = \frac{1}{2}(\sigma_A^3 + \sigma_B^3)$. We obtain Z at each annealing step and then use interpolation to determine Z at a fixed set of packing fractions. We report the weighted average equation of state $\bar{Z}$ as a function of packing fraction. Error bars are obtained by bootstrapping the weighted averages. Two independent ways of measuring the pressure are described in the next two subsections.

3.   Thermodynamic Pressure Estimator

Using Eq. (2), pressure and equation of state can be found through thermodynamic definition,

$$\frac{P}{T} = \frac{\partial S}{\partial V}. \qquad (16)$$

In our simulations the change in entropy for each annealing step is a constant $\log(1 - \epsilon)$ and the increment

in packing fraction is variable. Thus a thermodynamic estimator for the pressure at annealing step k is

$$\frac{P_k}{T} = -\frac{\varphi_k}{V_k}\frac{\Delta S}{\Delta \varphi_k} = -\frac{\varphi_k}{V_k}\frac{\log(1-\epsilon)}{\Delta \varphi_k}. \qquad (17)$$

Setting $T = 1$,

$$Z = -\frac{\pi\sigma^3}{6V_k}\frac{\log(1-\epsilon)}{\Delta \varphi_k} \qquad (18)$$

where V_k is the volume at step k and $\Delta\varphi_k = (\varphi_{k+1} - \varphi_{k-1})/2$ is the symmetric difference.

4. Dynamic Pressure Estimator

Event chain Monte Carlo offers a direct method of calculating the equation of state, as described in Ref. [30]. Consider an event chain in the $+x$ direction. At each collision between particles j and k, the distance between the centers of the particles projected on the direction of motion is $x_k - x_j$. Define the "lifted" distance of an event chain, $x_{\text{final}} - x_{\text{initial}}$ as

$$x_{\text{final}} - x_{\text{initial}} = \ell_c + \sum_{(k,j)}(x_k - x_j) \qquad (19)$$

where the sum is over all of the collision events in the chain. The equation of state is then given by an average of the lifted distance over all event chains [30],

$$Z = \left\langle \frac{x_{\text{final}} - x_{\text{initial}}}{\ell_c} \right\rangle_{\text{chains}}. \qquad (20)$$

The results we present below for the equation of state are from this dynamic, ECMC estimator.

We note that Eq. (20) is not a correct measure of Z when $\ell_c > L$. When using the dynamic event chain length described in the Appendix, Eq. (20) cannot be used for $\varphi \lesssim 0.15$

E. ρ_f, ρ_f^* and ρ_t

In order to assess equilibration we computed the population size measures ρ_f, ρ_f^*, and ρ_t as described in Secs. II E and II F. The mean square family size ρ_t is straightforward to measure in each run with Eq. (11) by keeping track of the family to which each replica belongs. These estimates can also be combined through a weighted average across independent runs.

To estimate the equilibrium population size ρ_f through Eq. (7), however, we need the variance of entropy over many independent runs. But because we use a variable annealing schedule, we do not have direct access to the variance of the entropy as a function of packing fraction. Instead we have variable packing fractions at each value

of the entropy and $\text{var}(\varphi)$ as a function of S_c. By considering the S_c vs. φ curves from each run as independent random functions we can make the first-order estimate

$$\text{var}(\tilde{S}_c) = \left(\frac{\partial S_c}{\partial \varphi}\right)^2 \text{var}(\varphi), \qquad (21)$$

with the same symmetric derivative estimator for $\partial S_c/\partial\varphi$ as in Eq. (17).

The situation is simpler for the weighted average equilibrium population size ρ_f^*, because the weighted average entropy $\bar{S}_c$ is calculated at fixed packing fractions, so $\text{var}(\bar{S}_c)$ can be calculated directly by bootstrapping and Eq. (8) gives ρ_f^*.

IV. RESULTS

In this section we present the results of our simulations. Our results address two objectives. The first objective is to study the properties of the population annealing algorithm as applied, for the first time, to fluid systems and determine whether the algorithm can produce equilibrated results in the glassy regime of bidisperse hard spheres. The second goal is to obtain precise equilibrated results for the equation of state at high densities and to study nonequilibrium behavior at even higher densities near jamming for the 50/50 binary mixture of hard spheres with diameter ratio 1.4:1. We study systems with $N = 60$ and $N = 100$ particles. For each of these cases we performed $M \approx 1000$ independent runs, each with population size $R = 10^6$. Table I shows the simulation parameters. Note that the population size (10^6) and the number ECMC sweeps per annealing step (one) are the same for both $N = 60$ and 100 particles but the number of annealing steps required to keep the culling fraction fixed increases linearly in the number of particles. This linear scaling is the result of the fact that the probability of an overlap for a given fractional compression increasing linearly in the increase in packing fraction. Together with the linear scaling of carrying out a single sweep of ECMC, we see that the naive complexity of the algorithm is $O(N^2)$. Of course, this scaling does not take into account how computational resources must scale with N to achieve equilibration. The question of equilibration is discussed in the following subsection.

