





RESEARCH ARTICLE | SEPTEMBER 27 2022

Multilevel simulation of hard-sphere mixtures

Paul B. Rohrbach ; Hideki Kobayashi; Robert Scheichl ; Nigel B. Wilding ; Robert L. Jack 



J. Chem. Phys. 157, 124109 (2022)

<https://doi.org/10.1063/5.0102875>



CrossMark

Articles You May Be Interested In

Multilevel applications in education studies

AIP Conference Proceedings (February 2015)

Multilevel geometry optimization

J. Chem. Phys. (February 2000)

Materials and technologies for multilevel metallizations in VLSI

AIP Conference Proceedings (April 1986)

500 kHz or 8.5 GHz?
And all the ranges in between.

Lock-in Amplifiers for your periodic signal measurements



Find out more



Multilevel simulation of hard-sphere mixtures

Cite as: *J. Chem. Phys.* **157**, 124109 (2022); doi: [10.1063/5.0102875](https://doi.org/10.1063/5.0102875)

Submitted: 13 June 2022 • Accepted: 26 August 2022 •

Published Online: 27 September 2022



View Online



Export Citation



CrossMark

Paul B. Rohrbach,^{1,a)}  Hideki Kobayashi,² Robert Scheichl,^{3,4}  Nigel B. Wilding,⁵  and Robert L. Jack^{1,6} 

AFFILIATIONS

¹ Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, United Kingdom

² Institute for Theoretical Physics, Georg-August University, Göttingen, Germany

³ Institute for Applied Mathematics, Heidelberg University, Im Neuenheimer Feld 205, 69120 Heidelberg, Germany

⁴ Department of Mathematical Sciences, University of Bath, Claverton Down, Bath BA2 7AY, United Kingdom

⁵ H. H. Wills Physics Laboratory, University of Bristol, Royal Fort, Bristol BS8 1TL, United Kingdom

⁶ Yusuf Hamied Department of Chemistry, University of Cambridge, Lensfield Road, Cambridge CB2 1EW, United Kingdom

^{a)} Author to whom correspondence should be addressed: pbr28@cam.ac.uk

ABSTRACT

We present a multilevel Monte Carlo simulation method for analyzing multi-scale physical systems via a hierarchy of coarse-grained representations, to obtain numerically exact results, at the most detailed level. We apply the method to a mixture of size-asymmetric hard spheres, in the grand canonical ensemble. A three-level version of the method is compared with a previously studied two-level version. The extra level interpolates between the full mixture and a coarse-grained description where only the large particles are present—this is achieved by restricting the small particles to regions close to the large ones. The three-level method improves the performance of the estimator, at fixed computational cost. We analyze the asymptotic variance of the estimator and discuss the mechanisms for the improved performance.

© 2022 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0102875>

I. INTRODUCTION

A key challenge in the molecular simulation of soft matter is posed by the separation of length-scales between its microscopic description and the existence or emergence of mesoscopic structure. In such cases, one often relies on coarse-grained (CG) descriptions of the system that (approximately) integrate out microscopic degrees of freedom^{1–3} to yield a tractable simplified model. Examples in the soft-matter context include polymers,^{4,5} biomolecules,^{6–8} and colloidal systems.^{9,10} Such CG descriptions are essential for multi-scale modeling approaches.^{11,12} However, they are not usually exact, and the associated coarse-graining errors are often difficult to assess.

Such CG models have been studied extensively in colloidal systems with depletion interactions.^{9,13,14} The typical example is a mixture of relatively large colloidal particles with much smaller non-adsorbing polymers that generate effective attractions between the colloids. This can drive de-mixing, crystallization, or gelation, depending on the context. Model systems in this context include the Asakura–Oosawa (AO) model⁹ where the CG model can even

be exact if the disparity between colloid and polymer radii is large enough. The theoretical tractability of the AO model arises from a simplified modeling assumption that polymer particles act as spheres that can interpenetrate.

Alternatively, one may consider a mixture where both the colloids and the depletant are modeled as hard spheres. From a theoretical perspective, this is an interesting model in its own right as it undergoes a fluid–fluid de-mixing phase separation for large enough size-disparities and concentrations.^{10,15–17} This happens despite the lack of attractive forces between the particles in the model and can be attributed to geometric packing effects of the big and small spheres.

Direct simulation of such mixtures is very challenging because of the large number of small particles. Accurate CG models are also available in this context,¹⁰ but the CG representations are not exact: their errors can be detected by accurate computer simulation of the full (FG) mixtures. Hence, such models are natural testing grounds for theories and simulation methods associated with coarse-graining.

In this context, we recently developed a method^{17,18} that links a CG description with the underlying fine-grained (FG) description. We call this the *two-level* method because the CG and FG models describe the same system with different levels of details. The method was validated by computations on the AO system,¹⁸ where it provided numerically exact results for the FG model, even in the regime where the CG description is not quantitatively accurate. The methodology was also applied to the hard sphere mixture,¹⁷ where it provided a quantitative analysis of the critical point associated with de-mixing of the large and small particles.

These previous results rely on the idea that properties of the FG model can be estimated in terms of some CG quantity, with an additive correction that accounts for the coarse-graining error. This is an importance sampling method, familiar in equilibrium statistical mechanics from free-energy perturbation theory,¹⁹ which involves reweighting between two thermodynamic ensembles. In the present context, the reweighting factors depend on the free energy of the small spheres, computed in a system where the large particles are held fixed. This free energy can be estimated by an annealing process based on Jarzynski's equality^{20–22} that slowly introduces small particles to fixed CG configurations.

In this paper, we present an extension of the two-level method that incorporates additional intermediate levels to improve the overall performance. Specifically, we introduce a step in the annealing process where small particles are partially inserted in regions close to big particles. Before finishing the small-particle insertion, we then replace weighted sets of configurations with unweighted ones, duplicating configurations with large weight and deleting ones with low weight. This resampling step allows us to make optimal use of the information available at the intermediate stage, focusing our subsequent computations on configurations that matter.

This general approach fits in the framework of sequential Monte Carlo (SMC).^{23–25} Such algorithmic ideas have been successfully applied in applications across disciplines under various names, including population Monte Carlo²⁶ or the go-with-the-winners strategy.²⁷ Examples in computational physics include the pruned-enriched Rosenbluth method for polymers,²⁸ the cloning method for rare events,²⁹ and diffusion quantum Monte Carlo.³⁰ We combine the SMC method with an additional variance reduction strategy. Instead of estimating the FG average directly, we combine a CG estimate with estimates of subsequent level differences using the previous levels as control variate. This is the idea behind multilevel Monte Carlo methods.^{31–33} The combination of a difference estimate with SMC has been previously investigated, for example, in Refs. 34–36. As in Ref. 17, we develop a general method alongside its application to highly size-asymmetric binary hard-sphere mixtures, which provide a challenging but well understood example to benchmark our algorithm.

This paper is organized as follows: In Sec. II, we introduce the hard-sphere mixture model. In Sec. III, we summarize the setup of the two-level method before presenting our extension to three (or more) levels. The three-level method requires an intermediate level for the hard-sphere mixture, whose details we discuss in Sec. IV. In Sec. V, we present a numerical test of the method and compare its performance against the two-level method, and in Sec. VI, we present convergence results. We conclude in Sec. VII.

II. HARD-SPHERE MIXTURE

Throughout this work, we illustrate the multilevel method with an example system, which is a mixture of large and small hard spheres at size ratio 10 : 1. This system is challenging for simulation because the big particles may display interesting collective behavior (in particular, a critical point), but the dominant computational cost for simulation comes from the large number of small particles.

However, despite our focus on this single example, we emphasize that the multilevel method is presented in a general way, which should also be applicable in other systems with a separation of length scales.

A. Hard sphere mixture

The example system is a mixture of big and small particles, whose diameters are σ_B and σ_S , respectively. We consider a periodic box $[0, L]^3$ of linear size L and we work in the grand canonical ensemble. (This choice is particularly relevant for analysis of de-mixing, where the number density of large particles is a suitable order parameter.³⁷)

In a given configuration, the numbers of big and small particles are N and n , respectively; the position of the i th big particles is \mathbf{R}_i , while the position of the j th small particle is \mathbf{r}_j . We denote the configurations of big and small particles by $\mathcal{C} = (N; \mathbf{R}_1, \dots, \mathbf{R}_N)$ and $\mathcal{F} = (n; \mathbf{r}_1, \dots, \mathbf{r}_n)$, respectively, and the full configuration is denoted $\mathcal{X} = (\mathcal{C}, \mathcal{F})$.

Since the particles are hard, the temperature plays no role in the following, and so we set the temperature as $k_B T = 1$ without any loss of generality. The equilibrium distribution of the mixture is described by a probability density

$$p_F(\mathcal{C}, \mathcal{F}) = \frac{1}{\Xi_F} e^{\mu_B N + \mu_S n - U_F(\mathcal{C}, \mathcal{F})}, \quad (1)$$

where the subscript F indicates that we refer to the FG model, μ_B and μ_S are the chemical potentials for the large and small particles, respectively, and Ξ_F is the grand canonical partition function. The particles are hard (non-overlapping), so the potential energy is

$$U_F(\mathcal{C}, \mathcal{F}) = \begin{cases} \infty, & \text{if any particles overlap,} \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

This p_F is normalized as $1 = \int p_F(\mathcal{C}, \mathcal{F}) d\mathcal{C} d\mathcal{F}$, the precise meaning of these integrals is given in Appendix A 1.

Within this setting, the dimensionless parameters of the model are the ratio of the particle diameters σ_B/σ_S , the system size parameter L/σ_B , and the two chemical potentials μ_S and μ_B . In practice, μ_S is more naturally parameterized by the associated reservoir volume fraction η_S^* , which we relate to μ_B via an accurate equation of state.³⁸

Our multi-level method is designed for accurate estimates of properties of the large particles. Specifically, we consider observable quantities of interest $A = A(\mathcal{C})$ that only depend on the large particles. (Examples are discussed in Sec. II C; see also Fig. 3.) Our aim is to compute the equilibrium average of A , that is,

$$\langle A \rangle_F = \int A(\mathcal{C}) p_F(\mathcal{C}, \mathcal{F}) d\mathcal{C} d\mathcal{F}. \quad (3)$$

Since $A(C)$ does not depend on \mathcal{F} , it is natural to define the marginal distribution for the big particles

$$p_F(C) = \int p_F(C, \mathcal{F}) d\mathcal{F} \quad (4)$$

so that $\langle A \rangle_F = \int A(C) p_F(C) dC$. A similar situation occurs in the context of statistics, where one seeks to analyze the behavior of a few quantities of interest in a high-dimensional system: in that context, the small-particle degrees of freedom in (3) would be referred to as *nuisance parameters*. This means that their values are not required to compute the quantity of interest, but their statistical properties strongly affect the average of this quantity.

B. Coarse-grained model

If samples for the marginal distribution $p_F(C)$ could be generated by an MC method for the big particles alone, this would make the system much more tractable by simulation. This is a central idea in coarse-grained modeling.² However, the complexity of packing of the small hard spheres means that $p_F(C)$ is a complex distribution, and it is not possible to sample it exactly. A great deal of effort has gone into developing CG models that approximate this distribution with high accuracy.^{10,16,39}

A suitable CG model is an equilibrium distribution with probability density

$$p_C(C) = \frac{1}{\Xi_C} e^{\mu_B N - U_C(C)}, \quad (5)$$

where Ξ_C is the partition function, and the CG (effective) interaction energy is

$$U_C(C) = N\Delta\mu + \sum_{i=1}^{N-1} \sum_{j=i+1}^N V_2(|\mathbf{R}_i - \mathbf{R}_j|), \quad (6)$$

where V_2 is a pairwise interaction potential. Averages with respect to the CG model are denoted as

$$\langle A \rangle_C = \int A(C) p(C) dC. \quad (7)$$

For a suitably chosen V_2 , the coarse distribution $p_C(C)$ can be an accurate approximation to $p_F(C)$. For the CG model in this work, we take the accurate potential $V_2 = V_{\text{RED}}$, developed by Roth, Evans, and Dietrich.¹⁰ Following Ref. 17, we choose $\Delta\mu$ such that the distributions of N coincide for FG and CG models.

C. Benchmark system: Parameters and observables

Throughout the paper, we benchmark our numerical methods by considering the hard-sphere mixture with fixed parameters, as follows: We take the ratio of particle sizes (σ_B/σ_S) = 10, the linear size of the periodic system is $L = 31\sigma_S$, and the small-particle (reservoir) volume fraction is $\eta_S^r = 0.2$. This volume fraction is large enough to generate a significant depletion attraction between the large particles, but not strong enough to cause de-mixing of the large and small particles.¹⁷

Aspects of the CG and FG models are illustrated in Figs. 1(a)–1(c) for these parameters. In particular, we show representative configurations of the CG and FG models, as well as a plot of the RED potential. While direct grand canonical Monte Carlo (GCMC) sampling of the full mixture is possible, in principle, it should be apparent from Fig. 1(c) that this would be intractable because insertion of large particles in such a fluid is hardly ever possible. Advanced MC methods^{39,40} might be applicable, but these tend

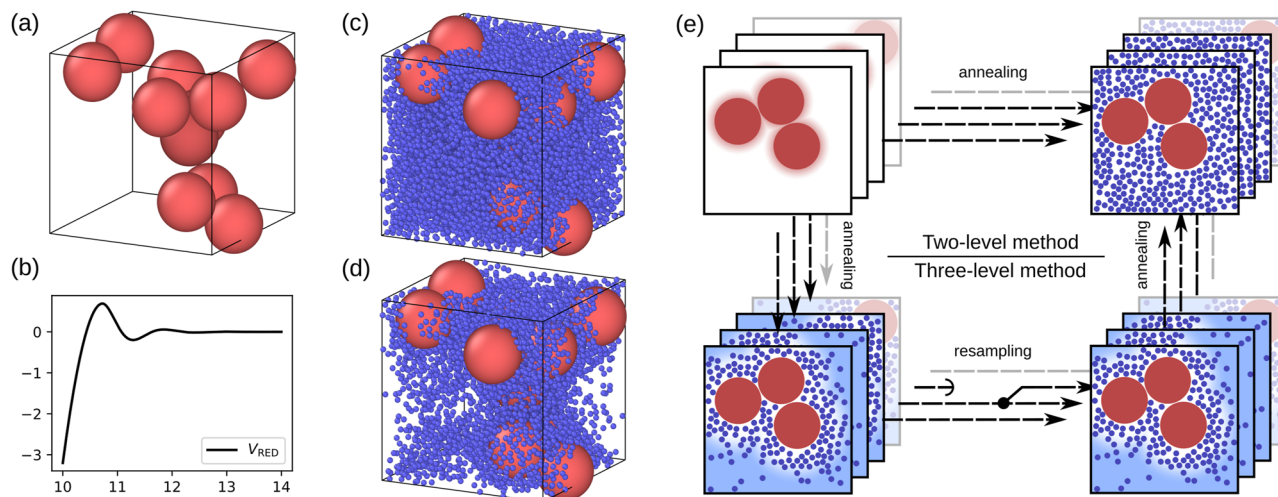


FIG. 1. An overview of the levels of the example system from Sec. II C and the structure of the three-level method. (a) A sample of the CG model p_C with $N = 11$ big particles. (b) The two-body potential V_{RED} used for the CG model. (c) A sample of the full FG model p_F that has the big particle configuration as (a). This system contains $n = 8842$ small particles. (d) A sample of the partially inserted model p_I used in the three-level algorithm. The small particles are primarily inserted around big particles, reducing the number to $n = 4473$. (e) A sketch of the two- and three-level method: starting with a population of CG configurations, we can directly compute importance weights by simulating an annealing process introducing the small particles (upper path, two-level method). Alternatively, we introduce a partially inserted intermediate level where we interrupt this annealing process and resample to boost relevant configurations (lower path, three-level method).




FIG. 2. Properties of a big-particle-only hard sphere model with RED potential when varying the effective chemical potential μ_B^{eff} as defined in the main text. (a) The average number of big particles $\langle N \rangle$. (b) The variance of the number of big particles $\text{var}(N)$, which is maximized around $\mu_B^{\text{eff}} = -7$.

to struggle when the volume fraction gets large. This motivates the development of two-level and multi-level methods.

