Grand canonical ensemble simulation studies of polydisperse fluids

Nigel B. Wilding
Department of Mathematical Sciences, The University of Liverpool, Liverpool L69 7ZL, U.K.

Peter Sollich
Department of Mathematics, King's College, University of London, Strand, London WC2R 2LS, UK

Abstract

We describe a Monte Carlo scheme for simulating polydisperse fluids within the grand canonical ensemble. Given some polydisperse attribute σ, the state of the system is described by a density distribution $\rho(\sigma)$ whose form is controlled by the imposed chemical potential distribution $\mu(\sigma)$. We detail how histogram extrapolation techniques can be employed to tune $\mu(\sigma)$ such as to traverse some particular desired path in the space of $\rho(\sigma)$. The method is applied in simulations of size-disperse hard spheres with densities distributed according to Schulz and log-normal forms. In each case, the equation of state is obtained along the dilution line, i.e. the path along which the scale of $\rho(\sigma)$ changes but not its shape. The results are compared with the moment-based expressions of Boublik et al (J. Chem. Phys. **54**, 1523 (1971)) and Salacuse and Stell (J. Chem. Phys. **77**, 3714 (1982)). It is found that for high degrees of polydispersity, both expressions fail to give a quantitatively accurate description of the equation of state when the overall volume fraction is large.

Typeset using REVTEX

I. INTRODUCTION

Statistical mechanics was originally formulated to describe the properties of systems of *identical* particles such as atoms or small molecules. However, many materials of industrial and commercial importance do not fit neatly into this framework. For example, the particles in a colloidal suspension are never precisely identical to one another, but have a range of radii (and possibly surface charges, shapes etc). This dependence of the particle properties on one or more continuous parameters is known as polydispersity.

To process a polydisperse colloidal material, one needs to know its phase behavior, i.e. the conditions of temperature and pressure under which a given structure is thermodynamically stable. The main obstacles to gaining this information arise from the effectively infinite number of particle species present in a polydisperse system. Labeling these by the continuous polydispersity attribute σ, the state of the system must be described by a density distribution $\rho(\sigma)$, rather than a finite number of density variables. The phase diagram is therefore infinite dimensional, a feature that poses serious problems to experiment and theory alike.

The chief difficulty faced in *experimental* studies of polydisperse systems is that the infinite dimensionality of the phase diagram precludes a complete mapping of the phase behavior. Instead one is forced to focus attention on particular low dimensional manifolds (slices) of the full diagram. Typically this involves determining the system properties along some desired trajectory through the space of $\rho(\sigma)$. Such a strategy is often pursued in experiments on colloidal suspensions [1,2], where the phase behavior is studied along a so-called *dilution line*. The experimental procedure for tracking this line involves adding a prescribed quantity of colloid of some known degree of polydispersity to a vessel of fixed volume V, the remaining volume being occupied by a solvent. The number distribution of colloidal particles $N(\sigma)$ determines the density distribution of the suspension, $\rho(\sigma) \equiv N(\sigma)/V$. Since in a given substance, the relative proportions of the number of particles of each σ are fixed, changing the amount of colloid added simply alters the *scale* of $\rho(\sigma)$, not its *shape*. Thus, by varying $N(\sigma)$ at fixed V (or vice versa) one traces out a locus in the phase diagram in which only the overall scale of $\rho(\sigma)$ changes.

As regards *theoretical* studies of phase behavior, these typically endeavor to calculate the system free energy as a function of a set of density variables. The difficulty in achieving this for a polydisperse system is that the free energy $f[\rho(\sigma)]$ is a functional of $\rho(\sigma)$, and therefore itself occupies an infinite dimensional space. This renders intractable the task of identifying phase boundaries. Recently, however, progress has stemmed from the observation that it is possible to approximate the full free energy by a so-called "moment free energy" containing the full ideal gas contribution plus an excess part that depends only on certain principal moments of the full excess free energy [3]. Doing so reduces the problem to a finite number of density variables and allows calculation of phase coexistence properties within a systematically refinable approximation scheme. Additionally the theory delivers (for the given free energy) exact results for the location of spinodals, critical points and the cloud and shadow curves. Use of this approach promises to enhance significantly our understanding of phase behavior in polydisperse systems.

In view of the approximations inherent in theoretical approaches, it is natural to consider deploying computer simulation to study the phase behavior of polydisperse colloidal fluids. Although simulations (like experiment) are restricted to studying limited regions of the phase diagram such as the dilution line, they have the advantage that can be used to investigate the *same* model systems as studied theoretically. Furthermore, they deliver (modulo finite-size effects) essentially exact results, providing invaluable benchmarks against which to test theoretical predictions. Sometimes too, the physical insight gleaned from simulations serves as the impetus for fresh theoretical advances.