A. Equilibration of Population Annealing

Figure 2 shows the characteristic population sizes ρ_f and ρ_t vs. packing fraction. These quantities are discussed in Secs. II E and II F, respectively, and they characterize the errors in population annealing. Specifically, when ρ_f is small compared to the population size, systematic errors are small and each run is well-equilibrated. Similarly, when ρ_t is small compared with the population size, statistical error are small. Figure 2 demonstrates

TABLE I: Simulation parameters for the results reported in Sec. IV. $\varphi_{\max}$ is the maximum packing fraction. Average wall clock time for a run corresponds to an OpenMP implementation with 32 threads.

System Size N	60	100
Population Size R	10^6	10^6
Max packing fraction $\varphi_{\max}$	0.63	0.63
Culling fraction ϵ	0.05	0.05
Avg # Annealing Steps	11120	18800
Average Time (hrs)	22	80
Memory Usage (Mb)	1536	2452
Number of Runs M	1069	986
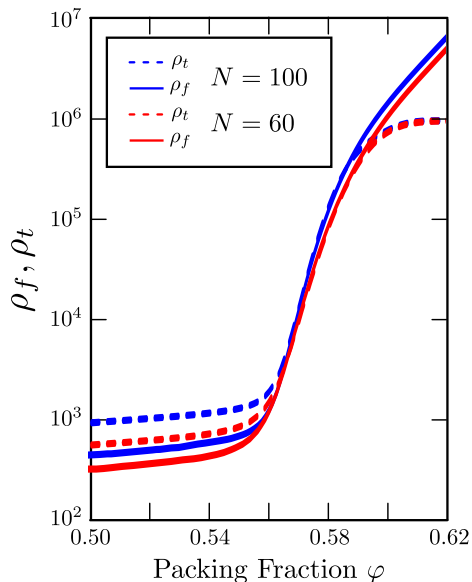
FIG. 2: (color online) Characteristic population sizes, ρ_f (solid lines) and ρ_t (dotted lines) vs. packing fraction φ for sizes $N = 60$ (red) and $N = 100$ (blue). See Secs. II E and II F for definitions.

that both measures are relatively flat and much smaller than the population size ($R = 10^6$) until $\varphi = 0.55$, but then both ρ_f and ρ_t grow quickly, increasing by about a factor of 100 between $\varphi = 0.55$ and 0.58. If we impose a relatively conservative equilibration criterion that $R > 100\rho_f$ then individual runs are found to be well-equilibrated for $\varphi < 0.575$ for both system sizes.

Note that ρ_t is approximately a factor of 2 larger than ρ_f at low density but in the region $0.56 \lesssim \varphi \lesssim 0.59$, the two quantities track each other closely (see Ref. [24]). Finally, for $\varphi \approx 0.59$ we find that $\rho_f \approx \rho_t \approx R$. Beyond this packing fraction, individual runs are clearly not in equilibrium. Since ρ_t is estimated from runs at population size R, it bounded by R while ρ_f is not bounded so the two quantities part ways above $\varphi \gtrsim 0.59$.

Considering the ratios of ρ_f for the two system sizes it seems clear that in the glassy regime $\varphi \geq 0.58$ it is more difficult to equilibrate larger systems. With only two

TABLE II: Measured quantities at selected packing fractions, φ for $N = 60$ spheres. $\bar{S}_c$ is the weighted average entropy defined in Eq. (5) and $\bar{Z}$ is the weighted average equation of state with one standard deviation errors. μ, σ, and λ are the parameters of the best fit exponentially modified Gaussian distribution (see text above Eq. (23)). A dash in the λ column indicates that a Gaussian distribution is the preferred fit. ρ_f and ρ_f^* are the equilibration population sizes for single runs and weighted averages defined in Eqs. (7) and (8), respectively.

φ	$\bar{S}_c$	$\bar{Z}$	μ	σ	λ	ρ_f	ρ_f^*
0.54	-343.71	16.2233(2)	-343.70	0.021	-	4.5×10^2	4.5×10^2
0.56	-379.66	18.832(1)	-379.70	0.034	-	1.2×10^3	1.2×10^3
0.58	-420.44	22.04(3)	-420.70	0.08	4.36	8.2×10^4	3.2×10^5
0.59	-442.99	23.99(14)	-443.90	0.16	1.88	3.7×10^5	8.2×10^6
0.60	-467.33	26.5	-469.60	0.29	1.04	1.1×10^6	9.4×10^7

TABLE III: Measured quantities at selected packing fractions, φ for $N = 100$ spheres. See Table II for details.