Figure 2 highlights properties of the distribution of the number of big particles for the CG model when varying the effective large-particle chemical potential $\mu_B^{\text{eff}} = \mu_B - \Delta\mu$. In particular, Fig. 2(b) shows that increasing μ_B^{eff} in the CG model leads to a non-monotonic behavior in the variance of the particle number N (analogous to the compressibility of the model). This maximum indicates that the system has a tendency for de-mixing at larger η_S^*




FIG. 3. Two quantities of interest for the binary hard-sphere system, computed for the CG and FG model from Sec. II C. (a) The big-particle pair correlation function $g(r)$. Apart from underestimating its value at the touch of two big particles, the CG approximation captures the behavior accurately. (b) The distribution of the number of big particles N . By the choice of $\Delta\mu$, the average number of big particles of the CG and FG models coincide. Both models are clearly well below the critical point of demixing.

(one expects a divergent compressibility at the critical point, if one exists). In the following, we fix μ_B at the value corresponding to this maximum—the relatively large fluctuations at this point are challenging for the multi-level model because the distributions $p_C(C)$ and $p_F(C)$ are broader, requiring good sampling. The corresponding CG system has an average of $N \approx 11.6$ big particles, occupying around 20% of the available volume.

For the specific quantities that we will compute for this mixture, Fig. 3 shows the expectations of the big-particle pair correlation function $g(r)$ and the distribution of the number of big particles $P(N)$. Results are shown for both CG and FG models (in the FG case, the results are computed using the two-level method). For both quantities of interest, the CG model provides an accurate but not exact description of the model. In particular, the CG model underestimates the pair correlation at the point where two big particles are in contact. The distributions of the number of big particles in Fig. 3(b) are both unimodal: both the FG and CG systems are well below the critical point of demixing.

Compared to the critical hard-sphere mixture discussed in Ref. 17, the system we consider here is smaller and has a lower volume fraction η_S^* of the small particles. This is still challenging for conventional Monte Carlo algorithms but can be simulated fast enough to evaluate the performance and compare the computational methods discussed here. Furthermore, the lower small-particle volume fraction helps with the construction of the intermediate level in Sec. IV, whose underlying approximation decays as η_S^* increases (see Appendix B).

III. MULTILEVEL SIMULATION

A. Overview

This section reviews the two-level method of Refs. 17 and 18 and then lays out its three-level extension. The presentation of the method is intended to be generic and applicable to a variety of systems. However, we first introduce the key ideas using the example and illustrations of Fig. 1 for the hard-sphere mixture.

The two-level method is constructed with the scale separation of the mixture in mind: it splits the simulation of the big and small spheres into two stages by first simulating a CG system of large particles alone and computing $\langle A \rangle_C$. Then, differences between $\langle A \rangle_C$ and $\langle A \rangle_F$ are computed by a reweighting (importance sampling) method. The weight factors for this computation are obtained by an annealing step, where the small particles are slowly inserted into the system, with the large particles held fixed [see Fig. 1(e)]. The advantage of this procedure is that large particle motion only happens in the CG simulation where the small particles are absent—there is no scale separation in this case so simulations are tractable. Similarly, insertion of the small particles happens in a background of fixed large particles, so these annealing simulations do not suffer long time scales associated with large-particle motion. This makes for tractable simulations in scale-separated systems, as long as the CG model is sufficiently accurate: see Refs. 17 and 18 for further discussion.

In practice, the simulation effort for two-level computations is dominated by the annealing step. The weighting factors are required to high accuracy, which means that the annealing must be done gradually. Moreover, the weights are subject to numerical uncertainties that tend to be large in systems with many small particles.

This limits the method to systems of moderate size, with moderate η_S^r (see Ref. 17).

We show in this work that such problems can be reduced by breaking the annealing process into several stages—this is the idea of the three-level method [Fig. 1(e)]. Specifically, we start (as before) with a population of configurations of the CG model. We perform a first annealing step where the small particles are added in regions that are close to large ones. The information from this step is used in a resampling process, which partially corrects the coarse-graining error by discarding some of the configurations from the population and duplicating others. (This idea is similar to go-with-the-winners.²⁷) Finally, the second annealing step inserts the small particles in the remaining empty regions, arriving at configurations of the FG model. Hence, the end point is the same as the two-level method, but the annealing route is different.

In practice, the effectiveness of the three-level method relies on a clear physical understanding of the intermediate (partially inserted) system in order to decide which configurations to discard in the resampling step. For the hard-sphere case, that issue will be discussed in Sec. IV; a more general discussion is given in Sec. VII. The remainder of this section describes the two- and three-level methods in more detail.

B. Two-level method

We review the two-level method of Refs. 17 and 18. For a general presentation, we assume that CG and FG models exist with configurations \mathcal{C} and $\mathcal{X} = (\mathcal{C}, \mathcal{F})$, respectively. In the case of hard spheres, \mathcal{C} and \mathcal{F} correspond to configurations of the large and small spheres, respectively.

The two-level method is an importance sampling⁴¹ (or reweighting) computation, closely related to the free-energy perturbation method of Zwanzig.¹⁹ We use the grand canonical Monte Carlo (GCMC) method to sample M_C configurations from p_C , and these are denoted by $\mathcal{C}^1, \mathcal{C}^2, \dots, \mathcal{C}^{M_C}$. Then, the CG average can be estimated as

$$\hat{A}_C = \frac{1}{M_C} \sum_{j=1}^{M_C} A(\mathcal{C}^j). \quad (8)$$

As the sampling is increased ($M_C \rightarrow \infty$), we have $\hat{A}_C \rightarrow \langle A \rangle_C$. However, if the coarse-graining error

$$\Delta = \langle A \rangle_F - \langle A \rangle_C \quad (9)$$

is significant, \hat{A}_C does not provide an accurate estimate of $\langle A \rangle_F$.

To address this problem, we use an annealing procedure based on Jarzynski's equality²⁰ that starts from a coarse configuration \mathcal{C} and populates the fine degrees of freedom \mathcal{F} ; at the same time, it generates a random weight $\hat{W}(\mathcal{C})$ with the property that

$$\langle \hat{W}(\mathcal{C}) \rangle_J = \frac{\xi p_F(\mathcal{C})}{p_C(\mathcal{C})}, \quad (10)$$

where the angle brackets with subscript J indicate an averaging over the annealing process (analogous to Jarzynski's equality²⁰), and ξ is a constant (independent of \mathcal{C}). The details of the annealing process are given in Appendix A. It is applied to a set of M_F coarse configurations, again denoted by $\mathcal{C}^1, \mathcal{C}^2, \dots, \mathcal{C}^{M_F}$, which are typically a subset of the M_C CG configurations above.

For later convenience, we define

$$\hat{W}^n(\mathcal{C}) = \hat{W}(\mathcal{C})/\xi. \quad (11)$$

In practical applications, the constant ξ is not known, but its effect can be controlled by defining the self-normalized weight

$$\hat{w}(\mathcal{C}^j) = \frac{\hat{W}(\mathcal{C}^j)}{\frac{1}{M_F} \sum_{i=1}^{M_F} \hat{W}(\mathcal{C}^i)}. \quad (12)$$

Since the \mathcal{C}^j are representative of p_C , the denominator in \hat{w} converges to ξ as $M_F \rightarrow \infty$ and so $\hat{w}(\mathcal{C}^j) \rightarrow \hat{W}^n(\mathcal{C}^j)$. Then, the estimator

$$\hat{A}_F = \frac{1}{M_F} \sum_{j=1}^{M_F} \hat{w}(\mathcal{C}^j) A(\mathcal{C}^j) \quad (13)$$

converges to $\langle A \rangle_F$ as $M_F \rightarrow \infty$. (In the case that \hat{W} is not random then this procedure recovers the free energy perturbation theory of Zwanzig.¹⁹)

The annealing process has one useful additional property: Let the joint probability density for the weight and the fine degrees of freedom be $\kappa(\hat{W}, \mathcal{F} | \mathcal{C})$, which is normalized as $\int \kappa(\hat{W}, \mathcal{F} | \mathcal{C}) d\mathcal{F} d\hat{W} = 1$. We show in Appendix A that

$$\int \hat{W} \kappa(\hat{W}, \mathcal{F} | \mathcal{C}) d\hat{W} = \frac{\xi p_F(\mathcal{C}, \mathcal{F})}{p_C(\mathcal{C})}. \quad (14)$$

This formula is the essential property of the annealing procedure, which is required for the operation of the method. Additionally integrating over \mathcal{F} shows that (14) ensures that (10) also holds. This means in turn that if $B = B(\mathcal{C}, \mathcal{F})$ is an observable quantity that depends on both coarse and fine degrees of freedom,

$$\hat{B}_F = \frac{1}{M_F} \sum_{j=1}^{M_F} \hat{w}(\mathcal{C}^j) B(\mathcal{C}^j, \mathcal{F}^j) \quad (15)$$

converges to $\langle B \rangle_F$ as $M_F \rightarrow \infty$.

This method can be easily improved without extra computational effort. The key idea^{31–33} is to estimate the FG average as the sum of the CG average and the coarse-graining error (9),

$$\langle A \rangle_F = \langle A \rangle_C + \Delta. \quad (16)$$

Then, use importance sampling to estimate Δ , as

$$\hat{\Delta} = \frac{1}{M_F} \sum_{j=1}^{M_F} (\hat{w}(\mathcal{C}^j) - 1) A(\mathcal{C}^j). \quad (17)$$

Finally, a suitable estimator for the FG average is obtained by combining the estimate of the coarse-graining error with the corresponding CG quantity,

$$\hat{A}_{F,\Delta} = \hat{A}_C + \hat{\Delta}. \quad (18)$$

This estimator converges to $\langle A \rangle_F$ in the limit where $M_C, M_F \rightarrow \infty$. As discussed in Ref. 18, the variance of the estimate $\hat{\Delta}$ is typically smaller than that of \hat{A}_F , and the CG estimate \hat{A}_C is cheap to compute

accurately. Thus, the combined difference estimator $\hat{A}_{F,\Delta}$ is typically more accurate at fixed computational cost.

The importance sampling methodology has a useful physical interpretation, which we explain for the example of the hard-sphere mixture. If we consider a fixed configuration of the large particles, the grand canonical partition function for the small particles is

$$\Xi[\mathcal{C}, \mu_S] = \int e^{\mu_S n - U_F(\mathcal{C}, \mathcal{F})} d\mathcal{F}. \quad (19)$$

As the system is annealed (the small particles are inserted), we estimate (19) by a free-energy method based on Jarzynski's equality²⁰ (see Appendix A for details). Since the annealing is stochastic, this yields an estimate of the partition function, which we denote by $\hat{\Xi}[\mathcal{C}, \mu_S]$. Moreover, this estimate is unbiased $\langle \hat{\Xi}[\mathcal{C}, \mu_S] \rangle_I = \Xi[\mathcal{C}, \mu_S]$. Hence, we can take

$$\hat{W}(\mathcal{C}) = \hat{\Xi}[\mathcal{C}, \mu_S] e^{U_C(\mathcal{C})}, \quad (20)$$

and using (4) and (5), we see that (10) holds, with $\xi = (\Xi_C / \Xi_F)$.

Physically, the CG model is constructed so that the Boltzmann factor $e^{-U_C(\mathcal{C})}$ is a good estimate of the small-particle partition function $\Xi[\mathcal{C}, \mu_S]$. If this is the case, the model is accurate. The two-level methodology uses estimates of the small-particle partition function (or, equivalently, their free energy) and compares it with the assumptions that were made about this quantity in the CG model. By analyzing the differences between these quantities, the differences between CG and FG models can be quantified. The effectiveness of this method for numerical simulation of the mixtures of large and small particles was discussed in Refs. 17 and 18.

The distribution of the importance weights $\hat{w}(\mathcal{C})$ impacts the accuracy of the resulting FG estimate \hat{A}_F . Additionally, it serves as a useful indicator of the accuracy of the CG model and the variance of the free energy computation. To give an example, we apply the two-level method to the example problem from Sec. II C. In Fig. 4, we show the empirical distribution of 18 000 weights of the example system that are computed using an accurate annealing process; we use these computations as the reference solution in Sec. V. This illustrates a situation where the two-level method is applicable, where no single sample dominates and only very few samples have a weight larger than 10.

If one considers less accurate CG models, the variance of the weights increases, and the tail of their distribution gets heavier. Eventually, one would reach a situation where a few samples dominate the weighted sum (13). For accurately computed weights $\hat{w}(\mathcal{C})$, such a breakdown of the two-level method indicates that the CG




FIG. 4. The empirical distribution of weights $\hat{w}(\mathcal{C})$ of the two-level method for 18 000 coarse samples $\mathcal{C} \sim p_C$ applied to the example system from Sec. II C.

model is not sufficiently accurate. This behavior provides a useful feedback loop that can be used to iterate on the CG model itself.¹⁸

C. Three-level method

We now present the three-level method for estimation of $\langle A(\mathcal{C}) \rangle_F$.

1. Coarse level

We start by generating M_0 samples of the CG model, denoted by $\mathcal{C}_0^1, \dots, \mathcal{C}_0^{M_0}$. The subscript indicates the step within the algorithm (which is 0 for the initial sampling of coarse configurations). The CG average of A can be estimated similarly to (8)

$$\hat{A}_C^{3L} = \frac{1}{M_0} \sum_{j=1}^{M_0} A(\mathcal{C}_0^j). \quad (21)$$

2. Intermediate level

In addition to the CG and FG models, the three-level method also relies on an intermediate set of configurations, which correspond in the hard-sphere mixture to the system where the small particles have been inserted in regions close to the large ones (see Fig. 5). This state is described by an equilibrium probability distribution

$$p_1(\mathcal{C}, \mathcal{F}) = \frac{1}{\Xi_1} e^{\mu_S N + \mu_S n - U_1(\mathcal{C}, \mathcal{F})}, \quad (22)$$

where $U_1(\mathcal{C}, \mathcal{F})$ is an interaction energy. Its construction for the hard-sphere mixture will be discussed in Sec. IV.

The first annealing step of the three-level algorithm applies the two-level method, with the FG distribution p_F replaced by p_1 . This part of the algorithm closely follows Sec. III B, and we give a brief discussion that mostly serves to fix notation. We start with a set of M_1 coarse configurations that are samples of p_C ; they are denoted by $\mathcal{C}_1^1, \mathcal{C}_1^2, \dots, \mathcal{C}_1^{M_1}$, where the subscript 1 indicates the intermediate stage of the three-level method. (These will typically be a subset of the configurations that were generated on the coarse level.)




FIG. 5. Schematic representation of the small particle insertion process during two stages of the three level method.

For each coarse configuration C_1 , we anneal the fine degrees of freedom of the system $\hat{\mathcal{F}}_1$ to arrive at the intermediate level and generate a random weight $\hat{W}_1(C_1)$ with the property

$$\langle \hat{W}_1(C_1) \rangle_I = \frac{\xi_1 p_I(C_1)}{p_C(C_1)}, \quad (23)$$

with a constant ξ_1 independent of C_1 . (For the hard spheres, we recall that particles are inserted preferentially in regions close to large ones, and this is illustrated in the top row of Fig. 5.)