One simulation approach for obtaining the thermodynamic properties of a polydisperse system is to simply mimic the experimental procedure. This can be achieved by employing a canonical ensemble (CE) formalism wherein a simulation box of fixed volume V is populated by a prescribed number of particles N, whose sizes are distributed according to the desired $N(\sigma)$. It practice, however, it transpires that the CE represents a far from optimal framework for simulating polydisperse fluids. The principal difficulty lies with the limited range of computationally accessible particle numbers which, in any simulation, is typically many orders of magnitude smaller than found in an experiment. The resulting finite-size effects are particularly pronounced in a CE simulation because the specific realization of the disorder $N(\sigma)$ is *fixed.* This suppresses large scale fluctuations in $\rho(\sigma)$ and potentially leads to sampling deficiencies [4]. Additionally, the CE suffers other drawbacks familiar from simulation studies of monodisperse fluids. For instance, relaxation times are extended because density fluctuations decay solely via diffusion; there is no direct access to information on chemical potentials; metastability and hysteresis hinder the study of phase transitions.

Experience with the simulation of monodisperse fluids has shown that use of the grand canonical ensemble (GCE) is highly effective in circumventing many of the aforementioned problems associated with the CE [5–7]. As we shall show, its application in the context of polydisperse fluids retains many of the benefits of the monodisperse case. Moreover, it provides the key to improved sampling of the density distribution $\rho(\sigma)$. This is because within the GCE framework $N(\sigma)$ fluctuates *as a whole*, thereby capturing fluctuations in $\rho(\sigma)$ on all simulation length scales. Notwithstanding these advantages, however, the GCE might appear, at first sight, unsuitable for the purpose of traversing a *particular* trajectory through the space of the density distribution $\rho(\sigma)$. This is because $\rho(\sigma)$ ostensibly lies outwith the direct control of the simulator, its form being instead determined by the imposed chemical potential distribution $\mu(\sigma)$. Nevertheless, it turns out to be possible to *tune* $\mu(\sigma)$ within a histogram extrapolation scheme, in such a way as to realize a specific desired form of $\rho(\sigma)$. We shall demonstrate that this dual use of the GCE and histogram extrapolation methods permits a chosen phase space path to be followed efficiently.

The layout of our paper is as follows. In sec. II A we formulate the statistical mechanics for a polydisperse fluid within the grand canonical ensemble. We then describe (sec. II B) the combined GCE plus histogram extrapolation methodology for tracking a particular path through the space of $\rho(\sigma)$. In section III we apply the method to the problem of obtaining the dilution line properties of three types of size-disperse hard sphere fluids. The chemical potential distribution of these fluids is determined as a function of volume fraction and the results compared with the predictions of two commonly used equations of state. Finally in section IV, we discuss our findings and their implications.

II. METHOD

A. Statistical Mechanics

We consider a classical fluid of polydisperse particles confined to a volume $V = L^d$. The system is assumed to be thermodynamically open, so that the particle-number distribution $N(\sigma)$ is a statistical quantity. The associated grand canonical partition function takes the form:

$$\mathcal{Z}_V = \sum_{N=0}^{\infty} \frac{1}{N!} \prod_{i=1}^{N} \left\{ \int_V d\vec{r}_i \int_0^{\infty} d\sigma_i \right\} \exp\left(-\beta \mathcal{H}_N\left(\{\vec{r},\sigma\}\right)\right) \tag{2.1}$$

with

$$\mathcal{H}_N\left(\{\vec{r},\sigma\}\right) = \Phi\left(\{\vec{r},\sigma\}\right) - \sum_{i=1}^{N} \mu(\sigma_i) \; . \tag{2.2}$$

Here N is the overall particle number, while $\beta = (k_B T)^{-1}$ and $\mu(\sigma)$ are respectively the prescribed inverse temperature and chemical potential distribution. $\{\vec{r},\sigma\}$ denotes the *configuration*, i.e. the complete set $(\vec{r_1},\sigma_1),(\vec{r_2},\sigma_2)\cdots(\vec{r_N},\sigma_N)$ of particle position vectors and polydisperse attributes. The corresponding configurational energy $\Phi\left(\{\vec{r},\sigma\}\right)$ is assumed to reside in a sum of pairwise interactions

$$\Phi\left(\{\vec{r},\sigma\}\right) = \sum_{i<j=1}^{N} \phi(\mid \vec{r}_i - \vec{r}_j \mid, \sigma_i, \sigma_j) \; , \tag{2.3}$$

where ϕ is the pair potential.