φ	$\bar{S}_c$	$\bar{Z}$	μ	σ	λ	ρ_f	ρ_f^*
0.54	-575.97	16.3455(1)	-576.00	0.025	-	6.1×10^2	6.1×10^2
0.56	-636.40	18.9956(7)	-636.40	0.035	-	1.2×10^3	1.2×10^3
0.58	-705.04	22.27(4)	-705.40	0.10	3.68	1.2×10^5	1.6×10^6
0.59	-743.19	24.51(6)	-744.40	0.21	1.52	5.3×10^5	1.6×10^7
0.60	-785.18	27.6	-787.80	0.40	0.88	1.4×10^6	5.8×10^7

system sizes, we cannot determine the scaling behavior of ρ_f with N.

Results from weighted averaging are equilibrated to higher densities than results from a single run. Figure 3 shows $\rho_f^*(R)$ the weighted average equilibration population size, together with ρ_f, as a function packing fraction. If we apply the conservative equilibration criterion that $MR \geq 100\rho_f^*(R)$ then results from weighted averaging are well-equilibrated for $\varphi < 0.587$ for $N = 100$ hard spheres and for $\varphi < 0.591$ for $N = 60$ hard spheres. Note that we believe that our results for the fluid equation of state remain reasonably accurate to somewhat higher packing fractions than these strict cut-offs.

A deeper understanding of weighted averaging can be gained by looking at the joint distribution of the entropy and equation of state estimators, $\tilde{S}_c$ and $\tilde{Z}$, as a function of packing fraction. Figure 4 displays these joint distributions as scatter plots for $N = 60$ particles and several packing fractions. Each point in these figures represents a single run of PA. The horizontal coordinate of the point is the relative configurational entropy estimator $\tilde{S}_c$ and the vertical coordinate is the equation of state estimator $\tilde{Z}$ for the run. The x and y coordinates of the red squares are the weighted average values $\bar{S}_c$ and $\bar{Z}$, respectively. Two important features of these plots are immediately evident: (1) there is an approximately inverse correlation between $\tilde{S}_c$ and $\tilde{Z}$ so that the weighted average value of the pressure is less than the ordinary average and (2) there are many outliers with large entropies and
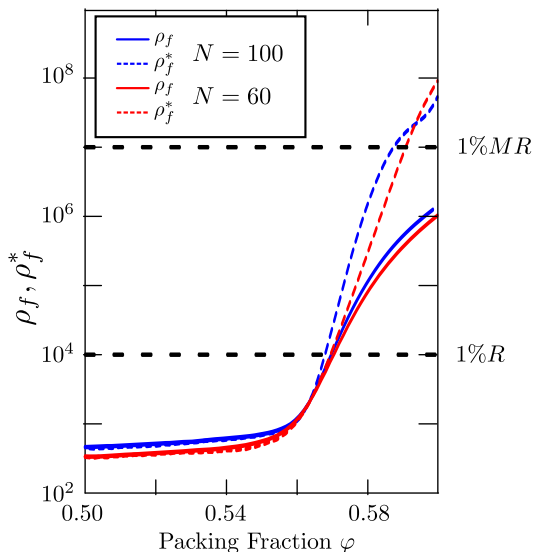
FIG. 3: (color online) Equilibration population sizes, ρ_f (solid lines) and ρ_f^* (dotted lines) vs. packing fraction φ for sizes $N = 60$ (red) and $N = 100$ (blue). See Sec. II E for definitions. Dotted horizontal lines at 10^4 and 10^7 show where the cut-offs for equilibration occur for single runs and weighted averages, respectively, based on the requirements $0.01R > \rho_f$ for single runs and $0.01MR > \rho_f^*$ for weighted averages.

small pressures and the distributions are clearly not bi-variate Gaussians. For $\varphi = 0.56$ and 0.58, these outliers have little effect on the weighted averages but for the $\varphi = 0.6$ and 0.62 they dominate the weighted averages. For packing fractions less that 0.55, the distributions are well-described by bi-variate Gaussians.

The tail of the entropy distribution determines the role of outliers in weighted averages. Figure 5 shows logarithmic plots of the entropy estimator $\tilde{S}_c$ distributions (blue curves) for $\varphi = 0.58$, and 0.60. The distributions have been smoothed with Gaussian kernels using Mathematica's *KernelMixtureDistribution*. It is clear that these distributions have exponential tails for large entropy rather than Gaussian tails. A good fit to these distributions, shown as dotted curves in Fig. 5 , is obtained using an exponentially modified Gaussian, which is defined as follow: Let X be Gaussian distributed and Y be exponentially distributed, then $Z = X + Y$ is described by an exponentially modified Gaussian distribution. An exponentially modified Gaussian requires three parameters: μ and σ are the mean and standard deviation of the Gaussian distribution, respectively, and λ is the characteristic decay rate of the exponential distribution. The probability density function is given by

$$p(x; \mu, \sigma, \lambda) = \frac{\lambda}{2} e^{\frac{1}{2}\lambda\left(\lambda\sigma^2 - 2x + 2\mu\right)} \mathrm{erfc}\left(\frac{\lambda\sigma^2 - x + \mu}{\sqrt{2}\sigma}\right).$$