As before, we define $\hat{W}_1^n(C_1) = \hat{W}_1(C_1)/\xi_1$. Again the constant ξ_1 is generally not known, so we define the self-normalized weight

$$\hat{w}_1(C_1^j) = \frac{\hat{W}_1(C_1^j)}{\frac{1}{M_1} \sum_{i=1}^{M_1} \hat{W}_1(C_1^i)}, \quad (24)$$

which converges to $\hat{W}_1^n(C_1^j)$ as $M_1 \rightarrow \infty$. Then, the estimator

$$\hat{A}_I^{3L} = \frac{1}{M_1} \sum_{j=1}^{M_1} \hat{w}_1(C_1^j) A(C_1^j) \quad (25)$$

converges to $\langle A \rangle_I$ as $M_1 \rightarrow \infty$. Similar to (14), the joint probability density $\kappa_1(\hat{W}_1, \hat{\mathcal{F}}_1 | C_1)$ of the weight and fine degrees of freedom at the intermediate level, defined by the annealing process, fulfills

$$\int \hat{W}_1 \kappa_1(\hat{W}_1, \hat{\mathcal{F}}_1 | C_1) d\hat{W}_1 = \frac{\xi_1 p_I(C_1, \hat{\mathcal{F}}_1)}{p_C(C_1)}. \quad (26)$$

Hence, similar to (15), we also obtain

$$\hat{B}_I^{3L} = \frac{1}{M_1} \sum_{j=1}^{M_1} \hat{w}_1(C_1^j) B(C_1^j, \hat{\mathcal{F}}_1^j), \quad (27)$$

which converges to $\langle B \rangle_I$ as $M_1 \rightarrow \infty$.

3. Fine level

At the end of the intermediate level, we have M_1 large-particle configurations. For each configuration C_1^j , the process of annealing

to the intermediate level also provided the weight $\hat{w}_1(C_1^j)$ and the small-particle configuration $\hat{\mathcal{F}}_1^j$. This information can be used to build a set of configurations that are representative of p_I . This procedure is called *resampling*, its validity in this example relies on the property (26) of the annealing procedure. This is the part of the method that is similar to population-based sampling approaches, such as SMC²⁸ or go-with-the-winners.²⁷ The idea is that one should focus the effort of the annealing process onto coarse configurations that are typical of the full system, and to discard those which are atypical, see Fig. 6 for a visualization of this step.

We write $\hat{\mathcal{X}}_1^j = (C_1^j, \hat{\mathcal{F}}_1^j)$ for the full configuration that is obtained by the annealing procedure at the intermediate level. The resampled configurations will be denoted by $\mathcal{X}_2^1, \mathcal{X}_2^2, \dots, \mathcal{X}_2^{M_2}$; they are representative of the intermediate level p_I . There are M_2 of them, and the subscript 2 indicates the final stage of the three-level method. The simplest resampling method (*multinomial resampling*) is that each \mathcal{X}_2^i is obtained by copying one of the $\hat{\mathcal{X}}_1^j$, chosen at random with probability $\hat{w}_1(C_1^j)$. In applications, one typically replaces this by a lower variance resampling scheme like residual resampling; see Ref. 42 for a comparison of commonly used variants.

We then perform the second annealing step that starts from an intermediate level configuration $\mathcal{X}_2 = (C_2, \mathcal{F}_2)$ and anneals the fine degrees of freedom from the intermediate to the fine level, yielding $\hat{\mathcal{F}}_2$ and a weight $\hat{W}_2(\mathcal{X}_2)$; details are given in Appendix A. For the hard sphere system, this involves further insertion of small particles to fill the system and generate realistic configurations of the full mixture. This procedure is shown in the bottom row of Fig. 5.

Since the starting point of the annealing procedure is \mathcal{X}_2 , the joint probability density of the annealing process $\kappa_2(\hat{W}_2, \hat{\mathcal{F}}_2 | \mathcal{X}_2)$ depends on both large and small particles. Therefore, the analog of (26) requires an additional average over the small particles of the starting configuration

$$\int \hat{W}_2 \kappa_2(\hat{W}_2, \hat{\mathcal{F}}_2 | C_2, \mathcal{F}_2) p_I(\mathcal{F}_2 | C_2) d\hat{W}_2 d\mathcal{F}_2 = \frac{\xi_2 p_{\mathbb{F}}(C_2, \hat{\mathcal{F}}_2)}{p_I(C_2)} \quad (28)$$




FIG. 6. Visualization of the resampling step. We start with a population of weighted configurations (top row), where the weighting is depicted by a star rating. The goal of the resampling step is to randomly transform the weighted population into an unweighted one that has, on average, the same empirical distribution. We achieve this by duplicating large-weight configurations and deleting small-weight configurations, yielding an unweighted population of configurations (bottom row).

for some constant ξ_2 . Note that $p_1(\mathcal{F}|\mathcal{C}) = p_1(\mathcal{C}, \mathcal{F})/p_1(\mathcal{C})$. Similar to (23), the weights $\hat{W}_2(\mathcal{X}_2)$ have the property

$$\int \langle \hat{W}_2(\mathcal{C}_2, \mathcal{F}_2) \rangle_{p_1(\mathcal{F}_2 | \mathcal{C}_2)} d\mathcal{F}_2 = \frac{\xi_2 p_F(\mathcal{C}_2)}{p_1(\mathcal{C}_2)}. \quad (29)$$

From here, we proceed as before. We define the normalized weight $\hat{W}_2^N(\mathcal{X}_2) = \hat{W}_2(\mathcal{X}_2)/\xi_2$ and its self-normalized estimate

$$\hat{w}_2(\mathcal{X}_2^j) = \frac{\hat{W}_2(\mathcal{X}_2^j)}{\frac{1}{M_2} \sum_{i=1}^{M_2} \hat{W}_2(\mathcal{X}_2^i)}. \quad (30)$$

Since the \mathcal{X}_2^j are representative of p_1 , it follows from (28) that observables of the coarse system A can be estimated as

$$\hat{A}_F^{3L} = \frac{1}{M_2} \sum_{j=1}^{M_2} \hat{w}_2(\mathcal{X}_2^j) A(\mathcal{C}_2^j), \quad (31)$$

which converges to $\langle A \rangle_F$ as $M_2 \rightarrow \infty$. Similar to (15), we can also obtain a consistent FG estimates of observable quantities B that depend both on coarse and fine degrees of freedom by

$$\hat{B}_F^{3L} = \frac{1}{M_2} \sum_{j=1}^{M_2} \hat{w}_2(\mathcal{X}_2^j) B(\mathcal{C}_2^j, \hat{\mathcal{F}}_2^j). \quad (32)$$

Following the same variance reduction strategy as in Sec. III B, we can define a difference estimator of the FG average, which is expected to have lower statistical uncertainty: let

$$\hat{\Delta}_I^{3L} = \frac{1}{M_1} \sum_{j=1}^{M_1} (\hat{w}_1(\mathcal{C}_1^j) - 1) A(\mathcal{C}_1^j), \quad (33)$$

$$\hat{\Delta}_F^{3L} = \frac{1}{M_2} \sum_{j=1}^{M_2} (\hat{w}_2(\mathcal{X}_2^j) - 1) A(\mathcal{C}_2^j). \quad (34)$$

Then,

$$\hat{A}_{F,\Delta}^{3L} = \hat{A}_C^{3L} + \hat{\Delta}_I^{3L} + \hat{\Delta}_F^{3L} \quad (35)$$

is a consistent estimator of $\langle A \rangle_F$, analogous to (18).

4. General features of the three-level method

A few comments on the three-level method are in order. First, there is a simple generalization to four or more levels by splitting the annealing procedure into more than two stages. As such, the method is an example of a sequential Monte Carlo (SMC) algorithm (which is sometimes more descriptively referred to as sequential importance sampling and resampling^{38,43–45}). We note from (10) that the weights obtained from the annealing step are random, and this is not the standard situation in SMC, but similar ideas have been previously studied in Refs. 46–49. Combining an SMC algorithm with a difference estimate as in (35) has been investigated in Refs. 34–36.

Second, we observe that the key distinction between the two- and three-level algorithms is the resampling step at the intermediate level. Without this, the three-level method reduces to a simple two-level method with an arbitrary stop in the middle of the annealing process. As noted above, the resampling process is designed to partially correct differences between the CG and FG models. This

relies on a good accuracy of the intermediate level (otherwise the wrong configurations might be discarded, which hinders numerical accuracy). On the other hand, we note that for sufficiently large numbers of samples M_0, M_1, M_2 , the method does provide accurate FG estimates, even if the CG and intermediate level models are not extremely accurate. The distinction between the different methods comes through the number of samples that are required to obtain accurate FG results.

Third, note that the ideal situation for difference estimation is that the three terms in (35) get successively smaller. That is, the coarse estimate is already close to $\langle A \rangle_F$, the intermediate-level estimate provides a large part of the correction, and the fine-level correction is small. In this case, it is natural to use a tapering strategy where the number of samples used at each level decreases,

$$M_0 > M_1 > M_2. \quad (36)$$

This allows a fixed computational budget to be distributed evenly between the various levels to minimize the total error.

IV. CONSTRUCTION OF THE INTERMEDIATE LEVEL

As noted above, the intermediate probability distribution p_1 must be designed carefully in order for the resampling part of the three-level method to be effective. We now describe how this is achieved for the hard sphere mixture.

To motivate the intermediate level, recall Fig. 5, and note that defining a suitable CG model is equivalent to an estimate of the small-particle free energy in the final (fully inserted) state. The physical idea of the intermediate level is that the free energy associated with the first stage of insertion may be hard to estimate (because of the complicated packing of the small particles around the large ones), but the free energy difference associated with the second stage should be easier (because it corresponds to insertion into large empty regions where the packing of the small particles is similar to that of a homogeneous fluid, whose free energy can be estimated based on analytic approximations). A combination of these ideas yields an intermediate level that represents the big particle statistics more accurately than the CG model. Similar ideas have been considered before in multi-scale simulation,^{4,50,51} in particular the problem of estimating the small-particle free energy has some similarities to estimation of solvation free energies (where the depletant here plays the role of a solvent).

We start by analyzing the small particles, so we fix the large particles in some configuration \mathcal{C} . The idea of the intermediate level is to first insert small particles only in a region close to the large particles \mathcal{C} , and then use this information to make the intermediate marginal distribution $p_1(\mathcal{C})$ match the FG marginal $p_F(\mathcal{C})$ as closely as possible. The structure of the intermediate level is depicted in the bottom row of Fig. 1(e), and an example configuration is shown in Fig. 1(d). We implement this idea by introducing an effective (one-body) potential that acts on the small particles. We first define

$$\text{dist}(\mathbf{r}, \mathcal{C}) = \min_{j=1, \dots, N} |\mathbf{r} - \mathbf{R}_j| \quad (37)$$

to be the distance from the point \mathbf{r} to the nearest large particle. Small-particle insertion is suppressed in regions far from large particles by a potential energy term

$$\tilde{U}(C, \mathcal{F}) = \sum_{j=1}^n E_C(\mathbf{r}_j), \quad (38)$$

where

$$E_C(\mathbf{r}) = \varepsilon(\text{dist}(\mathbf{r}, C)) \quad (39)$$

and the function

$$\varepsilon(r) = \begin{cases} 0, & r < \delta_{\text{free}}, \\ s \sin^2 \left[\frac{(r - \delta_{\text{free}})\pi}{2l} \right], & \delta_{\text{free}} \leq r < \delta_{\text{free}} + l, \\ s, & r \geq \delta_{\text{free}} + l \end{cases} \quad (40)$$

interpolates from zero (for small distances r) to the value s at large r . This function acts as a smoothed out step function, where δ_{free} is the position of the step and l its width. In Figs. 1(e) and 5, areas where $E_C(\mathbf{r}) > 0$ are indicated by blue shaded regions, in which the insertion of small particles is suppressed.

Then, define a grand-canonical probability distribution for the small particles in the partially inserted (intermediate) system as

$$\tilde{p}_1(\mathcal{F} | C) = \frac{1}{\Xi_1[C, \mu_S]} e^{\mu_S n - U_F(C, \mathcal{F}) - \tilde{U}(C, \mathcal{F})}. \quad (41)$$

This distribution is normalized as $\int \tilde{p}_1(\mathcal{F} | C) d\mathcal{F} = 1$. It depends on the three parameters $s, \delta_{\text{free}}, l$, as well as the underlying parameters of the hard sphere mixture model.

The next step is to construct the weights $\hat{W}_1(C)$. For consistency with (22), we write the intermediate-level distribution in the form

$$p_1(C, \mathcal{F}) = \frac{1}{\Xi_1} e^{\mu_B N + \mu_S n - U_F(C, \mathcal{F}) - \tilde{U}(C, \mathcal{F}) - \Phi^{\text{corr}}(C)}. \quad (42)$$

As discussed above, the term Φ^{corr} should be designed so that the respective coarse-particle marginals $p_1(C)$ and $p_F(C)$ match as closely as possible. Using (4), (19), and (41), we can show that a perfect match requires $\Phi^{\text{corr}}(C) = \Phi^{\text{ex}}(C)$ with

$$\Phi^{\text{ex}}(C) = \log \frac{\Xi_1[C, \mu_S]}{\Xi_F[C, \mu_S]} - \phi_0, \quad (43)$$

where ϕ_0 is an irrelevant constant. Since the Ξ s in (43) are partition functions, determination of Φ^{ex} reduces to computation of the free energy difference between the non-homogeneous small particle distributions of the partially and fully inserted system. We now explain how Φ^{corr} is defined, as an approximation to Φ^{ex} .

A. Square-gradient approximation of a non-homogeneous hard sphere fluid

As a preliminary step for estimating Φ^{ex} , we first consider the grand potential Φ for the small particles, in a system with no large particles, where the small particles feel an (arbitrary) smooth potential $\mathcal{E} = \mathcal{E}(\mathbf{r})$. The grand potential of this system is

$$\Phi[\mathcal{E}; \mu_S] = -\log \int e^{\mu_S n - U_F(\mathbf{0}, \mathcal{F}) - \sum_{j=1}^n \varepsilon(\mathbf{r}_j)} d\mathcal{F}, \quad (44)$$

where $\mathbf{0}$ indicates the large-particle configuration with no particles at all ($N = 0$).

If \mathcal{E} varies slowly in space, a simple approach to this integral is to assume that the system is locally the same as a homogeneous system in equilibrium—similar to the local density approximation.⁵² In this case,

$$\Phi[\mathcal{E}; \mu_S] \approx -\int_V p(\mu_S - \mathcal{E}(\mathbf{r})) d\mathbf{r}, \quad (45)$$

where p is the pressure, expressed as a function of the chemical potential.

However, this approximation is not sufficiently accurate for the current application. To this end, we include a correction to account for inhomogeneities, as a squared gradient term $\Phi[\mathcal{E}; \mu_S] \approx \Phi^{\text{sq}}[\mathcal{E}; \mu_S]$ with

$$\Phi^{\text{sq}}[\mathcal{E}; \mu_S] = -\int_V p(\mu_S - \mathcal{E}(\mathbf{r})) + g(\mu_S - \mathcal{E}(\mathbf{r})) |\nabla \mathcal{E}(\mathbf{r})|^2 d\mathbf{r}. \quad (46)$$

(Within a gradient expansion, this is the first correction that is consistent with rotational and inversion symmetry.)

We show in Appendix B 1 that g can be estimated as

$$g(\mu) = \frac{3\eta}{2\pi\sigma^3} \frac{\partial^2}{\partial q^2} S(\mu; q) \Big|_{q=0}, \quad (47)$$

where $S(\mu; q)$ is the structure factor of the small hard-sphere system. For a numerical estimate of this Φ^{sq} , we estimate the pressure p by the accurate equation of state from Ref. 38, and g is estimated from (47) using the structure factor from Ref. 53. A numerical example demonstrating the accuracy of this second order approximation for a non-homogeneous hard-sphere fluid can be found in Appendix B 2.

B. Definition of Φ^{corr}

We are now in a position to approximate Φ^{ex} in terms of Φ^{sq} . This (analytical) calculation is illustrated in Fig. 7. We require an estimate of Φ^{ex} , which is the free-energy difference between the partially inserted and fully inserted systems in panels (a) and (b). This is achieved as a sum of three free-energy differences. In the first step, the large particles are removed and the small-particle fluid is re-equilibrated, to fill up the remaining space, leading to panel (c). Then, the confining potential \tilde{U} is removed and the small particles fully inserted, leading to (d). Finally, the large particles are re-inserted and the small particles re-equilibrated again, leading to (b).

To make this precise, define $\Phi_c[\mathcal{E}; \mu_S]$ as the grand potential of the small particles in the potential \mathcal{E} , where the large particles are also included, with configuration C . Then, the desired free energy difference between panels (a) and (b) is

$$\Phi^{\text{ex}}(C) = \Phi_c[0; \mu_S] - \Phi_c[E_C; \mu_S], \quad (48)$$

where we took $\phi_0 = 0$.