The instantaneous number distribution is defined by

$$N(\sigma) \equiv \sum_{i=1}^{N} \delta(\sigma - \sigma_i) \; , \tag{2.4}$$

with σ the continuous polydispersity attribute and $N = \int N(\sigma) d\sigma$. We shall be concerned with the fluctuations in the associated density distribution

$$\rho(\sigma) \equiv N(\sigma)/V \; , \tag{2.5}$$

and the configurational energy density

$$u \equiv \Phi(\{\vec{r},\sigma\})/V \; .$$

The statistical behavior of these observables is completely described by their joint probability distribution [8]

$$p_V[\rho(\sigma), u] = \left\langle \delta\left(u - V^{-1}\Phi(\{\vec{r},\sigma\})\right) \prod_{\sigma} \delta\left(\rho(\sigma) - V^{-1}N(\sigma)\right) \right\rangle \; , \tag{2.6}$$

or more explicitly

$$\begin{aligned} p_V[\rho(\sigma), u] = \frac{1}{\mathcal{Z}_V} \sum_{N=0}^{\infty} \frac{1}{N!} \prod_{i=1}^{N} \left\{ \int_V d\vec{r}_i \int_0^{\infty} d\sigma_i \right\} \exp\left(-\beta \mathcal{H}_N\left(\{\vec{r},\sigma\}\right)\right) \\ \times\, \delta\left(u - V^{-1}\Phi(\{\vec{r},\sigma\})\right) \prod_{\sigma} \delta\left(\rho(\sigma) - V^{-1}N(\sigma)\right) \; . \end{aligned} \tag{2.7}$$

Integrating over the energy fluctuations yields the probability distribution function of the density distribution

$$p_V[\rho(\sigma)] = \int p_V[\rho(\sigma), u]du \; . \tag{2.8}$$

Our specific concern is with the average form of $\rho(\sigma)$, given by

$$\overline{\rho}(\sigma) = \int \rho(\sigma) p_V[\rho(\sigma)] d\rho(\sigma) \; . \tag{2.9}$$

Given a prescribed chemical potential distribution $\mu(\sigma)$ and temperature β, the form of $\overline{\rho}(\sigma)$ can be determined by simulation. Except in the ideal gas limit, however, no exact relationship between $\overline{\rho}(\sigma)$ and $\mu(\sigma)$ will generally be available. Thus one cannot (from a 'bare' GCE simulation), readily determine that $\mu(\sigma)$ corresponding to a particular desired *target* density distribution $\rho_t(\sigma)$. Subject to certain restrictions however, this can be achieved via use of histogram extrapolation.

The key idea of histogram extrapolation [9] is that a measured distribution $p_V[\rho(\sigma), u]$ accumulated at one set of model parameters $\mu(\sigma), \beta$ can be reweighted to yield estimates of the distribution appropriate to other parameters $\mu'(\sigma), \beta'$. In its simplest form the reweighting is given by

$$p'_V[\rho(\sigma), u|\mu'(\sigma), \beta'] = w p_V[\rho(\sigma), u|\mu(\sigma), \beta] \tag{2.10}$$

where the reweighting factor w takes the form

$$w = \exp\left(\sum_{i=1}^{N} [\beta'\mu'(\sigma_i) - \beta\mu(\sigma_i)] - V(\beta' - \beta) u\right) \tag{2.11}$$

By tuning the form of $\mu'(\sigma)$ and the value of β' within the reweighting scheme, it is possible to scan the space of $\overline{\rho}(\sigma)$, thereby "homing in" on the target density distribution. To this end it is expedient to define a *cost function* measuring the deviation of $\overline{\rho}(\sigma \mid \mu'(\sigma), \beta')$ from the target form

$$\Delta(\mu'(\sigma), \beta') \equiv \int \mid \overline{\rho}(\sigma) - \rho_t(\sigma) \mid \gamma(\sigma) d\sigma \; . \tag{2.12}$$

Here, for numerical convenience, we have incorporated a weight function $\gamma(\sigma)$ into our definition, the role of which (as described below) is to ensure that comparable contributions to the cost function arise from all sampled regions of the σ-domain. Within this framework, the task of determining those values of $\mu'(\sigma), \beta'$ that yield the target distribution $\rho_t(\sigma)$ reduces to that of functionally minimizing the cost function Δ with respect to $\mu'(\sigma), \beta'$. As we describe in the next section, this is achievable using standard algorithms.

B. Implementation

We have employed Monte Carlo (MC) simulation within the grand canonical ensemble to study the dilution line properties of systems of size-disperse hard spheres. This section details the principal aspects of our simulation and analysis procedure.

1. Program, data structure and acquisition

Our simulated system comprises a variable number of hard spheres contained within a periodic box of volume $V = L^d$. The dimensionless polydisperse variable σ was taken to be the particle *diameter* expressed in units of the mean diameter (see section III). An upper bound σ_c was placed on the permitted range of diameters, and the simulation volume was partitioned into an array of l^3 cells each of linear dimension σ_c, so that $L = l\sigma_c$. This strategy aids efficient identification of particle interactions by ensuring that interacting particles occupy either the same cell or directly neighboring ones.