(22)

The best fit values of the parameters for several values of φ are shown in Tables II and III, for $N = 60$ and 100 particles, respectively.
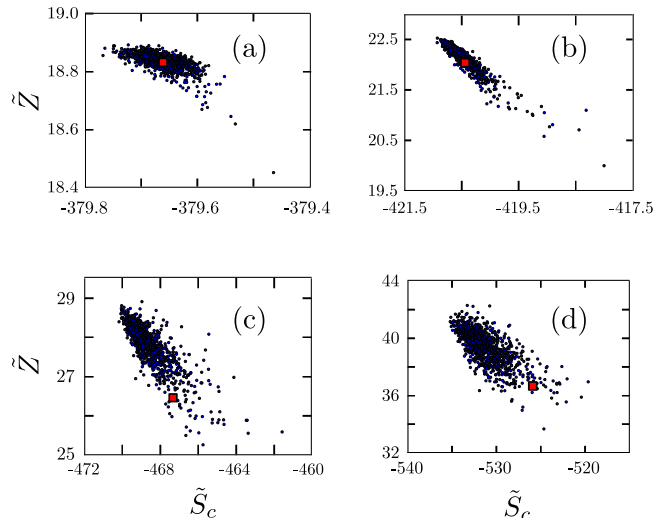

FIG. 4: (color online) Joint distribution of the entropy and equation of state estimators for $N = 60$ and several packing fractions φ. Each point represents a single run of population annealing. The x-coordinate of the point is the entropy estimator, $\tilde{S}_c$ and the y-coordinate is the equation of state estimator, $\tilde{Z}$. The weighted average values are shown as red squares.
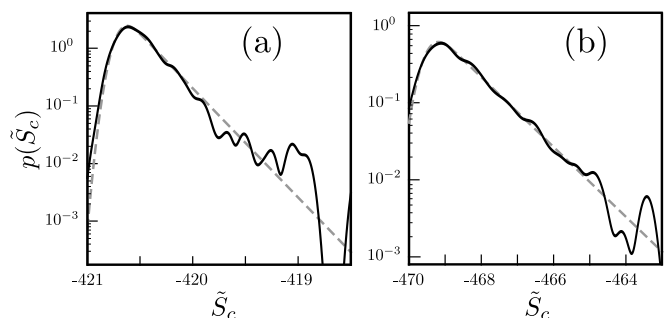

FIG. 5: Distributions of the configurational entropy estimator $\tilde{S}_c$ (solid curves) at packing fractions $\varphi = 0.58$ and $\varphi = 0.60$ together with a fit to an exponentially modified Gaussian distribution (dotted curves) for $N = 60$ particles. See Table II for fitting parameters.

If the entropy estimator is described by an exponentially modified Gaussian, the weighted average entropy $\bar{S}_c$, defined in Eq. (5), approaches, as $M \to \infty$, the integral of the distribution and the weighting factor,

$$\bar{S}_c = \log \int dx\, p(x; \mu, \sigma, \lambda) e^x. \qquad (23)$$

Carrying out the integral we find,

$$\bar{S}_c \to \mu + \frac{\sigma^2}{2} + \log\left(\frac{\lambda}{\lambda - 1}\right). \qquad (24)$$

When the exponential decay parameter λ is significantly larger than one, the effect of the exponential tail is small and the weighted average is a good approximation for the entropy and other observables so long as the number of trials is sufficiently large to explore the tail, as is the case here with $M \approx 1000$. However, if $\lambda \lesssim 1$ then the effect of the tail on the weighted average diverges and features of the entropy distribution that have not yet been explored, for example, a cut-off to the exponential tail, control the equilibrium averages. In the high packing fraction region where, $\lambda \lesssim 1$, the population annealing results do not yield useful information about the equilibrium properties of the system. From Tables II and III we see that for $N = 60$, $\lambda = 1$ is reached at about $\varphi = 0.60$ while for $N = 100$, $\lambda = 1$ is reached between 0.59 and 0.60. These estimates for where equilibration breaks down are consistent with the more conservative estimates based on ρ_f^*.

B.   Equation of State

Figure 6 shows the weighted average equation of state $\bar{Z}$ as a function of packing fraction φ for systems with $N = 60$ and $N = 100$ particles. The weighted average is performed using $M \approx 1000$ independent trials (see Table I), each with population size $R = 10^6$. We estimated the statistical errors by randomly resampling with replacement (bootstrapping) the collection of trials and recalculating the weighted average for the resampled trials. We note that this method makes use of the joint distribution of entropy and pressure to estimate systematic errors in the equation of state. The shaded regions around the weighted average curves represent 95% confidence intervals. If the original weighted average is dominated by a small number of trials with significantly greater entropies, the bootstrapped collection may leave out these trials, resulting in a highly skewed confidence interval. This effect is clearly seen in Fig. 6 at $\varphi \approx 0.61$ where the simulations have clearly fallen out of equilibrium.