From the definitions in Sec. IV A, the free energy difference between panels (c) and (d) is $\Phi[0; \mu_S] - \Phi[E_C; \mu_S]$, from (38) and (44). Our central approximation is that the free energy difference between panels (a) and (c) is (approximately) equal and opposite of the difference between (d) and (b) because the local environment of the large particles is the same in both cases. (The only differences are




FIG. 7. Illustration of the computation of Φ^{corr} , which is an estimate of the free energy difference Φ^{ex} between panels (a) and (b) [see (48)]. As described in the text, this difference is computed as a sum of three differences using the integration path [(a),(c),(d),(b)]. The free energy difference between (c) and (d) is $\Phi[0; \mu_S] - \Phi[E_C; \mu_S]$. We make the approximation that the differences between [(a),(c)] and [(d),(b)] are equal and opposite, and this should be accurate if the shaded blue regions in panel (a) are well-separated in space from the large particles. Combining this assumption with the square gradient approximation (46) yields $\Phi^{\text{corr}}(C)$ in (50) as a numerically tractable estimate of Φ^{ex} .

in regions far from any large particles.) At this level of approximation, the free energy differences between [(a),(b)] and [(c),(d)] are equal,

$$\Phi^{\text{ex}}(C) \approx \Phi[0; \mu_S] - \Phi[E_C; \mu_S]. \quad (49)$$

Finally, the right-hand side can be estimated by the square gradient approximation (46), yielding $\Phi^{\text{ex}}(C) \approx \Phi^{\text{corr}}(C)$ with

$$\Phi^{\text{corr}}(C) = \Phi^{\text{sq}}[0; \mu_S] - \Phi^{\text{sq}}[E_C; \mu_S]. \quad (50)$$

Operation of the three-level method requires numerical estimates of this Φ^{corr} , which includes the integral in (46). Moreover, its value is exponentiated when computing weight factors \hat{W} , so these numerical estimates are required to high accuracy. This is a non-trivial requirement because the integrand is constant on regions far from the big particles, but it varies much more rapidly when these particles are approached. In such situations, adaptive quadrature schemes are appropriate: we use the *cuhre* algorithm of the *cuba* library⁵⁴ that uses globally adaptive subdivision to refine its approximations in the relevant regions of space. Note, however, that while the choice of the numerical integrator influences the intermediate level, small errors in estimation of this integral will be corrected by the second annealing step, so such errors do not affect the consistency of our numerical estimators.

Given this choice of $\Phi^{\text{corr}}(C)$, the intermediate level distribution p_I of (42) has been completely defined, although it still depends on the three parameters $\delta_{\text{free}}, s, l$ that appear in the function $\varepsilon(r)$. We also note that given the approximations made, it is not expected that this p_I is optimal (its marginal $p_I(C)$ does not match $p_F(C)$ perfectly). Subsection IV C discusses the parameter choices, and some possibilities for correction factors that can be added to Φ^{corr} in order to address specific sources of error.

C. Variants of the intermediate-level distribution

In fixing the parameters $\delta_{\text{free}}, s, l$, several considerations are relevant. First, if s is too small or δ_{free} is too large, the potential E_C has little effect on the system and the small particles are not restricted to be close to the large ones. In this case, p_I ends up close to p_F , and there is little benefit from the intermediate level. On the other hand, the accuracy of Φ^{sq} is greatest when the gradient of the potential E_C is small, and this favors small s and large l, δ_{free} . In practice, it is also convenient if the two annealing stages insert similar numbers of particles, so that their computational costs are similar. For the example system in Sec. II C, we will present results for a suitable parameter set

$$\delta_{\text{free}} = 0.5\sigma_S, \quad s = 4.4, \quad l = 3.5\sigma_S. \quad (51)$$

We have also tested other values, a few comments are given below.

We will consider several variants of the intermediate level. We denote by $p_I^{(1)}$ the distribution defined by (42) and (50), with parameters (51). Figure 8 shows how the quantities of interest differ between the CG and FG models, and the corresponding differences between the intermediate level and the FG model. Here, $\Delta g(r)$ is the difference between $g(r)$ for the FG model and the distribution of interest (which is either the CG distribution p_C or one of the variants of the intermediate distribution). $\Delta P(N)$ is the corresponding difference in the probability that the system has N large particles.

For the value of $g(r)$ at contact, we see that the intermediate level $p_I^{(1)}$ corrects around half of the deviation between CG and FG models. However, the probability distribution of the N has the opposite situation that the intermediate level is *less* accurate than the CG

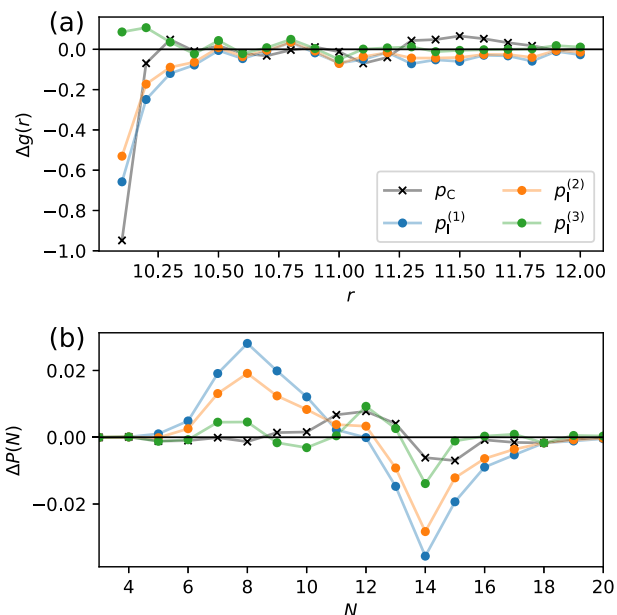


FIG. 8. Estimated accuracy of the CG and the intermediate levels from the main text for the example from Sec. II C. We show the difference between the respective CG and intermediate level estimates and the true FG estimate for the pair correlation function $g(r)$ in (a) and the distribution of big particles in (b). The FG estimates of these quantities of interest are shown in Fig. 3.

model. [This is partly attributable to the fact that $\Delta\mu$ in Eq. (6) has been chosen to make the CG model accurate.]

To explore the behavior of the intermediate level, we constructed two variants of p_I . The aim is to understand why $p_I^{(1)}$ has inaccuracies and to (partially) correct for them. There are two main approximations in the intermediate level $p_I^{(1)}$: the first is (49) and the second is that $\Phi[E_C, \mu_S]$ can be approximated by the square-gradient approximation (46). The first approximation neglects a significant physical phenomenon in these systems, which is a layering effect of the small particles around the large ones. This is illustrated in Fig. 9 by the radial distribution function g_{BS}^0 between large and small particles (measured in a system with a single large particle). One sees that there is typically an excess of small particles close to the large ones, followed by a deficit [$g_{BS}^0(r) < 1$], and a (weak) second layer.

For (49) to be accurate, the intermediate level should have enough small particles to capture this layering so that the particles being inserted in the second annealing stage are not strongly affected by the presence of the large particles. However, computational efficiency requires that δ_{free} is not too large, so these layers are not fully resolved at the intermediate level. To partially account for this effect, we make an *ad hoc* replacement of μ_S in (46) by an effective chemical potential $\mu_{\text{lay}}(\mathbf{r})$, which is chosen such that the corresponding reservoir volume fraction $\eta_S^{\text{lay}}(\mathbf{r})$ satisfies

$$\frac{\eta_S^{\text{lay}}(\mathbf{r})}{\eta_S^0} = g_{BS}^0(\text{dist}(\mathbf{r}, C)). \quad (52)$$

In estimating the free energy of the small particles that are inserted in the second level of annealing, this adjustment to Φ^{sq} helps to counteract the error made in (49), leading to an updated potential $\Phi^{\text{corr},2}$. The intermediate level constructed in this way is denoted by $p_I^{(2)}$. The results of Fig. 8 show that this variant is (somewhat) more accurate than $p_I^{(1)}$. However, the intermediate level still tends to have a smaller number of large particles than the full (FG) mixture.

To investigate this further, we took 800 representative CG configurations. For each one, we estimate the error associated with the approximation (50)

$$\Delta\Phi^{\text{corr}} = \Phi^{\text{ex}} - \Phi^{\text{corr},2}. \quad (53)$$

Results are shown in Fig. 10. One sees that the errors are of order unity (note that Φ^{ex} itself is of order 10^4 so this is a small relative




FIG. 9. The layering of small particles ($\sigma_S = 1$) around one big particle ($\sigma_B = 10$) at volume fraction $\eta_S = 0.2$. The displayed pair correlation function $g_{BS}^0(r)$ given the radius from the center of the big particle shows that small particles form layers of higher and lower concentration that vanish with increased distance from the big particle.




FIG. 10. The error of the free energy prediction for 800 configurations sampled from the coarse distribution p_C , grouped by their number of big particles. Each dot represents the difference of an estimate of the predicted small-particle free energy used in the intermediate level $p_I^{(2)}$ against an estimate of the full free energy. We correct for the noticeable trend in the error with a linear *ad-hoc* correction term, displayed in black.

error; see below); there is a systematic trend that $\Phi^{\text{corr},2}$ underestimates Φ^{ex} when N is large. To correct this error, we introduce an additional correction term to $\Phi^{\text{corr},2}$,

$$\Phi^{\text{corr},3}(C) = \Phi^{\text{corr},2}(C) + \alpha_{\text{corr}}N, \quad (54)$$

and denote the intermediate level constructed in this way by $p_I^{(3)}$.

A least squares fit to Fig. 10 suggests to take $\alpha_{\text{corr}} = 0.076$; in practice, this tends to over-correct the error in $\Phi^{\text{corr},2}(C)$, and we find better performance with a smaller value

$$\alpha_{\text{corr}} = 0.058. \quad (55)$$

However, the performance of the method depends only weakly on the specific choice of α_{corr} , which is discussed in Appendix C. For all following results, we define the intermediate level $p_I = p_I^{(3)}$ to use the potential $\Phi^{\text{corr},3}$.

D. Discussion of intermediate level

An important aspect of the three-level method is the self-consistency of the general approach. The intermediate level variants $p_I^{(1)}$ and $p_I^{(2)}$ were constructed on a purely theoretical basis. The corresponding results in Fig. 8 indicated good performance, but that the distribution of N had a systematic error. This error was quantified precisely in Fig. 10, which enabled an improvement to the intermediate level. In principle, this procedure could be repeated to develop increasingly accurate variants of p_I . That approach would be useful if (for example) one wanted to consider increasingly large systems, where the requirements for the accuracy of p_I become increasingly demanding.

One way to see the effect of system size is to note that Fig. 10 required the estimation of $\Phi_C[E_C, \mu_S]$ and $\Phi_C[0, \mu_S]$, whose values are of order 1×10^4 . Since the free energies are exponentiated in the weights for resampling, an absolute error of ± 1 is required on these free energies, while their absolute values are extensive in the system size. Hence, one sees that accurate estimates of the free energy are required: their relative error is required to be of the order of the inverse volume of the system.

V. NUMERICAL TESTS

In this section, we apply the three-level method to the example from Sec. II C using the intermediate level from Sec. IV, with the parameters defined in (51) and (55). The parameters and the annealing schedules are chosen such that, on average, the first and second steps have the same computational effort; see Appendix A for details.

It can be proven⁴⁹ that the three-level method provides accurate results, in the limit where the population sizes M_0, M_1, M_2 are all large. In particular, we expect the estimators $\hat{A}_F^{3L}, \hat{A}_{F,\Delta}^{3L}$ to all obey central limit theorems (CLTs), and the two-level estimators $\hat{A}_F, \hat{A}_{F,\Delta}$ behave similarly. Detailed results are given in Sec. VI. The important fact is that, for large populations, the variances of the estimators behave as

$$\text{Var}(\hat{A}) \approx \frac{1}{M} \Sigma, \quad (56)$$

where M is the relevant population size and Σ is called the asymptotic variance (it depends on the observable A and on which specific estimator is used). In general, the estimators may have a bias, which is also of order $1/M$. This means that the uncertainty in our numerical computations is dominated by the random error, whose typical size is $\sqrt{\Sigma/M}$, and the mean squared error is given by the variance $\text{MSE}(\hat{A}) = \text{Var}(\hat{A})$, to leading order.

This gives us an easy way to measure and compare the performance of the different estimators. Suppose that we require an estimate of A with a prescribed mean squared error. The associated computational cost can be identified with the population size M and is given by (56) as $M \approx \Sigma/\text{MSE}(\hat{A})$. Clearly, estimators with small Σ should be preferred. In practice, we do not compare computational costs at fixed error, instead we compare variances $\text{Var}(\hat{A})$ at fixed M . For any two algorithms (and assuming that M is large), the ratio of these variances approximates the ratio of the Σ 's, which can then be interpreted as a ratio of computational costs (at fixed MSE). Numerical results are presented in Sec. V B.

We note that the theoretical results for convergence do not require that the coarse or intermediate levels are accurate. However, one easily sees¹⁸ that serious inaccuracies in these levels lead to very large Σ . In such cases, one may require prohibitively large populations to obtain accurate results.

In this section, we demonstrate (for the example in Sec. II C) that we do not require very large populations for the three-level method and that the numerical results are consistent with (56). After that, we estimate the asymptotic variances for the two-level and three-level methods. We will find that introducing the third level improves the numerical performance, corresponding to a reduction in Σ .

To this end, we investigate the pair correlation $g(r)$ of the big particles. As seen in Fig. 8(a), the coarse approximation of $g(r)$ has a substantial error, especially when two big particles are in contact. To quantify this specific effect, we define the coordination number N_c , which is the number of large particles within a distance r_1 of a given large particle. (For a given configuration, this quantity is estimated as an average over the large particles. We take $r_1 \approx 10.73\sigma_S$ to be the first minimum of $g(r)$ of the CG model.) For our example, the coordination number for the FG and CG systems are given by

$$\langle N_c \rangle_F \approx 1.61, \quad \langle N_c \rangle_C \approx 1.56. \quad (57)$$

A. Accuracy of method

To illustrate the reliable performance of the method, we take a simple example with $M_0 = 4 \times 10^5$ and $M_1 = M_2$ (no tapering) and we focus on the difference estimator $\hat{A}_{F,\Delta}^{3L}$, which we expect to be the most accurate. The corresponding numerical estimate of $g(r)$ is denoted by $\hat{g}(r)$, binned using 40 equidistant bins at positions r_j between $r = 10$ and $r = 12$. Figure 11(a) shows estimates of the difference between $\hat{g}(r)$ and its true value, as the population size increases. (The FG result was estimated independently by the two-level method using a large value of $M_F = 18\,000$.) A population M_2 of several thousand is sufficient for an accuracy better than 0.5 in each bin of $g(r)$.

For smaller M_2 , fluctuations in the measured $\hat{g}(r)$ are apparent in Fig. 11(a). To estimate their size, we define the error for a single run of the three-level method by summing over the bins,

$$(\text{error})^2 = \sum_{j=1}^{40} |\hat{g}(r_j) - g(r_j)|^2. \quad (58)$$

Hence, one expects from (56) that this error decays with increasing population, proportional to $M_F^{-1/2}$. Figure 11(b) shows an estimate of (58), which is consistent with this expected scaling.