The grand canonical ensemble Monte Carlo (GCMC) algorithm employed has a Metropolis form [10] and invokes four types of operation: particle displacements, particle insertions, particle deletions and particle resizing; each is attempted with equal frequency. Specific to the polydisperse case is the resizing operation which entails attempting to change the diameter of a nominated particle by an amount drawn from a uniform random deviate constrained to lie in some prescribed range. This range (maximum diameter step-size) is chosen to provide a suitable balance between efficient sampling and a satisfactory acceptance rate at the prevailing number density. As regards the remaining types of moves, these proceed in a manner similar to the monodisperse case [10], except that for insertion attempts the new particle diameter is drawn with uniform probability from the range $\sigma \in [0, \sigma_c]$.

As primary *input*, the program takes the chemical potential distribution $\mu(\sigma)$, which is required for the accept/reject Monte Carlo lottery. This distribution is stored in the form of a histogram, constructed by dividing the truncated interval $0 < \sigma < \sigma_c$ into a prescribed number M of sub-intervals or "bins" [11]. Thus all particles whose diameters fall within the scope of a given bin are associated with the same value of the chemical potential.

The principal *observable* of interest is the joint probability distribution $p_V[\rho(\sigma), u]$ (c.f. Eq. 2.8). Operationally it is impractical to construct the full form of this distribution in a simulation because to do so would entail forming a histogram over histograms–the memory storage costs of which would be prohibitive. The procedure adopted, therefore, was to sample the *instantaneous* density histogram $\rho(\sigma)$ (discretized in the same manner as $\mu(\sigma)$), and append successive measurements of this quantity to a file [12]. The set of all such samples constitutes a *list* representing a sequential history of the individual data measurements [7]. This data list is post-processed by an analysis program which reads in each of the individual list entries and averages over all of them to construct a histogram approximation to the average density distribution $\bar{\rho}(\sigma)$. If desired, the analysis program additionally implements histogram reweighting of the data in order to enable extrapolation to neighboring values of the model parameters. This extrapolation is achieved by assigning a weight to each list entry of the form given by equation 2.11. The complete set of weights permits construction of the reweighted histogram.

2. Tracking a phase space path

The strategy we have adopted is to traverse the phase space path of interest in a stepwise fashion, utilizing histogram extrapolation to proceed from one step to the next. In the general case of particles having a finite potential, the phase space path may involve changes in the temperature as well the form of $\rho(\sigma)$. For simplicity of illustration, however, let us presume that the path is isothermal (i.e β = constant) and that the form of the chemical potential distribution, $\mu^{(0)}(\sigma)$ say, is known at some arbitrary point $\rho^{(0)}(\sigma)$ along the path. The procedure is then as follows. From simulations at the known state point, data for $p_V[\rho(\sigma), u]$ can be accumulated directly. Histogram reweighting is then applied to this data to extrapolate some distance along the path to a new point $\rho^{(1)}(\sigma)$, *and* to provide an estimate of the corresponding form of $\mu^{(1)}(\sigma)$. The latter quantity is then employed in a fresh simulation, the results of which are extrapolated to a point further along the path, and so on. By iterating this procedure $\rho^{(i)}(\sigma) \rightarrow \rho^{(i+1)}(\sigma)$, $\mu^{(i)}(\sigma) \rightarrow \mu^{(i+1)}(\sigma)$, one traces out the entire phase space path.

The implementation of the extrapolation stage necessitates a prior choice for the step size, that is the difference between the measured $\rho(\sigma)$ and the next target $\rho_t(\sigma)$. The magnitude of this difference should be chosen to be as large as possible, consistent with remaining within the range of reliable extrapolation. A good indicator that this is in fact the case is that the individual densities of the target distribution $\rho_t(\sigma)$ each overlap with the range of typical fluctuations appearing in the simulation data $\rho(\sigma)$ [13].

Once a suitable step size has been determined, the extrapolation procedure proceeds by minimizing (within the reweighting scheme) the cost function Δ introduced in eq. 2.12. For all but the lowest densities, this task is complicated by the existence of strong coupling between the μ variables, deriving from the fact that the number density for each σ depends on the *whole* chemical potential distribution. Fortunately, efficient algorithms for performing multi-dimensional functional minimization are widely available [14] and, at least for the cases we have considered, appear to operate effectively. The sole difficulty encountered was that, on occasion, the minimization failed to fully converge for σ values in the wings of $\rho_t(\sigma)$. To remedy this problem, a weight function $\gamma(\sigma)$ was incorporated in the cost function (c.f. eq. 2.12), the purpose of which is to enhance the contribution to Δ from σ values for which $\rho_t(\sigma)$ is small. Good results were obtained by setting $\gamma(\sigma) \propto [\rho_t(\sigma)]^{-1}$ [15].