The solid line in Fig. 6 is the BMCSL equation of state (see Eq. (13)), which is reasonably accurate for low densities but significantly underestimates Z for high densities. Of course, the BMCSL equation predicts finite pressure up to $\varphi = 1$ so that deviations of this kind are inevitable but it is not clear where they become significant. Figure
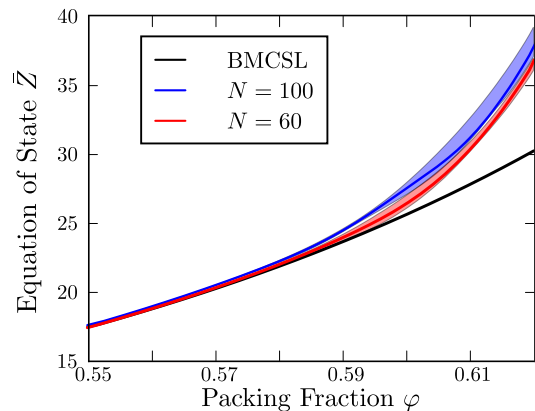

FIG. 6: (color online) The weighted average equation of state, $\bar{Z}$ with bootstrapped 95% confidence intervals for $N = 60$ (red) and $N = 100$ (blue) particles. The solid black line is the BMCSL equation of state.
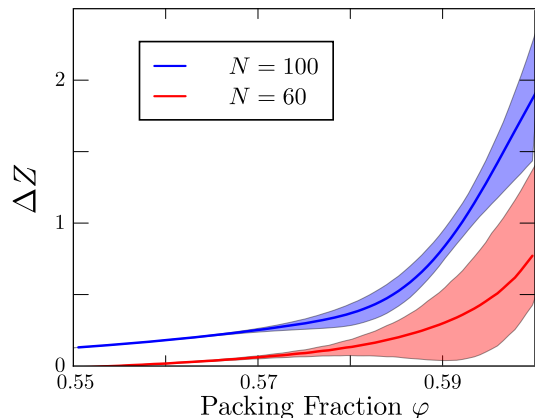

FIG. 7: (color online) Deviations ΔZ between the simulation results and the BMCSL equation of state as a function of packing fraction φ with bootstrapped 95% confidence intervals for $N = 60$ (red) and $N = 100$ (blue) particles.

7 shows the difference $\Delta Z = \bar{Z} - Z_{CS}$ between the equation of state measured in the simulations and the BMCSL prediction. For $N = 60$, significant deviations begin at $\varphi = 0.56$ while for $N = 100$ small deviations persist to much lower packing fractions. Large differences from BMCSL begin at $\varphi = 0.58$.

Since we are interested in the (metastable) equilibrium behavior of the fluid and not possible ordered states, we checked for ordering with the pair correlation function $g(r)$. The pair correlation function is shown in Fig. 8 for $N = 100$ particles at packing fraction $\varphi = 0.58$. The pair correlation function for $N = 60$ particles is indistinguishable from $g(r)$ for $N = 100$ particles. We see sharp peaks only at the three contact distances for this binary mixture. If the system typically formed phase-separated crystals or other ordered states we would expect additional sharp structure in $g(r)$. In addition, Fig. 8 is obtained from a weighted average over runs but we observed no noticeable difference between between the
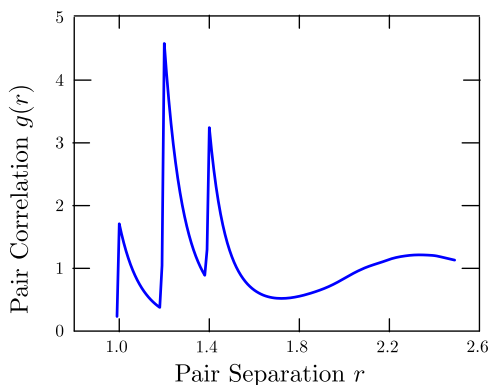
FIG. 8: Pair correlation function $g(r)$ for $N = 100$ particles at packing fraction $\varphi = 0.58$ as a function of separation r. The three peaks correspond to the separations at contact for this binary mixture.
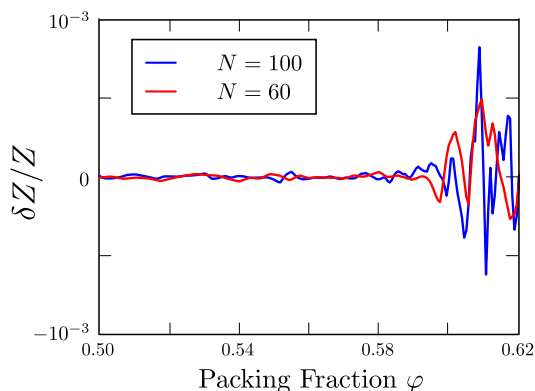

FIG. 9: (color online) Relative difference $\delta Z/Z$ between the dynamic and thermodynamic estimates of the equation of state, Eqs (20) and (17), respectively, for $N = 60$ (red) and $N = 100$ (blue) particles.

weighted and unweighted averages. This result suggests that the entropy of these systems is not strongly correlated with the $g(r)$. Although we cannot rule out the possibility that a small fraction of configurations in the ensemble are ordered, we conclude that the predominant configurations sampled by the algorithm are amorphous.