B. Measurements of variances Σ

We now investigate whether the three-level method does indeed improve on the performance of the (simpler) two-level

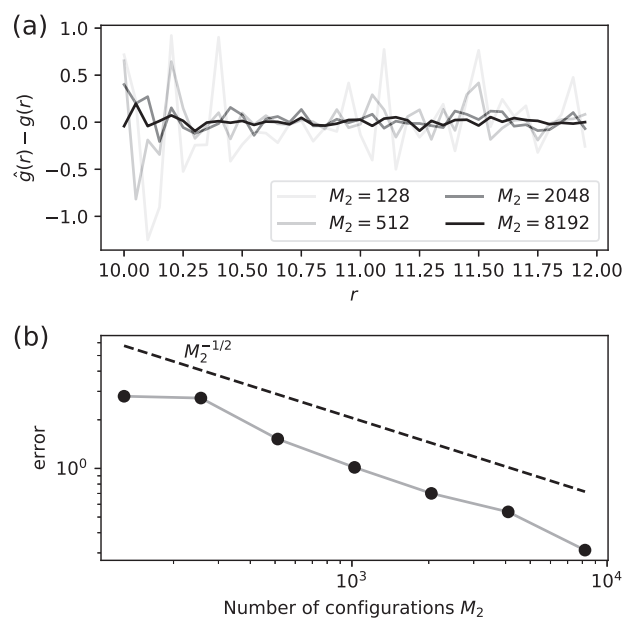


FIG. 11. Estimating the pair-correlation function $g(r)$ from Fig. 3(a) with the two- and three-level method. (a) The difference of three-level estimates $\hat{g}(r)$ with increasing numbers $M_1 = M_2$ number of particles against a reference value of $g(r)$ that was computed with the two-level method. (b) The error of the binned values in (a) as defined in (58). The dotted black line displays the expected asymptotic Monte Carlo convergence rate of $M_2^{-1/2}$.

method of Refs. 17 and 18. The key question is whether the resampling step is effective in focusing the computational effort on the most important configurations of the big particles.

We recall from above that removing the resampling step from the three-level method leads to a two-level method, where the annealing process is paused at the intermediate level and then restarted again. In order to test the effect of resampling, we compare these two schemes, keeping the other properties of the algorithm constant, including the annealing schedule. (To test the overall performance, one might also optimize separately the annealing schedules for the two-level and three-level algorithms and compare the total computational time for the two methods to obtain a result of fixed accuracy. However, such an optimization would be very challenging, so instead we focus on the role of resampling.)

As a very simple quantity of interest, we take the co-ordination number N_c . We run the whole algorithm N_{runs} independent times, and we estimate N_c for each run. This can be done using several different estimates of N_c . These are (i) the two-level estimates \hat{A}_F and $\hat{A}_{F,\Delta}$ from (13) and (18); (ii) the corresponding three-level estimates \hat{A}_F^{3L} and $\hat{A}_{F,\Delta}^{3L}$ of (31) and (35), in which we also vary the ratio $M_1 : M_2$, to see the effects of tapering.

All comparisons are done with a fixed total computational budget. We have chosen parameters such that the first and second annealing stages have the same (average) computational cost. This means we need to hold $M_T = (M_1 + M_2)/2$ constant during tapering. The two-level method takes $M_F = M_T$ (because the single step of annealing in the two-level method has the same cost as the two annealing steps of the three-level method). For the coarse level estimates \hat{A}_C and \hat{A}_C^{3L} (which are used in computation of $\hat{A}_{F,\Delta}$ and $\hat{A}_{F,\Delta}^{3L}$), the CG computations are cheap so we take $M_0 = M_C = 6 \times 10^6$. This is large enough that the numerical errors on these coarse estimates are negligible in comparison to the errors from higher levels.

For each version of the algorithm, we measure the sample variance of the N_{runs} estimates. Results are shown in Fig. 12 for $N_{\text{runs}} = 60$ and $M_T = 500$. The error bars are computed by the bootstrap method.⁵⁵ It is useful that the variance of all these estimators is expected to be proportional $1/M_T$: this means that reducing the variance by a factor of α requires that the computational effort is increased by the same factor. Hence, the ratio of variances of two




FIG. 12. The sample variance of $N_{\text{runs}} = 60$ independent estimates of the coordination number N_c . We compare results using a two-level method as well as three-level methods, with and without tapering. Furthermore, we give the results for the final-level (left) and difference (right) variants of the estimator. The error bars are computed via bootstrap; their interpretation is, however, not obvious as the different estimators are highly correlated; see the main text.

estimators is a suitable estimate for the ratio of their computational costs.

When carrying out these runs, each estimator was computed by performing annealing on the same set of coarse configurations, to ensure direct comparability. (More precisely, we take a set of 700 representative configurations that are used for the method with $M_1 : M_2 = 7 : 3$, and other versions of the method used a subset of these 700.) In addition, it is possible to share some of the annealing runs when computing the different estimators (while always keeping the 60 different runs completely independent). This freedom was exploited as far as possible, which reduces the total computational effort. However, it does mean that the calculations of the different estimators are not at all independent of each other.

C. Performance: Discussion

All three-level estimators have a reduced standard deviation compared to their two-level equivalents, demonstrating the usefulness of the intermediate resampling step. In all cases, the difference estimate outperforms its equivalent final-level estimate; this effect is stronger for the three-level estimate, providing evidence that the intermediate stop additionally improves the quality of the control variate in the difference estimates.

The effect of introducing tapering from $M_I = 600$ to $M_F = 400$ is difficult to assess, given the statistical uncertainties in this example. The variance of the tapered final-level estimator is very close to the non-tapered one, despite averaging over fewer configurations. This is possible since we start with more samples in the CG model that improves the sampling at the intermediate step, where we then resample to keep relevant configurations. As the results for the 700–300 tapering show, the tapering rate needs to be chosen carefully as a too aggressive rate can decrease the performance quickly.

Overall, the numerical tests in this section provide strong evidence of the benefit of the intermediate resampling. For our example, switching from a two-level to a three-level difference estimator substantially reduces the variance, from around 0.0029 for the two-level method to 0.0016 at a fixed computational budget. As discussed just below (56), the ratio of these numbers can be interpreted as the ratio of costs for the two- and three-level method: the conclusion for this case is that including the intermediate level reduces the cost by ~45%. This demonstrates a significant speedup in this specific case, which provides a proof-of-principle of the approach.

VI. CONVERGENCE OF THE MULTILEVEL METHOD

In Sec. V, we have seen that the three-level method outperforms the two-level method in numerical tests, both for the final-level as well as the difference version of the estimator. In this section, we provide convergence results for both algorithms and compare their asymptotic performance as the number of configurations goes to infinity.

The proof is general, but it does require some assumptions on the models of interest. First, for every allowed CG configuration [that is, configurations \mathcal{C} with $p_C(\mathcal{C}) > 0$], we assume that the quantity of interest A is bounded. In addition, the probability density $p_C(\mathcal{C})$ must be non-zero whenever $p_I(\mathcal{C})$ is non-zero, and similarly $p_I(\mathcal{C}, \mathcal{F})$ must be non-zero whenever $p_F(\mathcal{C}, \mathcal{F})$ is non-zero.

A. Two-level method

The two-level method has been previously analyzed in Ref. 18. We summarize its key properties. It was noted in Sec. III B that $\hat{A}_F \rightarrow \langle A \rangle_F$ as $M_F \rightarrow \infty$ (specifically, this is convergence in probability⁵⁶). We also expect a CLT for this quantity: as in Eq. (56), the distribution of the error ($\hat{A}_F - \langle A \rangle_F$) converges to a Gaussian with mean zero and variance Σ_F/M_F . We will derive a formula for this variance, which will be compared later with the corresponding quantity for the three-level model.

For compact notation, it is convenient to define the re-centered quantity of interest

$$A^r(C) = A(C) - \langle A \rangle_F. \quad (59)$$

A significant contribution to Σ_F comes from the randomness of the annealing procedure, and this can be quantified as

$$v(C) = \text{Var}_J[\hat{W}^n(C)], \quad (60)$$

where the variance is again with respect to the annealing procedure (from coarse to fine). Then, following Ref. 18, it can be shown that

$$\Sigma_F = \langle A^r(C)^2 [w(C)^2 + v(C)] \rangle_C, \quad (61)$$

where $w(C) = \langle \hat{W}^n(C) \rangle_J = p_F(C)/p_C(C)$, so one identifies $w(C)^2 + v(C)$ as the mean square weight obtained from the annealing procedure. Similarly, the estimator $\hat{\Delta}$ that appears in the difference estimate $\hat{A}_{F,\Delta}$ also obeys a CLT, with variance $\text{Var}(\hat{\Delta}) \approx \Sigma_{F,\Delta}/M_F$, where

$$\Sigma_{F,\Delta} = \left\langle A^r(C)^2 [w(C)^2 + v(C) - 1] \right\rangle_C + \text{Var}_C(A) - \text{Var}_F(A). \quad (62)$$

As discussed in Ref. 18, if the computational cost of the coarse model is low, M_C can be taken large enough that the variance of the coarse estimator \hat{A}_C is negligible, in which case (18) implies $\text{Var}(\hat{A}_{F,\Delta}) \approx \text{Var}(\hat{\Delta})$, and hence

$$\text{Var}(\hat{A}_{F,\Delta}) \approx \frac{1}{M_F} \Sigma_{F,\Delta}. \quad (63)$$

Comparing (61) and (62)—which give the variances of \hat{A}_F and $\hat{A}_{F,\Delta}$, respectively—the term $v(C)$ in (61) is replaced by $v(C) - 1$ in (62), which reduces the variance of the estimator. We expect, in general, that $\text{Var}_C(A)$ and $\text{Var}_F(A)$ should be similar in magnitude, in which case these terms in (62) should have little effect. Hence, one expects that the estimator $\hat{A}_{F,\Delta}$ has lower variance than \hat{A}_F . This is consistent with the results of Fig. 12.

B. Three-level method

The results (61) and (63) are based on the property that each estimator is a sum of (nearly) independent random variables, which means that we can immediately apply standard Monte Carlo convergence results.⁴¹ This is not possible for the three-level method, since the resampling step correlates the configurations. This makes the analysis of SMC-type algorithms challenging, but widely applicable results are available.^{43,57,58} The three-level method in Sec. III C is an implementation of a random-weight SMC method that has been analyzed in Ref. 49.

To analyze the variance of the three-level method, we require results analogous to (61), which depend on the mean square weights

associated with the annealing procedure. To this end, define the average of the final level weight

$$w_2(\mathcal{X}_2) = \langle \hat{W}_2^n(\mathcal{X}_2) \rangle_J, \quad (64)$$

which fulfills (29). Similar to (60), the variance of this weight is

$$v_2(\mathcal{X}_2) = \text{Var}_J[\hat{W}_2^n(\mathcal{X}_2)]. \quad (65)$$

The averages in these equations are with respect to the second annealing step (from intermediate to fine level), starting at configuration \mathcal{X}_2 (see Sec. III C 3).

For the contribution to the asymptotic variance of the first annealing step, it is important to consider a product of weight factors: $\hat{W}_1^n(C_1)w_2(\hat{\mathcal{X}}_1)$. The first factor in this product is the random weight \hat{W}_1^n that is obtained by annealing from the coarse to the intermediate level, leading to the intermediate configuration is $\hat{\mathcal{X}}_1 = (C_1, \hat{\mathcal{F}}_1)$. The second factor is the averaged weight $w_2(\hat{\mathcal{X}}_1)$ from (64) associated with the second (subsequent) annealing step. Combining (26) and (29), the average of the product is

$$\langle \hat{W}_1^n(C_1)w_2(\hat{\mathcal{X}}_1) \rangle_J = p_F(C_1)/p_C(C_1) = w(C_1), \quad (66)$$

and the corresponding variance is

$$v_1(C_1) = \text{Var}_J[\hat{W}_1^n(C_1)w_2(\hat{\mathcal{X}}_1)]. \quad (67)$$

Hence, $w(C_1)^2 + v_1(C_1)$ is the mean square value of $\hat{W}_1^n(C_1)w_2(\hat{\mathcal{X}}_1)$ with respect to the annealing process: this turns out to be a relevant quantity for the asymptotic variance.

The number of configurations M_1, M_2 can be varied between steps of the three-level method. We formulate the asymptotic variance in the average number of configurations

$$M_T = \frac{1}{2}(M_1 + M_2). \quad (68)$$

If the two annealing steps have comparable cost, we can then directly compare the variances for different tapering rates at fixed M_T . Define also

$$c = \frac{M_1}{2M_T}, \quad \bar{c} = 1 - c = \frac{M_2}{2M_T}. \quad (69)$$

Then, a direct application of Theorem 2.1 of Ref. 49 gives a CLT for \hat{A}_F^{3L} : for large M_T , we have

$$\text{Var}(\hat{A}_F^{3L}) \approx \frac{1}{M_T} \Sigma_F^{3L} \quad (70)$$

with asymptotic variance

$$\Sigma_F^{3L} = \frac{1}{2c} \Sigma_{F,1}^{3L} + \frac{1}{2\bar{c}} \Sigma_{F,2}^{3L} \quad (71)$$

with

$$\begin{aligned} \Sigma_{F,1}^{3L} &= \langle A^r(C)^2 [w(C)^2 + v_1(C)] \rangle_C, \\ \Sigma_{F,2}^{3L} &= \langle A^r(C)^2 [w_2(\mathcal{X})^2 + v_2(\mathcal{X})] \rangle_J. \end{aligned} \quad (72)$$

The physical interpretation of these formulas will be discussed in Subsection VI C.

Computing the asymptotic variance of the three-level difference estimator $\hat{A}_{F,\Delta}^{3L}$ is more difficult, since it involves difference of non-trivially correlated samples. For some examples of multi-level difference estimators, upper bounds on the asymptotic variance have been developed in Refs. 35 and 36. A detailed analysis of these bounds in the context of our algorithm is beyond the scope of this paper.

C. Discussion of CLTs

To understand the differences between the two- and three-level method, we compare the asymptotic variances of their corresponding final level estimators Σ_F in (61) and Σ_F^{3L} in (71). The variance of the three-level method has two contributions $\Sigma_{F,1}^{3L}$ and $\Sigma_{F,2}^{3L}$; they are the variances of two-level methods from the coarse to the fine model and the intermediate to the fine model, respectively. The first term $\Sigma_{F,1}^{3L}$ is therefore directly related to Σ_F , where the variance of the importance weight $v(C)$ has been replaced by $v_1(C)$.

In order to make quantitative comparisons, we again consider the three-level method without intermediate resampling. As discussed in Sec. V B, this is a two-level method with a specific annealing process that consists of the concatenation of the two annealing processes of the three-level method. For the concatenated annealing process, we have

$$\hat{W}^n(C) = \hat{W}_1^n(C) \hat{W}_2^n(\hat{\mathcal{X}}), \quad (73)$$

where $\hat{\mathcal{X}} = (C, \hat{\mathcal{F}})$ is generated by the first annealing stage. This means that

$$v(C) = \text{Var}_1[\hat{W}_1^n(C) \hat{W}_2^n(\hat{\mathcal{X}})], \quad (74)$$

where the variance is now over the randomness of both annealing processes. Comparing (74) to (67), we see that $v_1(C)$ computes the variance of the same importance weight, but after averaging over the second annealing stage in (64). We can apply Jensen's inequality⁵⁶ to show that

$$v(C) \geq v_1(C). \quad (75)$$

By definitions (61) and (71), this directly implies

$$\Sigma_{F,1}^{3L} \leq \Sigma_F. \quad (76)$$

For the case without tapering $c = 1/2$, the three-level method, therefore, trades a reduction in the variance of the importance weights from coarse to fine in $\Sigma_{F,1}^{3L}$ for the addition of a term $\Sigma_{F,2}^{3L}$ that corresponds to the variance of a two-level method going from the intermediate to the fine level. The possibility of tapering, i.e., $c \neq 1/2$, further allows us to optimize the distribution of computation effort between the two stages, which is particularly useful if $\Sigma_{F,2}^{3L} \ll \Sigma_{F,1}^{3L}$.

For our application to hard sphere mixture example in Sec. II C, the annealing process is computationally expensive and the resulting weights are noisy. We are therefore in the situation where the variance $v(C)$ contributes substantially to the overall variance, and where we have constructed an intermediate in Sec. IV that improves on the CG model. Following the discussion above, this is the setting where we expect the three-level method to improve upon a two-level method, which is confirmed by the numerical results in Sec. V. Further discussion of the effect of resampling on random-weight SMC methods can be found in Ref. 49.