It should be stressed that the methodology set out above presumes the availability of a form of $\mu(\sigma)$ for some starting point on the phase space path of interest. This may be obtained straightforwardly if the path passes through a region of low density where reliable analytical estimates for $\mu(\sigma)$ can be employed. Otherwise the method must be bootstrapped by other means. One simple but effective approach for achieving this operates as follows. Starting from some initial guess for the desired $\mu(\sigma)$ (e.g. $\mu^0(\sigma) = \ln[\rho_t(\sigma)]$) a series of short simulations are carried out in which $\mu(\sigma)$ is iterated according to

$$\mu^{(m+1)}(\sigma) = \mu^{(m)}(\sigma) + \left(\frac{\ln(\rho_t(\sigma))}{\ln(\rho^{(m)}(\sigma))}\right)^{\delta}, \tag{2.13}$$

where $0 < \delta < 1$ is a damping factor, the value of which may be tuned to optimize convergence. Although one could certainly envisage more sophisticated and efficient schemes, we have found this method to operate satisfactorily.

III. DILUTION LINE STUDIES OF POLYDISPERSE HARD SPHERE FLUIDS

A. System and simulation details

We have obtained the dilution line properties (cf. sec. I) of size-disperse hard sphere [16] fluids with diameter distribution $N(\sigma)$ assigned one of two forms:

(i) Schulz distribution.

(ii) Log-normal distribution.

These distributions are conveniently expressed in terms of a normalized size function $n(\sigma) = N(\sigma)/N$. For the Schulz, this takes the form

$$n(\sigma) = \frac{1}{z!}\left(\frac{z+1}{\overline{\sigma}}\right)^{z+1} \sigma^z \exp\left[-\left(\frac{z+1}{\overline{\sigma}}\right)\sigma\right] , \tag{3.1}$$

where $\overline{\sigma}$ is the average particle diameter and z is a parameter controlling the width of the distribution. For the log-normal distribution, the size function is given by

$$n(\sigma) = \frac{1+W^2}{\overline{\sigma}\sqrt{2\pi\ln(1+W^2)}} \exp\left(-\frac{[\ln(\sigma/\overline{\sigma}) + (3/2)\ln(1+W^2)]^2}{2\ln(1+W^2)}\right) , \tag{3.2}$$

with W the standard deviation in units of $\overline{\sigma}$. Note that both of these distribution are normalized, that is $\int_0^\infty n(\sigma)d\sigma = 1$, and vanish as $\sigma \to 0$, implying a natural lower limit to σ. By contrast, there is no finite upper limit and consequently, for simulation purposes, it was necessary to impose an upper bound (cut off) σ_c (see also sec. II B).

We have studied *three* distinct size distributions—two of the Schulz form and one of the log-normal form. For the Schulz distribution, width-parameter values of $z = 15$ and $z = 5$ were considered, with cutoff values $\sigma_c = 3$ and $\sigma_c = 4$ respectively. For the log-normal distribution, the single case $W = 2.5$ with $\sigma_c = 12$ was studied. In each instance we set $\overline{\sigma} = 1$. Cell array sizes (cf. sec. II B 1) of linear dimension $l = 3, 4, 5$ were used. In absolute dimensionless units ($L = l\sigma_c$) these correspond to $L = 9, 12, 15$ respectively for the $z = 15$ Schulz, to $L = 12, 16, 20$ for the $z = 5$ Schulz, and to $L = 36, 48, 60$ for the log-normal distribution. The histogram discretization parameter (cf. sec. II B) was set to $M = 75$ and $M = 100$ for the $z = 15$ and $z = 5$ Schulz distributions respectively, and to $M = 120$ for the log-normal distribution.

The average density distribution can be expressed in terms of the normalized size function multiplied by the overall number density $\rho_0 = N/V$, i.e.

$$\overline{\rho}(\sigma) \equiv \rho_0 n(\sigma) . \tag{3.3}$$

The procedure for tracking the dilution line utilized this overall density as a measure of the location along the line. The tracking procedure was initiated at a small value of ρ_0 by approximating the chemical potential distribution according to the ideal gas relation $\mu(\sigma) = \ln\overline{\rho}(\sigma)$ [17]. Histogram extrapolation was applied to the resulting simulation data in order to refine the initial estimate of $\mu(\sigma)$. Thereafter the strategy described in sec. II B 2 was implemented to follow the dilution line to higher densities. At each step this entailed setting a target form for $\overline{\rho}(\sigma)$ corresponding to a value of ρ_0 larger than that used at the previous step. The cost function measuring the discrepancy between $\overline{\rho}(\sigma)$ and the target was then minimized to yield an estimate for the appropriate $\mu(\sigma)$. On efficiency grounds, this minimization was performed in two stages; an initial approximation to $\mu(\sigma)$ was obtained from a one-dimensional minimization in which the activity distribution $\exp(\mu(\sigma))$ was multiplied by an overall factor. Thereafter the complete functional minimization was performed to yield a more accurate form for $\mu(\sigma)$.