The above results for the equation of state have been obtained using the dynamic measure of the pressure. Figure 9, shows the relative discrepancy, $\delta Z/Z$ between the thermodynamic and dynamic measures of pressure, given in Eqs. (17) and (20), respectively. We find that $\delta Z/Z \approx 10^{-4}$ until $\varphi \approx 0.59$. The good agreement between these two independent measures is an important validation of the algorithm and gives us confidence that estimates of the equation of state based on Eq. (17) are quite accurate.

TABLE IV: Parameters and bootstrapped 95% confidence interval for fitting the equation of state to the free volume form of Eq. (25) for the range of $0.61 < \varphi < 0.625$.

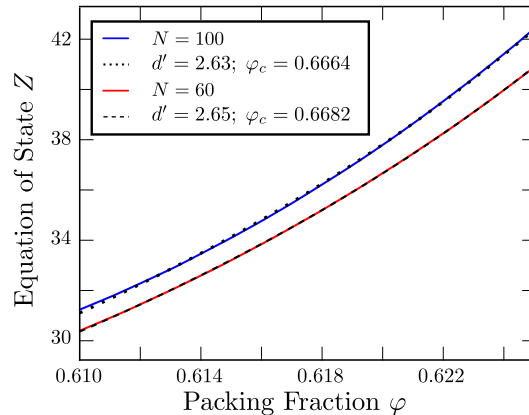| System Size N	d'	φ_c
60	$2.65^{+0.24}_{-0.11}$	$0.668^{+0.005}_{-0.002}$
100	$2.63^{+0.62}_{-0.08}$	$0.666^{+0.011}_{-0.002}$

FIG. 10: High density non-equilibrium equation of state, $\bar{Z}$ (solid curves) along with fits (dotted curves) to the free volume form, Eq. (25) for $N = 60$ (red) and $N = 100$ (blue) particles.

C. High Density Behavior

Though the simulations have fallen out of equilibrium for $\varphi \gtrsim 0.6$, population annealing continues to work as a nonequilibrium protocol to achieve high density packings and nearly jammed states. We fit our equation of state data in the range $0.61 < \varphi < 0.625$ to the "free volume" form,

$$Z = \frac{d'\varphi_c}{\varphi_c - \varphi}. \qquad (25)$$

Table IV gives the parameters of the fit and Fig. 10 shows the simulation data along with fits. It should be noted that in this range of packing fractions the weighted average is dominated by the single highest entropy run so that it is likely that the estimate for φ_c would increase as the number of runs or population size increases. It would be interesting to study how φ_c changes as either population size or number of runs increases. Our results agree reasonably well with Odriozola and Berthier [8] who find $d' = 2.82$ and $\varphi_c = 0.669$, from fits in the range $0.61 < \varphi \lesssim 0.65$.

To explore the behavior of population annealing as a protocol for achieving nearly jammed states, we allowed ten runs to go to much higher packing fraction. These runs terminated after exceeding a maximum run time and reached packing fractions in the range $0.660 < \varphi_c < 0.663$. We fit each run independently to Eq. (25) and obtained φ_c in the range $0.662 < \varphi_c < 0.664$. These

values of φ_c are smaller than the estimates shown in Table IV because they represent only 10 rather than 1000 independent experiments.

V. DISCUSSION

We have developed population annealing Monte Carlo for fluid systems and applied it to a glass-forming binary mixture of hard spheres. We find that population annealing is a promising method for computational studies of fluids in the high density regime. It is a highly parallelized algorithm that is well-suited to simulations on large computer clusters either by using weighted averaging and a large number of independent runs, as we have done here, or by carrying out a smaller number of very large population runs with a massively parallelized implementation of the algorithm. The advantage of the former approach is that we can learn much about the equilibration of weighted averages from the statistics of multiple runs but this advantage is balanced by the lesser efficiency of weighted averaging as opposed to a single large population run.

We measured the equation of state in the glassy region of a 50/50 mixture of hard spheres with diameter ratio 1.4:1 and have obtained precise results that are in reasonable agreement with previous simulations using parallel tempering [8]. We find good agreement with the BM-CSL equation of state up to a packing fraction of 0.58 but strong deviations above that packing fraction. Although this is also the region where equilibration becomes much for difficult, we believe the simulations are reasonably well equilibrated up to a packing fraction of 0.60. We also studied the equation of state at higher densities where population annealing serves as a non-equilibrium protocol for generating jammed states.