VII. CONCLUSIONS

We have introduced a three- and multilevel extension of the two-level simulation method first discussed in Ref. 18. We have applied this method to a highly size-asymmetric binary hard-sphere system. As shown in the numerical test in Sec. V and theoretical results in Sec. VI, the introduction of intermediate resampling that distinguishes the two-level from the three-level method can lead to substantial improvements in performance by reducing the variance in importance weights and by allowing efficient allocation of resources between levels via tapering.

A. Hard sphere model

In the application to binary hard-sphere systems, the introduction of an intermediate level required us to construct a semi-analytic estimate of the free energy of a system with partially inserted small particles. For this, we have combined a highly accurate square-gradient theory with pre-computed *ad hoc* corrections, yielding an intermediate level that substantially improves the accuracy of the investigated quantities of interest compared to the initial coarse level. Furthermore as we show in Appendix C, the three-level method appears robust with respect to slight deviations of the intermediate level.

Compared to our numerical example, Ref. 17 applied the two-level method to larger and more dense systems than considered here, to investigate the critical point of demixing. This was achieved by replacing the two-body CG model with RED potential used in this publication by a highly accurate two- and three-body potential. The computation of accurate effective potentials entails a substantial upfront computational cost (compared to our construction of the intermediate level), but for the hard sphere mixtures, this results in a CG level that is more accurate than our intermediate level. Despite the challenges of keeping the variance of the importance weights under control for large systems, this turned out to be more efficient overall.

B. Design principles for other potential applications

We have emphasized throughout that our three-level modeling approach is generally applicable whenever a suitable intermediate level can be constructed. We can identify two main scenarios where this might be attempted. The first scenario is illustrated by the binary hard sphere mixture, which is a two-scale system by construction (there are two species). In this situation, there is no obvious intermediate level, and a careful construction is required, to design one. Our results show that this strategy is possible—it is worthwhile in this example because the system is very challenging to characterize by other methods, so the effort of constructing the intermediate level is worthwhile.

The second scenario—where we may expect a multilevel method to be particularly useful—is that a multi-scale system admits a true hierarchy of coarse-graining, such as a system of long-chain polymers. We can coarse-grain a polymer chain by representing groups of monomers by their center of masses with suitable effective interactions.^{5,59} By varying the number of monomers per group, we get a hierarchy of CG models that could be targeted by a multilevel method. For such methods to be efficient in such a scenario, we require high accuracy of the CG models, and an efficient

annealing process to introduce the finer degrees of freedom analogous to the introduction of the small spheres in the hard sphere mixture. Fulfilling these requirements is still challenging and requires considerable physical insight about the specific polymer system of interest, but the hierarchical structure of the system hints that a suitable method might be fruitfully extended to more than three levels, with commensurately increased performance gains.

In both scenarios, careful thought is required to apply the three-level (or multi-level) methods: our approach is far from being a black-box method. Still, the results presented here show that it can be applied in a practical (challenging) computational problem.

A separate limitation of multilevel methods is that the population of unique coarse configurations is fixed from the start and reduces with each subsequent resampling step. This is closely related to the sample depletion effect commonly observed effect in particle filtering, and SMC methods in general.^{25,60} For the multilevel method, we can address this by following each resampling step with a number of Markov Chain Monte Carlo (MCMC) steps to decorrelate duplicated configurations and further explore the system at the current level of coarse-graining.⁶⁰ While such an approach is not feasible for the hard-sphere system where intermediate MCMC is limited by the cost of computing the required approximations, we expect this to be beneficial, for example, whenever intermediate physical systems are described in terms of effective, few-body interactions.

We end with a comment on the implementation of these methods. The introduction of intermediate levels increases the complexity of the code required to simulate the systems. It requires adding an intermediate stage to the annealing process and computing the required integrals (see Sec. IV). Additionally, when implementing the algorithm for the use on compute clusters, the resampling step requires the communication between all nodes. However, we emphasize that while these extra steps require some extra programming, none of the additional steps of the three-level method have added significant computational cost in our example.

To conclude, our results show that the multilevel method can effectively make use of intermediate levels when available, leading to improvements in performance at fixed computational cost. We look forward to further applications of multilevel methods in physical simulations.

ACKNOWLEDGMENTS

This work was supported by the Leverhulme Trust through Grant No. RPG-2017-203.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Paul B. Rohrbach: Investigation (lead); Methodology (lead); Validation (equal); Writing – original draft (lead). **Hideki Kobayashi:** Investigation (supporting); Methodology (supporting). **Robert Scheichl:** Conceptualization (equal); Funding acquisition (equal);

Writing – review & editing (equal). **Nigel B. Wilding:** Conceptualization (equal); Funding acquisition (equal); Writing – review & editing (equal). **Robert L. Jack:** Conceptualization (equal); Funding acquisition (equal); Supervision (equal); Writing – original draft (supporting).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

APPENDIX A: ENSEMBLE DEFINITIONS, AND ESTIMATION OF PARTITION FUNCTIONS

1. Grand canonical ensemble

We define the grand canonical ensemble of the hard sphere mixture discussed in Sec. II. Recall that $k_B T = 1$. For the system of interest, the equilibrium average of a quantity of interest $A(C)$ in (3) is defined as

$$\langle A \rangle_F = \frac{1}{\Xi_F} \sum_{N=0}^{\infty} \sum_{n=0}^{\infty} \frac{e^{\mu_B N + \mu_S n}}{n! N! \sigma_B^{3N} \sigma_S^{3n}} \times \int A(C) e^{-U_F(C, \mathcal{F})} d\mathbf{R}_1 \cdots d\mathbf{R}_N d\mathbf{r}_1 \cdots d\mathbf{r}_n, \quad (\text{A1})$$

where each particle position is integrated over the periodic domain $[0, L]^3$. For ease of notation, we introduce the integration measures $dC, d\mathcal{F}$ that include the prefactors accounting for the indistinguishability of particles that appear in (A1), which then becomes

$$\langle A \rangle_F = \frac{1}{\Xi_F} \int A(C) e^{\mu_B N + \mu_S n - U_F(C, \mathcal{F})} dC d\mathcal{F}, \quad (\text{A2})$$

consistent with (1) and (3). By definition, we require that p_F is normalized as $\int p_F(C, \mathcal{F}) dC d\mathcal{F} = 1$, so we have

$$\Xi_F = \sum_{N=0}^{\infty} \sum_{n=0}^{\infty} \frac{e^{\mu_B N + \mu_S n}}{n! N! \sigma_B^{3N} \sigma_S^{3n}} \int e^{-U_F(C, \mathcal{F})} d\mathbf{R}_1 \cdots d\mathbf{R}_N d\mathbf{r}_1 \cdots d\mathbf{r}_n. \quad (\text{A3})$$

The relevant quantities of the CG model $\langle A \rangle_C$ are defined analogously.

2. Estimation of the partition function

The implementation of the two- and three-level methods requires the computation of the small-particle partition function that appears in the importance weights. We use a method based on Jarzynski's equality^{20,21} that yields an unbiased estimator; see also Ref. 18. In the statistics literature, this is also known as annealed importance sampling.²² We first give a short summary of the method in Appendix A 2 a and then discuss how the annealing processes are implemented for the two- and three-level method in Appendix A 2 b. The parameters used for the numerical tests are given in Appendix A 2 c.

a. Theoretical details

We derive an annealing process that inserts the small particles \mathcal{F} for a fixed big particle configuration C . This process produces weighted configurations that correctly characterize the FG distribution. We closely follow the results from Appendix A of Ref. 18; see also Refs. 20–22 and 61.

Let $p_s(\mathcal{C}, \mathcal{F})$ and $p_e(\mathcal{C}, \mathcal{F})$ be two probability distributions for the FG model of the form

$$p_\alpha(\mathcal{C}, \mathcal{F}) = \frac{1}{\Xi_\alpha} e^{\Phi_\alpha(\mathcal{C}, \mathcal{F})}, \quad \alpha \in \{s, e\}. \quad (\text{A4})$$

The corresponding marginal distributions are $p_\alpha(\mathcal{C}) = \int p_\alpha(\mathcal{C}, \mathcal{F}) d\mathcal{F}$. The distributions p_s, p_e are the start and end point of an annealing process, with a sequence of intermediate distributions

$$p_k(\mathcal{C}, \mathcal{F}) = \frac{1}{\Xi_k} e^{\Phi_k(\mathcal{C}, \mathcal{F})}, \quad k = 0, \dots, K, \quad (\text{A5})$$

where $p_0 = p_s$ and $p_K = p_e$.

Let \mathcal{C} be a sample from $p_s(\mathcal{C})$: this configuration remains fixed during the annealing process. We anneal the small particles, as follows: first sample an initial small particle configuration \mathcal{F} from $p_s(\mathcal{F} | \mathcal{C})$, the conditional distribution of p_s . This distribution is p_0 so write $\mathcal{F}_0 = \mathcal{F}$ and set $k = 1$: then apply a sequence of MC steps with transition kernel $q_{c,k}(\mathcal{F}_{k-1} \rightarrow \mathcal{F}_k)$ that is in detailed balance with the small particle distribution $p_k(\mathcal{F} | \mathcal{C})$. Iterate this process for $k = 1, \dots, K - 1$: this yields a sequence of small-particle configurations $(\mathcal{F}_0, \mathcal{F}_1, \dots, \mathcal{F}_{K-1})$. The big-particle configuration \mathcal{C} stays fixed throughout this process.

The relevant results of this procedure are the final small-particle configuration $\hat{\mathcal{F}} = \mathcal{F}_{K-1}$ and an annealing weight

$$\hat{W}_A = e^{\sum_{k=1}^K [\Phi_k(\mathcal{C}, \mathcal{F}_{k-1}) - \Phi_{k-1}(\mathcal{C}, \mathcal{F}_{k-1})]}. \quad (\text{A6})$$

Given the initial coarse configuration $(\mathcal{C}, \mathcal{F})$, the MC steps define a probability distribution over the weight \hat{W}_A and the final small particle configuration $\hat{\mathcal{F}}$, which we denote by

$$\kappa(\hat{W}_A, \hat{\mathcal{F}} | \mathcal{C}, \mathcal{F}). \quad (\text{A7})$$

Given the initial configuration $(\mathcal{C}, \mathcal{F})$, averages with respect to the annealing process are denoted by $\langle \cdot \rangle_J$.

We now show that this annealing process produces weighted samples of p_e , up to a constant. More specifically,

$$\int \hat{W}_A \kappa(\hat{W}_A, \hat{\mathcal{F}} | \mathcal{C}, \mathcal{F}) p_s(\mathcal{F} | \mathcal{C}) d\hat{W}_A d\mathcal{F} = \frac{\Xi_e p_e(\mathcal{C}, \hat{\mathcal{F}})}{\Xi_s p_s(\mathcal{C})}. \quad (\text{A8})$$

This implies that averaging over the start distribution p_s and the annealing process yields

$$\langle \hat{W}_A B(\mathcal{C}, \hat{\mathcal{F}}) \rangle_{J, p_s} = \frac{\Xi_e}{\Xi_s} \langle B(\mathcal{C}, \mathcal{F}) \rangle_{p_e} \quad (\text{A9})$$

for any function $B = B(\mathcal{C}, \mathcal{F})$, which may depend on both big and small particles.

To show (A8), we compute the average over the annealing process explicitly

$$\begin{aligned} & \int \hat{W}_A \kappa(\hat{W}_A, \hat{\mathcal{F}} | \mathcal{C}, \mathcal{F}) p_s(\mathcal{F} | \mathcal{C}) d\hat{W}_A d\mathcal{F} \\ &= \int e^{\sum_{k=1}^K (\Phi_k(\mathcal{C}, \mathcal{F}_{k-1}) - \Phi_{k-1}(\mathcal{C}, \mathcal{F}_{k-1}))} \prod_{k=1}^{K-1} q_{c,k}(\mathcal{F}_{k-1} \rightarrow \mathcal{F}_k) \\ & \quad \times p_s(\mathcal{F}_0 | \mathcal{C}) d\mathcal{F}_0 \cdots d\mathcal{F}_{K-2}. \end{aligned} \quad (\text{A10})$$

By rearranging the factors in the exponential, the right-hand side of (A10) becomes

$$\begin{aligned} & \int \prod_{k=1}^{K-1} q_{c,k}(\mathcal{F}_{k-1} \rightarrow \mathcal{F}_k) e^{\Phi_k(\mathcal{C}, \mathcal{F}_{k-1}) - \Phi_k(\mathcal{C}, \mathcal{F}_k)} \\ & \quad \times e^{\Phi_K(\mathcal{C}, \mathcal{F}_{K-1})} \frac{p_s(\mathcal{F}_0 | \mathcal{C})}{e^{\Phi_0(\mathcal{C}, \mathcal{F}_0)}} d\mathcal{F}_0 \cdots d\mathcal{F}_{K-2}. \end{aligned} \quad (\text{A11})$$

Detailed balance of the Markov kernels $q_{c,k}$ implies

$$q_{c,k}(\mathcal{F}_{k-1} \rightarrow \mathcal{F}_k) e^{\Phi_k(\mathcal{C}, \mathcal{F}_{k-1}) - \Phi_k(\mathcal{C}, \mathcal{F}_k)} = q_{c,k}(\mathcal{F}_k \rightarrow \mathcal{F}_{k-1}), \quad (\text{A12})$$

and by definition

$$e^{\Phi_K(\mathcal{C}, \mathcal{F}_{K-1})} \frac{p_s(\mathcal{F}_0 | \mathcal{C})}{e^{\Phi_0(\mathcal{C}, \mathcal{F}_0)}} = \frac{\Xi_e p_e(\mathcal{C}, \mathcal{F}_{K-1})}{\Xi_s p_s(\mathcal{C})}. \quad (\text{A13})$$

Using (A12) and (A13), Eq. (A11) simplifies to

$$\int \prod_{k=1}^{K-1} q_{c,k}(\mathcal{F}_k \rightarrow \mathcal{F}_{k-1}) d\mathcal{F}_0 \cdots d\mathcal{F}_{K-2} \frac{\Xi_e p_e(\mathcal{C}, \mathcal{F}_{K-1})}{\Xi_s p_s(\mathcal{C})}. \quad (\text{A14})$$

Since $q_{c,k}(\mathcal{F}_k \rightarrow \mathcal{F}_{k-1})$ is a normalized probability density for \mathcal{F}_{k-1} , we can perform the integrals in (A14) one by one, yielding (A8).

b. Application to the two- and three-level method

This section describes the details of the annealing processes used in the two- and three-level method. We first discuss its implementation for the two-level method before showing how to split this process into two stages for the three-level method.

The two-level method starts with samples \mathcal{C} of the CG model p_C . We describe the annealing process,¹⁸ which produces a weight $\hat{W}(\mathcal{C})$ and small particle configuration $\hat{\mathcal{F}}$ that fulfills (14). Since we have no initial small particle distribution, we cannot directly apply the results of Appendix A 2 a and need to proceed in two steps. Let

$$p_{\mu_s}(\mathcal{F} | \mathcal{C}) = p_F(\mathcal{F} | \mathcal{C}) = \frac{1}{\Xi[\mathcal{C}, \mu_s]} e^{\mu_s n - U_F(\mathcal{C}, \mathcal{F})} \quad (\text{A15})$$

be the distribution of small particles around a fixed big-particle configuration \mathcal{C} , where we now explicitly note the dependence on the small particle chemical potential μ_s . Computing the unnormalized importance weight $\hat{W}(\mathcal{C})$ from (20) requires an estimate of the partition function of the small particles $\Xi[\mathcal{C}, \mu_s]$. For a system with a small value of the chemical potential $\mu_0 \ll \mu_s$, we can directly estimate this quantity

$$\Xi[\mathcal{C}, \mu_0] = \frac{1}{\mathbb{P}_{p_{\mu_0}(\cdot | \mathcal{C})}(n = 0)} \quad (\text{A16})$$

as it is the reciprocal probability of having zero small particles in a system with fixed \mathcal{C} . For small enough μ_0 , this value is close to 1 and can be estimated quickly by a GCMC simulation that decorrelates quickly due to the low density of small particles. Since we can compute this value to a very low variance at negligible cost, we consider our estimate of $\Xi[\mathcal{C}, \mu_0]$ to be exact and we neglect the influence of its fluctuations on the overall variance of the method. Furthermore, we assume that we can generate samples from the low chemical potential distribution of small particles $p_{\mu_0}(\mathcal{F} | \mathcal{C})$.