Although (within the specific context of the tracking procedure), ρ_0 provides a convenient measure of the location on the dilution line, it conveys little information regarding the degree of packing within a polydisperse system. We shall therefore find it convenient to quote values for the overall volume fraction of the system given by

$$\eta \equiv \int_0^{\sigma_c} d\sigma \frac{\pi}{6} \sigma^3 \bar{\rho}(\sigma)\,. \tag{3.4}$$

It is this quantity (rather than ρ_0) which features in the presentation of our results, to which we now turn.

B. Results

Owing to the computational complexity of the simulations, we have obtained the complete dilution line for each fluid only for the $l = 3$ cell array size. For the larger system sizes, only a few spot measurements were made along the dilution line, each of which was bootstrapped by appeal to the measured $\mu(\sigma)$ obtained from the $l = 3$ systems.

The dilution lines for the three fluids were tracked to the highest computationally accessible volume fraction, terminating only once the relaxation timescale became excessively large. The maximum volume fractions attained were $\eta = 0.445$ and $\eta = 0.40$ for the Schulz density distributions with $z = 15$ and $z = 5$ respectively, and $\eta = 0.33$ for the log-normal case. Snapshot configurations for all three fluids are shown in fig. 1 for η values slightly below the maximum attained in each case. We mention in passing that, at least for the cases of the two Schulz distributions, the maximum volume fractions reached are somewhat larger than would be readily attainable for GCMC simulations of monodisperse hard spheres. The latter, of course, become highly inefficient at large η because the acceptance rate for particle insertions falls rapidly as the free volume diminishes. Whilst the same is true in the polydisperse context for insertions of large particles, space can often be found for placing a small particle. This facilitates fluctuations in the overall particle number N, while resizing operations (whose efficiency at large η is typically greater than that of inserting large particles) ensure continued proper sampling of the density distribution.

The measured forms for $\mu(\sigma)$ as a function of η differ qualitatively between the Schulz and the log-normal density distributions, and accordingly we discuss them separately. Beginning with the Schulz case, figs. 2 and 3 show the measured $\overline{\rho}(\sigma)$ and the corresponding $\mu(\sigma)$ for the two fluids at a selection of volume fractions along their respective dilution lines. The range of σ values shown is that for which the simulations delivered data of reasonable statistical quality. One notes that for both the $z = 15$ and $z = 5$ cases, the tails marking the large–σ vestiges of the distributions fall considerably short of the respective cutoffs σ_c. Indeed, we have verified that in the course of the simulations, no particles of diameters approaching the cutoff diameter occurred on the simulation timescale, implying that our data is unaffected by its imposition. As regards the forms of the chemical potential distributions, one sees that for small η they display a maximum near the peak in $\overline{\rho}(\sigma)$—behavior that is of course mandated at sufficiently low η by the known properties of the ideal gas limit. For larger η, however, the peak is lost and $\mu(\sigma)$ increases monotonically. The increase of $\mu(\sigma)$ in the regime of large σ indicates that the excess chemical potential (measuring the work associated with inserting a sphere) grows faster with σ than the decrease in the ideal contribution associated with the decay of $\overline{\rho}(\sigma)$.

Turning now to the log-normal case, fig. 4(a) demonstrates that $\overline{\rho}(\sigma)$ decays extremely slowly with increasing σ. The peak in the distribution therefore occurs at much smaller σ than for the two Schulz forms although the average diameter $\overline{\sigma}$ is identical in all three cases. As a practical consequence of the slow decay (and notwithstanding the imposition of a very large value of σ_c), ranges of particle diameters extending up to the cutoff value were observed in the system. Indeed it was not feasible within the computational constraints to utilize a cutoff for which this did not occur, and hence truncation effects are always significant in this system—a point to which we shall return below. The measured forms of $\mu(\sigma)$ for the log-normal fluid are shown in fig. 4(b). In contrast to the Schulz distributions, they display (for all accessible η) a narrow maximum close to the peak in $\overline{\rho}(\sigma)$ at small ρ. Thereafter with increasing σ there is a slow fall to a broad minimum, whereafter $\mu(\sigma)$ increases strongly.