We must emphasize that our results at high densities likely describe the metastable equilibrium fluid branch of the equation of state and not the true equilibrium, which presumably consists either of a phase separated crystalline state or, for small N, perhaps a non-random best packing of N spheres. The definition of metastable equilibrium depends on the protocol used to sample the distribution so, in the region where the true equilibrium ensemble contains a significant contribution from non-random states, our results may differ from those obtained from other algorithms. It would be interesting to investigate the true equilibrium states of the binary mixture studied here for finite N.

This work is the first application of population annealing to classical fluids and there are undoubtably many avenues for improvement. We have implemented the algorithm in the NVT ensemble but it is important to study population annealing in the NPT ensemble. It would also be interesting to explore other annealing schedules. For example, a variable number of Monte Carlo sweeps with more sweeps in the region of the glass transition may improve efficiency.

Appendix A: Dynamic chain length

Based on the average number of particles displaced in each event chain, we define an ECMC sweep to be a number of event chain moves that will move approximately N particles on average. We arbitrarily choose to divide this work so that an average of $\sqrt{N}$ particles are displaced in each event chain. Then $\sqrt{N}$ event chains constitutes an ECMC sweep of the system.

The event chain length, ℓ_c must be chosen dynamically, which can be done using Eq. (20), re-written in the form,

$$Z = 1 + \frac{\left\langle \sum_{\text{chains}} (x_k - x_j) \right\rangle_{\text{chains}}}{\ell_c}. \qquad (A1)$$

The "lifting distance" $x_j - x_i$ between two particles of diameter σ with bond orientation (θ, φ) at contact is

$$x_j - x_i = \sigma \sin\theta \cos\varphi. \qquad (A2)$$

In a fluid, this bond orientation should be random, so the average projected distance is

$$\langle x_j - x_i \rangle = \frac{1}{2\pi} \int_0^\pi \int_{-\pi/2}^{\pi/2} d\Omega \, \sigma \sin\theta \cos\varphi = \frac{\sigma}{2}. \qquad (A3)$$

This will be somewhat less accurate in an fcc solid, but it turns out to be close enough to give a reasonable estimate. If an event chain consists of n_c collisions,

$$\left\langle \sum_{\text{chains}} (x_j - x_i) \right\rangle_{\text{chains}} \approx n_c \sigma/2, \qquad (A4)$$

and we find,

$$Z \approx \frac{\ell_c + n_c \sigma/2}{\ell_c}. \qquad (A5)$$

We would like $\bar{n}_c \approx \sqrt{N}$. For density φ_i then we use the equation of state $Z(\varphi_{i-1})$ at the previous density to estimate $\ell_c(\varphi_i)$ as,

$$\ell_c(\varphi_i) \approx \frac{\sigma\sqrt{N}/2}{Z(\varphi_{i-1}) - 1}. \qquad (A6)$$

Empirically, we found a somewhat better approximation to $n_c = \sqrt{N}$ using

$$\ell_c(\varphi_i) = \frac{\sigma\sqrt{N/2}}{Z(\varphi_{i-1}) - 1}. \qquad (A7)$$

[1] R. H. Swendsen and J.-S. Wang, *Replica Monte Carlo simulations of spin glasses*, Phys. Rev. Lett. **57**, 2607 (1986).

[2] K. Hukushima and K. Nemoto, *Exchange Monte Carlo method and application to spin glass simulations*, J. Phys. Soc. Jpn. **65**, 1604 (1996).

[3] H. G. Katzgraber, M. Palassini, and A. P. Young, *Monte Carlo simulations of spin glasses at low temperatures*, Phys. Rev. B **63**, 184422 (2001).

[4] R. A. Banos, A. Cruz, L. A. Fernandez, J. M. Gil-Narvion, A. Gordillo-Guerrero, M. Guidetti, A. Maiorano, F. Mantovani, E. Marinari, V. Martin-Mayor, et al., *Nature of the spin-glass phase at experimental length scales*, Journal of Statistical Mechanics: Theory and Experiment **2010**, P06026 (2010).

[5] B. Yucesoy, H. G. Katzgraber, and J. Machta, *Evidence of non-mean-field-like low-temperature behavior in the Edwards-Anderson spin-glass model*, Phys. Rev. Lett. **109**, 177204 (2012).

[6] T. Okabe, M. Kawata, Y. Okamoto, and M. Mikami, *Replica-exchange Monte Carlo method for the isobaric-isothermal ensemble*, Chemical Physics Letters **335**, 435 (2001).

[7] G. Odriozola, *Replica exchange Monte Carlo applied to hard spheres*, J. Chem. Phys. **131**, 144107 (2009).