Starting with a sample \mathcal{F} of the initial small particle distribution $p_{\mu_0}(\mathcal{F} | \mathcal{C})$, we can now apply steps of the annealing process defined in Appendix A 2 a. We define the steps of the annealing process by slowly increasing the chemical potential of the small particles μ_k in $K + 1$ steps, from μ_0 to $\mu_K = \mu_S$ while keeping the CG distribution fixed. More specifically, we simulate an annealing process for the sequence of probability distributions

$$p_{\mathcal{C}, \mu_k}(\mathcal{C}, \mathcal{F}) = p_{\mathcal{C}}(\mathcal{C}) p_{\mu_k}(\mathcal{F} | \mathcal{C}), \quad k = 0, \dots, K, \quad (\text{A17})$$

yielding an annealing weight \hat{W}_A and a fine-particle configuration $\hat{\mathcal{F}}$ as described previously.

Averaging over the initial distribution of small particles and the annealing process and using (A8), (A15), and (A17) yields

$$\int \hat{W}_A \kappa(\hat{W}_A, \hat{\mathcal{F}} | \mathcal{C}, \mathcal{F}) p_{\mu_0}(\mathcal{F} | \mathcal{C}) d\hat{W}_A d\mathcal{F} = \frac{\Xi[\mu_S, \mathcal{C}]}{\Xi[\mu_0, \mathcal{C}]} p_{\text{F}}(\hat{\mathcal{F}} | \mathcal{C}). \quad (\text{A18})$$

Combining this with (19), we have

$$\frac{\Xi[\mu_S, \mathcal{C}]}{\Xi[\mu_0, \mathcal{C}]} p_{\text{F}}(\hat{\mathcal{F}} | \mathcal{C}) = \frac{\Xi_{\text{F}}}{\Xi_{\text{C}}} \frac{p_{\text{F}}(\mathcal{C}, \hat{\mathcal{F}})}{p_{\text{C}}(\mathcal{C})} \frac{e^{-U_{\text{C}}(\mathcal{C})}}{\Xi[\mu_0, \mathcal{C}]}. \quad (\text{A19})$$

Thus, we scale the weight that is produced by the annealing process

$$\hat{W}(\mathcal{C}) = \Xi[\mu_0, \mathcal{C}] e^{U_{\text{C}}(\mathcal{C})} \hat{W}_A. \quad (\text{A20})$$

For this weight, the annealing process fulfills (14) when we include the sampling from the distribution of the initial small particles $\mathcal{F} \sim p_{\mu_0}(\mathcal{F} | \mathcal{C})$ as part of the annealing process.

For the three-level method, we split the annealing process discussed above into two consecutive steps. The first part follows exactly the same steps as above, where the annealing process increases the chemical potential μ_k from a small value μ_0 to μ_S . The only difference is that we include the potential \tilde{U} of the intermediate distribution: in place of (A15), we have

$$p_{\mu_S}(\mathcal{F} | \mathcal{C}) = p_{\text{I}}(\mathcal{F} | \mathcal{C}) \quad (\text{A21})$$

so that small particle insertion is suppressed in regions far from large particles. As before, the annealing process results in a (scaled) weight $\hat{W}_1(\mathcal{C})$ and small particle configuration $\hat{\mathcal{F}}_1$ that now fulfills (26).

For the second step of the three-level method, we need to define an annealing process that fulfills (28). We start with a sample $\mathcal{X}_2 = (\mathcal{C}_2, \mathcal{F}_2)$ from the intermediate level p_1 . Since this configuration already contains small particles, we can directly apply the results of Appendix A 2 a to anneal from $p_{\text{I}}(\mathcal{C}, \mathcal{F})$ to $p_{\text{F}}(\mathcal{C}, \mathcal{F})$. This is achieved by a sequence of intermediate annealing distributions that increase the parameter δ of the potential (40) so that the volume available to the small particles is slowly increased. This is done in $K + 1$ steps from the parameter $\delta_0 = \delta_{\text{free}}$ (the intermediate level) to a final value

$$\delta_K = \max_{\mathbf{r}} \text{dist}(\mathbf{r}, \mathcal{C}_2), \quad (\text{A22})$$

at which point the suppression potential does not affect any point in the domain. Then, the intermediate level with δ_K corresponds to the fine-level distribution, up to the correction factor $e^{\Phi_{\text{corr}}(\mathcal{C})}$ in (42) that only depends on the big particles. Following from (A8),

this annealing process with scaled weight $\hat{W}_2(\mathcal{C}) = \hat{W}_A e^{-\Phi_{\text{corr}}(\mathcal{C})}$ and final small particle configuration $\hat{\mathcal{F}}_2$ fulfills the property (28).

c. Annealing schedules for simulations

The importance weights produced by the annealing process are unbiased, and this feature is independent of the details of the annealing schedules. In this sense, the algorithm is valid for any schedule. However, the variance of the computed importance weights depends strongly on the choice of schedule.

The initial chemical potential μ_0 is chosen such that, in a system with no big particles, and there would be an average of $\bar{n}_0 = 0.01$ small particles present, that is, the initial reservoir volume fraction is $\eta_{\text{S},0}^r = 0.01/L^3$.

For the first stage annealing process, we increase the chemical potential in steps $\Delta\mu_k$ such that the average change in the number of small particles would be $\delta\bar{n}_k = 0.2$, in a system where no big particles were present. For the second stage, we increase δ_k in fixed steps $\Delta\delta_k = \sigma_{\text{S}}/20\,000$. In both cases, we run one GCMC sweep between each step.

To compute accurate FG reference results used, for example, in Figs. 4 and 11, we apply the two-level method using the same annealing strategy as for the first stage of the three-level method but with $\delta\bar{n}_k = 0.05$ for increased accuracy. Note that, for the numerical tests in Figs. 12 and 15 that directly compare the performance of the two- and three-level method, the two-level method uses the same annealing schedule as the three-level method outlined earlier. The only difference is the lack of resampling.

APPENDIX B: DETAILS OF THE INTERMEDIATE LEVEL

1. Perturbative approximation of non-homogeneous hard-sphere fluid

This section derives (47) of the main text. To this end, consider a homogeneous hard sphere fluid at chemical potential μ and add a perturbing potential

$$\mathcal{E}(\mathbf{r}) = a \sin(\mathbf{q} \cdot \mathbf{r}) \quad (\text{B1})$$

in (44). In a finite periodic system, \mathbf{q} should be a reciprocal lattice vector. We aim to estimate the free energy difference between the perturbed and homogeneous system. For this, we follow the steps of the local density approximation discussed in Ref. 52; see also Chap. 6 of Ref. 62. We approximate this difference as

$$\begin{aligned} \delta\Phi[\mathcal{E}] &= \Phi - \Phi[\mathcal{E}] \\ &\approx \int \mathbf{p}(\mu - \mathcal{E}) - \mathbf{p}(\mu) + \mathbf{g}(\mu - \mathcal{E}) |\nabla \mathcal{E}|^2 d\mathbf{r}. \end{aligned} \quad (\text{B2})$$

To compute \mathbf{g} , we need to investigate both sides of Eq. (B2). Starting with the right-hand side, we assume that \mathbf{g} is smooth; therefore, we can approximate it for small a by a constant $\mathbf{g}(\mu - \mathcal{E}) \equiv \mathbf{g}(\mu)$. To compute the integral of the pressure difference in (B2), we expand around μ ,

$$\begin{aligned} \int \mathbf{p}(\mu - \mathcal{E}) - \mathbf{p}(\mu) d\mathbf{r} &= \int -\mathbf{p}'(\mu) \mathcal{E}(\mathbf{r}) + \frac{\mathbf{p}''(\mu)}{2} \mathcal{E}(\mathbf{r})^2 d\mathbf{r} + \mathcal{O}(a^3) \\ &= \frac{V}{4} a^2 \mathbf{p}''(\mu) + \mathcal{O}(a^3), \end{aligned} \quad (\text{B3})$$

and integrate the gradient correction term

$$\begin{aligned} \int \mathbf{g}(\mu) |\nabla \mathcal{E}|^2 \mathbf{dr} &= q^2 a^2 \mathbf{g}(\mu) \int \cos^2(\mathbf{q} \cdot \mathbf{r}) \mathbf{dr} \\ &= \frac{V}{2} q^2 a^2 \mathbf{g}(\mu). \end{aligned} \quad (\text{B4})$$

Overall, we obtain

$$\int \mathbf{p}(\mu - \mathcal{E}) - \mathbf{p}(\mu) + \mathbf{g}(\mu - \mathcal{E}) |\nabla \mathcal{E}|^2 \mathbf{dr} \approx \frac{V}{2} a^2 \left(q^2 \mathbf{g}(\mu) + \frac{1}{2} \mathbf{p}''(\mu) \right). \quad (\text{B5})$$

To investigate the left-hand side of (B2), we look at perturbations of the free energy in the limit of small a . We can express the derivatives of the free energy in a as equilibrium averages of the perturbed system. For $a = 0$,

$$\frac{\partial \Phi[\mathcal{E}]}{\partial a} = -\langle \rho_{\mathbf{q}} \rangle_{\mu} = 0, \quad (\text{B6})$$

$$\frac{\partial^2 \Phi[\mathcal{E}]}{\partial a^2} = -\langle |\rho_{\mathbf{q}}|^2 \rangle_{\mu}, \quad (\text{B7})$$

with $\rho_{\mathbf{q}} = \sum_{j=1}^n \sin(\mathbf{q} \cdot \mathbf{r}_j)$. The first average is zero by translational invariance. For the second average, use that $\langle |\rho_{\mathbf{q}}|^2 \rangle_{\mu} = \langle n \rangle_{\mu} S(\mu; q)/2$, where $S(\mu; q)$ is the structure factor of the fluid⁶² at chemical potential μ . So,

$$\frac{\partial^2 \Phi[\mathcal{E}]}{\partial a^2} = -\frac{\langle n \rangle_{\mu}}{2} S(\mu; q) = -\frac{3V}{\pi \sigma_S^3} \eta S(\mu; q). \quad (\text{B8})$$

Now differentiate (B2) twice with respect to a , yielding

$$\frac{3V}{\pi \sigma_S^3} \eta S(\mu; q) \approx V \left[q^2 \mathbf{g}(\mu) + \frac{1}{2} \mathbf{p}''(\mu) \right], \quad (\text{B9})$$

where the left-hand side used (B8) and the right-hand side was approximated with (B5) before differentiation.

Finally, differentiate with respect to q and send $|q| \rightarrow 0$: we can identify the second order term of the square-gradient approximation as

$$\mathbf{g}(\mu) = \frac{3\eta}{2\pi \sigma_S^3} \frac{\partial^2}{\partial q^2} S(\mu; q) \Big|_{q=0}, \quad (\text{B10})$$

which is (47).

2. Accuracy of the square-gradient approximation

The intermediate level that is constructed in Sec. IV relies on the approximation (46) for the free energy of a non-homogeneous hard-sphere fluid. This section discusses the accuracy of this approximation for a system that only contains small particles in an external potential.

We consider a grand-canonical ensemble of small hard-spheres ($\sigma_S = 1$) in a periodic box $V = [0, L]^3$ of length $L = 10\sigma_S$ without any big particles. We perturb this system by a one-dimensional cosine potential with m periods

$$\mathcal{E}_{m,\cos}(r) = 2 \left[\cos\left(\frac{2m\pi}{10} r\right) + 1 \right]. \quad (\text{B11})$$

We apply this potential to the first component r_1 of the small particle positions $\mathbf{r} = (r_1, r_2, r_3)$. The potential $\mathcal{E}_{m,\cos}$ has a fixed maximal strength of 4. We vary the steepness of the potential by varying the number of periods of the cosine $m = 1, 2, 3$; as m increases, the potential changes more rapidly. We expect this to make our approximations increasingly inaccurate, as it is constructed under the assumption that the derivatives of the external potential are small.

We have computed the predicted free energy (46) in this system, as well as the cruder approximation (45), which lacks the square gradient approximation. We compare these values with the true free energy, computed via thermodynamic integration, and investigate the dependence of the accuracy of the approximation on the small particle volume fraction η_S^c and the steepness of the external potential. The results of this computation are shown in Fig. 13.

Figure 13(a) shows that the absolute value of the free energy depends weakly on the number of periods m in the external potential. As expected, the error of the approximation methods increases substantially in Figs. 13(b)–13(d), as the number of periods m increases, as does the steepness of the cosine potential $\mathcal{E}_{m,\cos}$. The absolute prediction error also increases in the volume fraction η_S^c . For all m considered here, the square-gradient method (46) substantially outperforms the prediction (45). For $m = 1$, where the external potential varies the slowest, the square-gradient approximation is nearly exact, which confirms the accuracy of the square-gradient factor $\mathbf{g} = \mathbf{g}(\mu)$ in Appendix B 1. In addition to the increasing error on increasing m , the relative improvement of the square-gradient method over the simple approximation decreases. This indicates that higher order terms in the derivative of the potential start to become more important.

The choice of the 1d cosine potential $\mathcal{E}_{m,\cos}$ here is motivated by the use of a half-period of the cosine to introduce the suppression potential in the construction of the intermediate level in Sec. IV. For the numerical examples in Sec. V, the suppression potential has a (half-period) length of $l = 3.5$ and a maximal strength of $s = 4.4$. In terms of the maximal squared gradient that appears in this potential, it lies between the cases $m = 1$ and $m = 2$. This indicates that our square-gradient approximation is appropriate for the use of predicting the free energy of the partially inserted system in the binary hard-sphere example.

APPENDIX C: INFLUENCE OF THE AD HOC CORRECTION FACTOR

The construction of the intermediate level in Sec. IV included an *ad hoc* correction term that was identified using preliminary computations. In this section, we take another look at this parameter and discuss its importance for the performance of the three-level method. For this, we consider four values for the correction factor

$$\alpha_{\text{corr}} \in \{0, 0.04, 0.058, 0.076\} \quad (\text{C1})$$

between $\alpha_{\text{corr}} = 0$ (no *ad hoc* correction, which means $p_1 = p_1^{(2)}$), and $\alpha_{\text{corr}} = 0.076$ (corresponding to a linear least-square fit to the weights in Fig. 10). The effect of the choice of correction factor is displayed in Fig. 14 where we show the differences between the intermediate levels and the FG estimate for the two quantities of interest (see Fig. 3

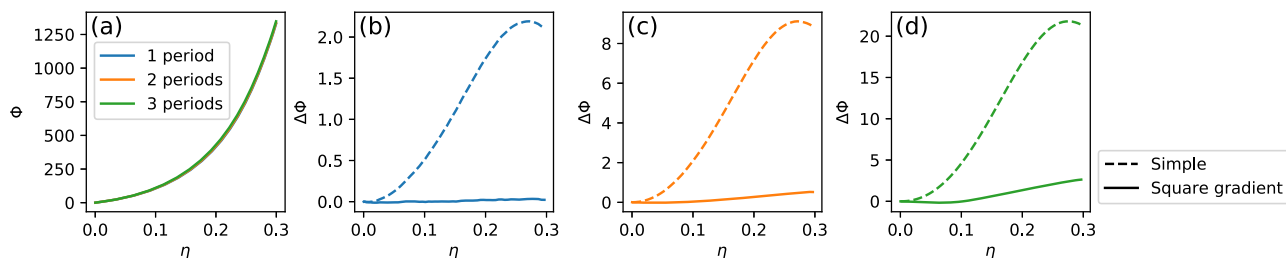


FIG. 13. A numerical test of the approximation accuracy of the square-gradient method for computing free energies of non-homogeneous hard-sphere fluids. (a) The free energy Φ of the hard-sphere fluid described in Appendix B 2 with a cosine potential with 1, 2, and 3 periods within the simulation box as a function of the volume fraction η . (b)–(d) The difference between the free energy and its pressure integral approximation in (46) for a cosine potential with one (b), two (c) or three (d) periods, with (solid line) and without (dashed lines) the square-gradient term.

for their FG averages). The *ad-hoc* correction has a noticeable influence on the quantities of interest, especially for the pair correlation function $g(r)$ when the two particles are almost touching ($r \approx 10$). This is in the relevant region for the coordination number N_c , which was measured as part of the performance test in Sec. V B.