Having outlined the main qualitative features of the relationship between $\overline{\rho}(\sigma)$ and $\mu(\sigma)$, it is instructive to perform a detailed comparison between our measurements and the predictions of analytical equations of state (EOS) appearing in the literature. For hard spheres, two commonly used equations are that due to Boublik, Mansoori, Carnahan, Starling and Leland (BMCSL) [18,19] based on the Carnahan-Starling equation for monodisperse hard spheres, and that due to Salacuse and Stell [20], based on the Percus-Yevick (PY) theory. Both are reproduced in appendix A, and express $\mu(\sigma)$ in terms of an expansion to third order in σ, with coefficients given in terms of the first three moments of $\overline{\rho}(\sigma)$. We have compared the predictions of the EOS for each of the three fluids studied, with the finite-size simulation data at three values of η, namely a low, a moderate and a high value. We describe our findings for each fluid in turn.

The results for the Schulz distribution with $z = 15$ at the low volume fraction $\eta = 0.056$ are shown in fig. 5(a). At first sight there is good agreement between the $L = 9$ and $L = 12$ simulation data and both the BMCSL and PY equations of state over the entire region of σ. However, closer inspection reveals appreciable discrepancies between theory and simulation, not visible on the scale of $\mu(\sigma)$. These are apparent if one suppresses the dominant ideal gas contribution to expose the *excess* chemical potential, given by (see appendix A)

$$\mu_{\rm ex}(\sigma) = \mu(\sigma) - \ln[\rho_0 n(\sigma)]\,. \tag{3.5}$$

This quantity is plotted in the inset of fig. 5(a), from which one sees that compared to the simulation results, the EOS slightly underestimate $\mu(\sigma)$.

Fig. 5(b) shows the results for the $z = 15$ Schulz distribution at the moderate volume fraction $\eta = 0.257$. Again there is good agreement between the $L = 9$ and $L = 12$ simulation data suggesting that finite-size effects are insignificant. Here, however, discrepancies between the BMCSL and PY equations of state are larger than at the lower value of η, being evident on the scale of the *absolute* chemical potential. One sees that both EOS significantly underestimate $\mu(\sigma)$ for large σ, although they agree quite well with one another. A similarly high level of agreement between the data from the $L = 9$ and $L = 12$ system sizes is manifest at the higher volume fraction $\eta = 0.426$, fig. 5(c). Here the PY equation of state is seen to fare somewhat better than the BMCSL equation although both underestimate $\mu(\sigma)$ substantially towards the upper end of the σ range.

A similar picture emerges for the Schulz distribution with $z = 5$ (fig. 6). Again both equations of state underestimate $\mu(\sigma)$, even at the lowest volume fraction, although the PY equation corresponds significantly more closely to the simulation results at high volume fractions than does the BMCSL equation. We could again discern no evidence for appreciable finite-size effects in the simulation results.

In seeking to compare the results for the log-normal system with the EOS predictions it is essential to bear in mind the importance of truncation effects. The moments of a truncated log-normal distribution can differ dramatically from those of the full distribution even for large values of σ_c. In order to facilitate a fair comparison with theory, the analytic form of $\mu(\sigma)$ must therefore be calculated using the moments of the *same* truncated distribution as employed in the simulations. The results of performing this comparison are presented in fig. 7. At low volume fraction, there is good agreement between the EOS predictions and the simulation results, while at high volume fraction the EOS underestimate the measured $\mu(\sigma)$ considerably, with the degree of discrepancy increasing towards the tail of the distribution. Once again we could discern no evidence of finite-size effects within the statistical uncertainties of our data.

The results presented above indicate that the BMCSL and PY equations fail to provide a quantitatively accurate description of the chemical potential distribution, particularly when the volume fraction is large. It is instructive to examine the implications of this finding for calculations of *densities*, which depends very sensitively (indeed exponentially so) on the chemical potential. To this end we have studied the degree to which the form of $\mu(\sigma)$ calculated via the BMCSL equation from some prescribed $\rho(\sigma)$, actually yields this density distribution when input to a simulation. The results of performing this comparison are shown in figures 8(a) – 8(c) for each of the three fluids at a high volume fraction. In each instance, the solid line shows the input density distribution $\rho(\sigma)$ from which $\mu(\sigma)$ is calculated. The data points show the simulation results obtained using this form of $\mu(\sigma)$ for 3 different system sizes. As the figures clearly demonstrate, the measured form of $\rho(\sigma)$ deviate substantially from the prediction.

Finally in this section, we examine the moment structure of the excess chemical potential $\mu_{\rm ex}(\sigma)$ given by eq. 3.5. Both the BMCSL and PY equations assume that $\mu_{\rm ex}(\sigma)$ is expressible in terms of a cubic polynomial in σ. Bartlett has reached the same conclusion using geometrical arguments inspired by scaled particle theory [21]. We have investigated this proposal by fitting our data to the expression

$$\mu_{\rm ex}(\sigma) = -\ln(1-\eta) + \alpha_1\sigma + \alpha_2\sigma^2 + \alpha_3\sigma^3 \,, \tag{3.6}$$

where the constant term is fixed by the requirement that in the limit $\sigma \to 0$ the probability of inserting a sphere is proportional to $1 - \eta$. We find that all our data are fitted very well by this expression; fig. 9(a) shows a typical fit for the case of the Schulz fluid with $z = 5$ at $\eta = 0.377$. Also shown, (fig. 9(b)) are the fit coefficients $\alpha_1, \alpha_2, \alpha_3$ for various values of η, together with the predictions of the BMCSL equation of state. One sees that at high volume fraction, the BMCSL underestimates all coefficients, the relative discrepancy being marginally larger for the α_1 coefficient than for the others.