[8] G. Odriozola and L. Berthier, *Equilibrium equation of state of a hard sphere binary mixture at very large densities using replica exchange Monte Carlo simulations*, J. Chem. Phys. **134**, 054504 (2011).

[9] E. P. Bernard, W. Krauth, and D. B. Wilson, *Event-chain Monte Carlo algorithms for hard-sphere systems*, Phys. Rev. E **80**, 5 (2009).

[10] M. Isobe and W. Krauth, *Hard-sphere melting and crystallization with event-chain Monte Carlo*, J. Chem. Phys. **143**, 084509 (2015).

[11] 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).

[12] K. Hukushima and Y. Iba, in *The Monte Carlo Method In The Physical Sciences: Celebrating the 50th Anniversary of the Metropolis Algorithm*, edited by J. E. Gubernatis (AIP, 2003), vol. 690, pp. 200–206.

[13] J. Machta, *Population annealing with weighted averages: A Monte Carlo method for rough free-energy landscapes*, Phys. Rev. E **82**, 026704 (2010).

[14] W. Wang, J. Machta, and H. Katzgraber, *Population annealing: Theory and application in spin glasses*, Phys. Rev. E **92**, 063307 (2015).

[15] W. Wang, J. Machta, and H. Katzgraber, *Evidence against a mean-field description of short-range spin glasses revealed through thermal boundary conditions*, Phys. Rev. B **90**, 184412 (2014).

[16] W. Wang, J. Machta, and H. G. Katzgraber, *Chaos in spin glasses revealed through thermal boundary conditions*, Phys. Rev. B (2015).

[17] W. Wang, J. Machta, and H. G. Katzgraber, *Comparing Monte Carlo methods for finding ground states of Ising spin glasses: Population annealing, simulated annealing, and parallel tempering*, Phys. Rev. E **92**, 013303 (2015).

[18] L. Berthier and T. Witten, *Glass transition of dense fluids of hard and compressible spheres*, Phys. Rev. E **80**, 021502 (2009).

[19] G. Brambilla, D. El Masri, M. Pierno, L. Berthier, L. Cipelletti, G. Petekidis, and A. B. Schofield, *Probing the equilibrium dynamics of colloidal hard spheres above the mode-coupling glass transition*, Phys. Rev. Lett. **102**, 085703 (2009).

[20] T. Boublk, *Hard-sphere equation of state*, J. Chem. Phys. **53**, 471 (1970).

[21] G. A. Mansoori, N. F. Carnahan, K. E. Starling, and T. W. Leland, *Equilibrium thermodynamic properties of the mixture of hard spheres*, J. Chem. Phys. **54**, 1523 (1971).

[22] C. Jarzynski, *Equilibrium free-energy differences from nonequilibrium measurements: A master-equation approach*, Phys. Rev. E **56**, 5018 (1997).

[23] A. B. Hopkins, F. H. Stillinger, and S. Torquato, *Densest binary sphere packings*, Phys. Rev. E **85**, 021130 (2012).

[24] C. Amey and J. Machta, *Optimized population annealing for spin glasses*, in preparation.

[25] A. Doucet, N. de Freitas, and N. Gordon, eds., *Sequential Monte Carlo Methods in Practice* (Springer-Verlag, New York, 2001).

[26] G. Brambilla, D. E. Masri, M. Pierno, L. Berthier, and L. Cipelletti, *Probing the equilibrium dynamics of colloidal hard spheres above the mode-coupling glass transition*, Phys. Rev. Lett. **102**, 085703 (2009).

[27] M. Hermes and M. Dijkstra, *Thermodynamic signature of the dynamic glass transition in hard spheres*, J. Phys.: Consens. Matter **22**, 104114 (2010).

[28] C. O'Hern, S. A. Langer, A. J. Liu, and S. R. Nagel, *Random packings of frictionless particles*, Phys. Rev. Lett. **88**, 075507 (2002).

[29] M. Engel, J. A. Anderson, S. C. Glotzer, M. Isobe, E. P. Bernard, and W. Krauth, *Hard-disk equation of state: First-order liquid-hexatic transition in two dimensions with three simulation methods*, Phys. Rev. E **87**, 042134 (2013).

[30] M. Michel, S. C. Kapfer, and W. Krauth, *Generalized event-chain Monte Carlo: Constructing rejection-free global-balance algorithms from infinitesimal steps*, J. Chem. Phys. **140**, 054116 (2014).

[31] E. Bernard, Ph.D. thesis, Universite Pierre et Marie Curie (2011).

[32] E. P. Bernard and W. Krauth, *Two-step melting in two dimensions : First-order liquid-hexatic transition*, Phys. Rev. Lett. **107**, 155704 (2011).

[33] Note that this choice of rescaling introduces a tiny deviation from the hard sphere distribution because one pair of particles out in one replica is *always* in contact immediately after the rescaling. The resulting error is far smaller than other systematic and statistical errors in the simulation.