To determine how much the three-level method depends on α_{corr} , we repeated this performance test for the different values in (C1). The measured sample variances for $N_{\text{runs}} = 60$ independent realizations of the estimators are shown in Fig. 15. As before, the results for different α_{corr} are highly correlated because they share the same configurations. The (bootstrap) standard errors of the values in Fig. 15 are comparable to the ones shown in Fig. 12; the same caveats apply and we have omitted them here for clarity of presentation.

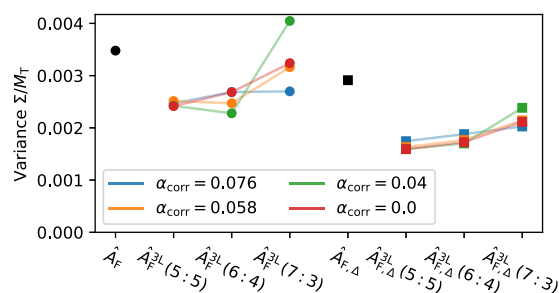


FIG. 15. The sample variance of 60 independent estimates of the coordination number N_c for different values of the *ad hoc* correction factor α_{corr} (see also Fig. 12).

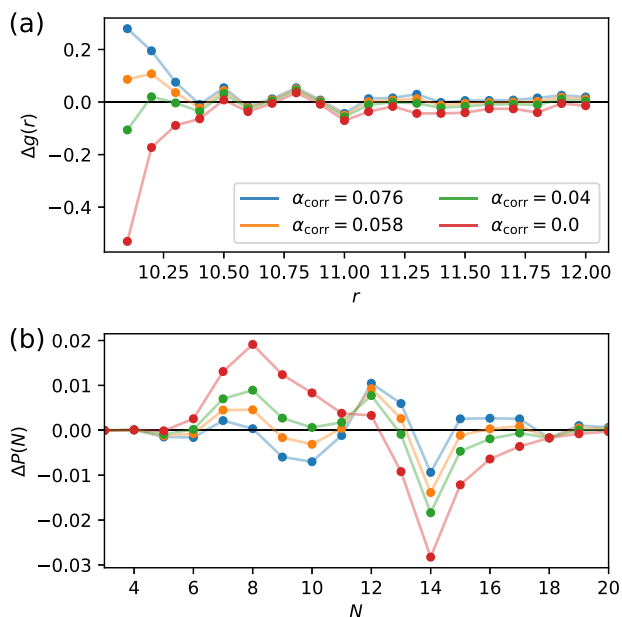


FIG. 14. The difference between the estimates at the intermediate and final level for (a) the pair correlation function $g(r)$ and (b) the distribution of the number of big particles for different values of the *ad-hoc* correction factor α_{corr} .

The main takeaway from Fig. 15 is that for our example, the three-level estimators without tapering outperform the corresponding two-level estimators, independent of the choice of α_{corr} . That is the method appears to be robust with respect to modifications of the intermediate level, even if the mean quantity of interest differs significantly between intermediate distributions.

Given the low number of samples, the exact variance figures should not be over-interpreted. Nevertheless, the trends in Fig. 15 illustrate two aspects of the SMC methodology that are relevant to applications. First, the details of the intermediate level become more important when we increase the tapering rate, as one can, for example, see by comparing the four different values for the final-level estimator \hat{A}_F^{3L} without tapering and with 7 : 3 tapering. Second, this effect is dampened for the difference estimators. The general robustness of the difference estimator in the example considered here supports our assertion that for appropriately defined levels, it should be the preferred estimator when applying the three-level method in practice.

REFERENCES

- C. N. Likos, “Effective interactions in soft condensed matter physics,” *Phys. Rep.* **348**, 267–439 (2001).
- W. G. Noid, J.-W. Chu, G. S. Ayton, V. Krishna, S. Izvekov, G. A. Voth, A. Das, and H. C. Andersen, “The multiscale coarse-graining method. I. A rigorous bridge between atomistic and coarse-grained models,” *J. Chem. Phys.* **128**, 244114 (2008).

- ³C. Peter and K. Kremer, "Multiscale simulation of soft matter systems—from the atomistic to the coarse-grained level and back," *Soft Matter* **5**, 4357–4366 (2009).
- ⁴M. Praprotnik, L. Delle Site, and K. Kremer, "A macromolecule in a solvent: Adaptive resolution molecular dynamics simulation," *J. Chem. Phys.* **126**, 134902 (2007).
- ⁵C. Pierleoni, B. Capone, and J.-P. Hansen, "A soft effective segment representation of semidilute polymer solutions," *J. Chem. Phys.* **127**, 171102 (2007).
- ⁶T. E. Ouldridge, A. A. Louis, and J. P. K. Doye, "Structural, mechanical, and thermodynamic properties of a coarse-grained DNA model," *J. Chem. Phys.* **134**, 085101 (2011).
- ⁷B. M. Mladek, J. Fornleitner, F. J. Martinez-Veracoechea, A. Dawid, and D. Frenkel, "Procedure to construct a multi-scale coarse-grained model of DNA-coated colloids from experimental data," *Soft Matter* **9**, 7342–7355 (2013).
- ⁸A. J. Pak and G. A. Voth, "Advances in coarse-grained modeling of macromolecular complexes," *Curr. Opin. Struct. Biol.* **52**, 119–126 (2018).
- ⁹S. Asakura and F. Oosawa, "On interaction between two bodies immersed in a solution of macromolecules," *J. Chem. Phys.* **22**, 1255–1256 (1954).
- ¹⁰R. Roth, R. Evans, and S. Dietrich, "Depletion potential in hard-sphere mixtures: Theory and applications," *Phys. Rev. E* **62**, 5360 (2000).
- ¹¹M. Karplus, "Development of multiscale models for complex chemical systems: From H + H₂ to biomolecules (Nobel lecture)," *Angew. Chem., Int. Ed.* **53**, 9992–10005 (2014).
- ¹²A. Warshel, "Multiscale modeling of biological functions: From enzymes to molecular machines (Nobel lecture)," *Angew. Chem., Int. Ed.* **53**, 10020–10031 (2014).
- ¹³H. N. W. Lekkerkerker, W. C.-K. Poon, P. N. Pusey, A. Stroobants, and P. B. Warren, "Phase behaviour of colloid + polymer mixtures," *Europhys. Lett.* **20**, 559–564 (1992).
- ¹⁴W. C. K. Poon, "The physics of a model colloid–polymer mixture," *J. Phys.: Condens. Matter* **14**, R859 (2002).
- ¹⁵T. Biben and J.-P. Hansen, "Phase separation of asymmetric binary hard-sphere fluids," *Phys. Rev. Lett.* **66**, 2215 (1991).
- ¹⁶M. Dijkstra, R. van Roij, and R. Evans, "Phase diagram of highly asymmetric binary hard-sphere mixtures," *Phys. Rev. E* **59**, 5744 (1999).
- ¹⁷H. Kobayashi, P. B. Rohrbach, R. Scheichl, N. B. Wilding, and R. L. Jack, "Critical point for demixing of binary hard spheres," *Phys. Rev. E* **104**, 044603 (2021).
- ¹⁸H. Kobayashi, P. B. Rohrbach, R. Scheichl, N. B. Wilding, and R. L. Jack, "Correction of coarse-graining errors by a two-level method: Application to the Asakura-Oosawa model," *J. Chem. Phys.* **151**, 144108 (2019).
- ¹⁹R. W. Zwanzig, "High-temperature equation of state by a perturbation method. I. Nonpolar gases," *J. Chem. Phys.* **22**, 1420–1426 (1954).
- ²⁰C. Jarzynski, "Nonequilibrium equality for free energy differences," *Phys. Rev. Lett.* **78**, 2690 (1997).
- ²¹G. E. Crooks, "Path-ensemble averages in systems driven far from equilibrium," *Phys. Rev. E* **61**, 2361 (2000).
- ²²R. M. Neal, "Annealed importance sampling," *Stat. Comput.* **11**, 125–139 (2001).
- ²³N. J. Gordon, D. J. Salmond, and A. F. M. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," in *IEE Proceedings F-Radar and Signal Processing* (IET, 1993), Vol. 140, pp. 107–113.
- ²⁴P. Del Moral, *Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications* (Springer, 2004), Vol. 88.
- ²⁵A. Doucet and A. M. Johansen, "A tutorial on particle filtering and smoothing: Fifteen years later," in *Handbook of Nonlinear Filtering*, edited by D. Crisan and B. Rozovsky (Oxford University Press, 2011), Vol. 12, p. 3.
- ²⁶Y. Iba, "Population Monte Carlo algorithms," *Trans. Jpn. Soc. Artif. Intell.* **16**, 279–286 (2001).
- ²⁷P. Grassberger, "Go with the winners: A general Monte Carlo strategy," *Comput. Phys. Commun.* **147**, 64–70 (2002).
- ²⁸H.-P. Hsu and P. Grassberger, "A review of Monte Carlo simulations of polymers with PERM," *J. Stat. Phys.* **144**, 597–637 (2011).
- ²⁹C. Giardinà, J. Kurchan, V. Lecomte, and J. Tailleur, "Simulating rare events in dynamical processes," *J. Stat. Phys.* **145**, 787–811 (2011).
- ³⁰P. J. Reynolds, D. M. Ceperley, B. J. Alder, and W. A. Lester, Jr, "Fixed-node quantum Monte Carlo for molecules," *J. Chem. Phys.* **77**, 5593–5603 (1982).
- ³¹M. B. Giles, "Multilevel Monte Carlo path simulation," *Oper. Res.* **56**, 607–617 (2008).
- ³²V. H. Hoang, C. Schwab, and A. M. Stuart, "Complexity analysis of accelerated MCMC methods for Bayesian inversion," *Inverse Probl.* **29**, 085010 (2013).
- ³³T. J. Dodwell, C. Ketelsen, R. Scheichl, and A. L. Teckentrup, "A hierarchical multilevel Markov chain Monte Carlo algorithm with applications to uncertainty quantification in subsurface flow," *SIAM/ASA J. Uncertainty Quantif.* **3**, 1075–1108 (2015).
- ³⁴A. Jasra, K. Kamatani, K. J. H. Law, and Y. Zhou, "Multilevel particle filters," *SIAM J. Numer. Anal.* **55**, 3068–3096 (2017).
- ³⁵A. Beskos, A. Jasra, K. Law, R. Tempone, and Y. Zhou, "Multilevel sequential Monte Carlo samplers," *Stochastic Processes Appl.* **127**, 1417–1440 (2017).
- ³⁶P. Del Moral, A. Jasra, and K. J. H. Law, "Multilevel sequential Monte Carlo: Mean square error bounds under verifiable conditions," *Stochastic Anal. Appl.* **35**, 478–498 (2017).
- ³⁷A. D. Bruce and N. B. Wilding, "Computational strategies for mapping equilibrium phase diagrams," *Adv. Chem. Phys.* **127**, 1 (2003).
- ³⁸J. Kolafa, S. Labík, and A. Malijevský, "Accurate equation of state of the hard sphere fluid in stable and metastable regions," *Phys. Chem. Chem. Phys.* **6**, 2335–2340 (2004).
- ³⁹D. J. Ashton, N. B. Wilding, R. Roth, and R. Evans, "Depletion potentials in highly size-asymmetric binary hard-sphere mixtures: Comparison of simulation results with theory," *Phys. Rev. E* **84**, 061136 (2011).
- ⁴⁰D. J. Ashton and N. B. Wilding, "Grand canonical simulation of phase behaviour in highly size-asymmetrical binary fluids," *Mol. Phys.* **109**, 999–1007 (2011).
- ⁴¹C. P. Robert and G. Casella, *Monte Carlo Statistical Methods* (Springer, 2004), Vol. 2.
- ⁴²R. Douc and O. Cappé, "Comparison of resampling schemes for particle filtering," in *ISPA 2005. Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis* (IEEE, 2005), pp. 64–69.
- ⁴³O. Cappé, E. Moulines, and T. Rydén, *Inference in Hidden Markov Models*, Springer Series in Statistics (Springer-Verlag, Berlin, Heidelberg, 2005).
- ⁴⁴P. Del Moral, A. Doucet, and A. Jasra, "Sequential Monte Carlo samplers," *J. R. Stat. Soc., Ser. B: Stat. Methodol.* **68**, 411–436 (2006).
- ⁴⁵E. L. Ionides, C. Bretó, and A. A. King, "Inference for nonlinear dynamical systems," *Proc. Natl. Acad. Sci. U. S. A.* **103**, 18438–18443 (2006).
- ⁴⁶P. Fearnhead, O. Papaspiliopoulos, and G. O. Roberts, "Particle filters for partially observed diffusions," *J. R. Stat. Soc., Ser. B: Stat. Methodol.* **70**, 755–777 (2008).
- ⁴⁷P. Fearnhead, O. Papaspiliopoulos, G. O. Roberts, and A. Stuart, "Random-weight particle filtering of continuous time processes," *J. R. Stat. Soc., Ser. B: Stat. Methodol.* **72**, 497–512 (2010).
- ⁴⁸C. Naesseth, F. Lindsten, and T. Schon, "Nested sequential Monte Carlo methods," in *International Conference on Machine Learning* (PMLR, 2015), pp. 1292–1301.
- ⁴⁹P. B. Rohrbach and R. L. Jack, "Convergence of random-weight sequential Monte Carlo methods," [arXiv:2208.12108](https://arxiv.org/abs/2208.12108) (2022).
- ⁵⁰H. Rafii-Tabar, L. Hua, and M. Cross, "A multi-scale atomistic-continuum modelling of crack propagation in a two-dimensional macroscopic plate," *J. Phys.: Condens. Matter* **10**, 2375 (1998).
- ⁵¹M. Praprotnik, L. Delle Site, and K. Kremer, "Adaptive resolution molecular-dynamics simulation: Changing the degrees of freedom on the fly," *J. Chem. Phys.* **123**, 224106 (2005).
- ⁵²R. Evans, "The nature of the liquid–vapour interface and other topics in the statistical mechanics of non-uniform, classical fluids," *Adv. Phys.* **28**, 143–200 (1979).
- ⁵³M. L. de Haro and M. Robles, "The structure factor and equation of state of hard-sphere fluids," *J. Phys.: Condens. Matter* **16**, S2089 (2004).
- ⁵⁴T. Hahn, "Cuba—A library for multidimensional numerical integration," *Comput. Phys. Commun.* **168**, 78–95 (2005).
- ⁵⁵B. Efron, "Nonparametric estimates of standard error: The jackknife, the bootstrap and other methods," *Biometrika* **68**, 589–599 (1981).

⁵⁶D. Williams, *Probability with Martingales* (Cambridge University Press, 1991).

⁵⁷R. Douc and E. Moulines, "Limit theorems for weighted samples with applications to sequential Monte Carlo methods," *Ann. Stat.* **36**, 2344–2376 (2008).

⁵⁸H. P. Chan and T. L. Lai, "A general theory of particle filters in hidden Markov models and some applications," *Ann. Stat.* **41**, 2877–2904 (2013).

⁵⁹G. D'Adamo, R. Menichetti, A. Pelissetto, and C. Pierleoni, "Coarse-graining polymer solutions: A critical appraisal of single- and multi-site models," *Eur. Phys. J.: Spec. Top.* **224**, 2239–2267 (2015).

⁶⁰D. Crisan and A. Doucet, "Convergence of sequential Monte Carlo methods," Signal Processing Group, Department of Engineering, University of Cambridge, Technical Report CUEDIF-INFENGR38, No. 1, 2000.

⁶¹H. Oberhofer and C. Dellago, "Efficient extraction of free energy profiles from nonequilibrium experiments," *J. Comput. Chem.* **30**, 1726–1736 (2009).

⁶²J.-P. Hansen and I. R. McDonald, *Theory of Simple Liquids: With Applications to Soft Matter* (Academic Press, 2013).