IV. DISCUSSION AND CONCLUSIONS

In summary, we have presented a grand canonical simulation method for studying polydisperse fluids. The method utilizes histogram extrapolation techniques to track efficiently an arbitrary path of interest through the space of the density distribution $\rho(\sigma)$. We have applied it to the specific problem of obtaining the dilution line properties of size-disperse hard spheres. It should also prove useful in studying more general species of polydisperse fluids and their phase transitions, both in the bulk and confined geometries. We intend to report on such extensions in future communications.

Previous simulation studies of polydisperse fluids have generally operated within a semi-grand canonical framework [22] in which a fixed number of particles are studied either at constant pressure (see eg. ref. [23]) or within a Gibbs Ensemble MC scheme [24]. In common with the present work, these studies utilized a fluctuating particle size distribution, realized by means of MC resizing moves controlled by a chemical potential distribution. In contrast to our approach, however, thermodynamic properties were studied as a function of the shape of the activity distribution $\exp(\mu(\sigma))$; no constraints were placed on the conjugate density distribution, which was consequently free to adopt whichever functional form minimized the free energy for the imposed $\mu(\sigma)$. In view of this, one might question the extent to which the simulation results reflect the actual situation in realistic systems. One situation in which the lack of a constrained density distribution might be relevant is the interesting issue of the influence of polydispersity on the freezing of hard spheres [25]. This was investigated in ref. [23], using semi-grand canonical MC and Gibbs-Duhem integration. It would be of considerable interest to harness the approach described in the present work to investigate the effects on the freezing behaviour of varying the width of the density distribution whilst constraining its shape to some physically realistic form.

In the present study, attention was focused on fluids having a high degree of polydispersity. The motivation for this choice was two-fold. First, models exhibiting a wide density distribution provide a suitably testing challenge against which to assess the effectiveness of our method. Second, there exist in the literature a number of interesting predictions concerning the role of depletion forces in highly polydisperse systems. For instance, it has been suggested by several authors that attractive depletion forces might engender novel phase transitions in polydisperse hard spheres [26–28]. Specifically, Sear has suggested that a fluid of hard spheres having a log-normal size distribution will be unstable with respect to crystallization of the large particles at *all* finite volume fractions [28]. By contrast, Cuesta [27] has predicted that a sufficiently wide log-normal distribution will exhibit fluid-fluid phase separation at some finite density. It seems likely that for a truncated size distribution a phase transition will not occur for arbitrarily small volume fraction because of the absence of the largest particles which mediate the greatest depletion forces. Thus our simulation results for the truncated distribution are able neither to conclusively confirm nor refute these predictions. It suffices to say that we observed no evidence of crystallization in the particular truncated log-normal fluid studied, up to a volume fraction of $\eta = 0.33$. Similarly, within the range of accessible volume fractions, no evidence for phase transitions was observed in either of the two Schulz fluids studied.

Turning finally to the comparison between our simulation results and the predictions of the polydisperse equations of state, we find that neither the BMCSL nor PY equations offer a quantitatively accurate description of the thermodynamics of hard spheres for large polydispersity and at high volume fraction [29]. Both equations underestimate $\mu(\sigma)$ at all fluid densities over the entire range of σ, implying that they underestimate the Helmholtz free energy density $f = \int_0^{\rho_0} \mu(\sigma|\rho_0' n(\sigma)) d\sigma d\rho_0'$ and hence overestimate the stability of the fluid. The magnitude of this overestimate becomes more pronounced the greater the volume fraction. Interestingly we find that in this regime the PY equation performs appreciably better than the BMCSL equation despite the fact (cf. appendix A) that the latter derives from a monodisperse hard sphere equation of state which has been found to be superior to the PY approximation. It remains to be seen to what extent the overestimate of fluid stability impinges on the results of existing calculations of depletion-force induced phase separation based on the BMCSL and PY approximations [26,27]. In any case, our results should provide a useful testing ground for any future improvements to the existing polydisperse equations of state [30].

ACKNOWLEDGMENTS

NBW acknowledges helpful discussions with Marcus Müller. This work was supported by The University of Liverpool Research Development Fund and the UK CCP5 group. PS acknowledges financial support through EPSRC grant GR/R52121.

FIGURES

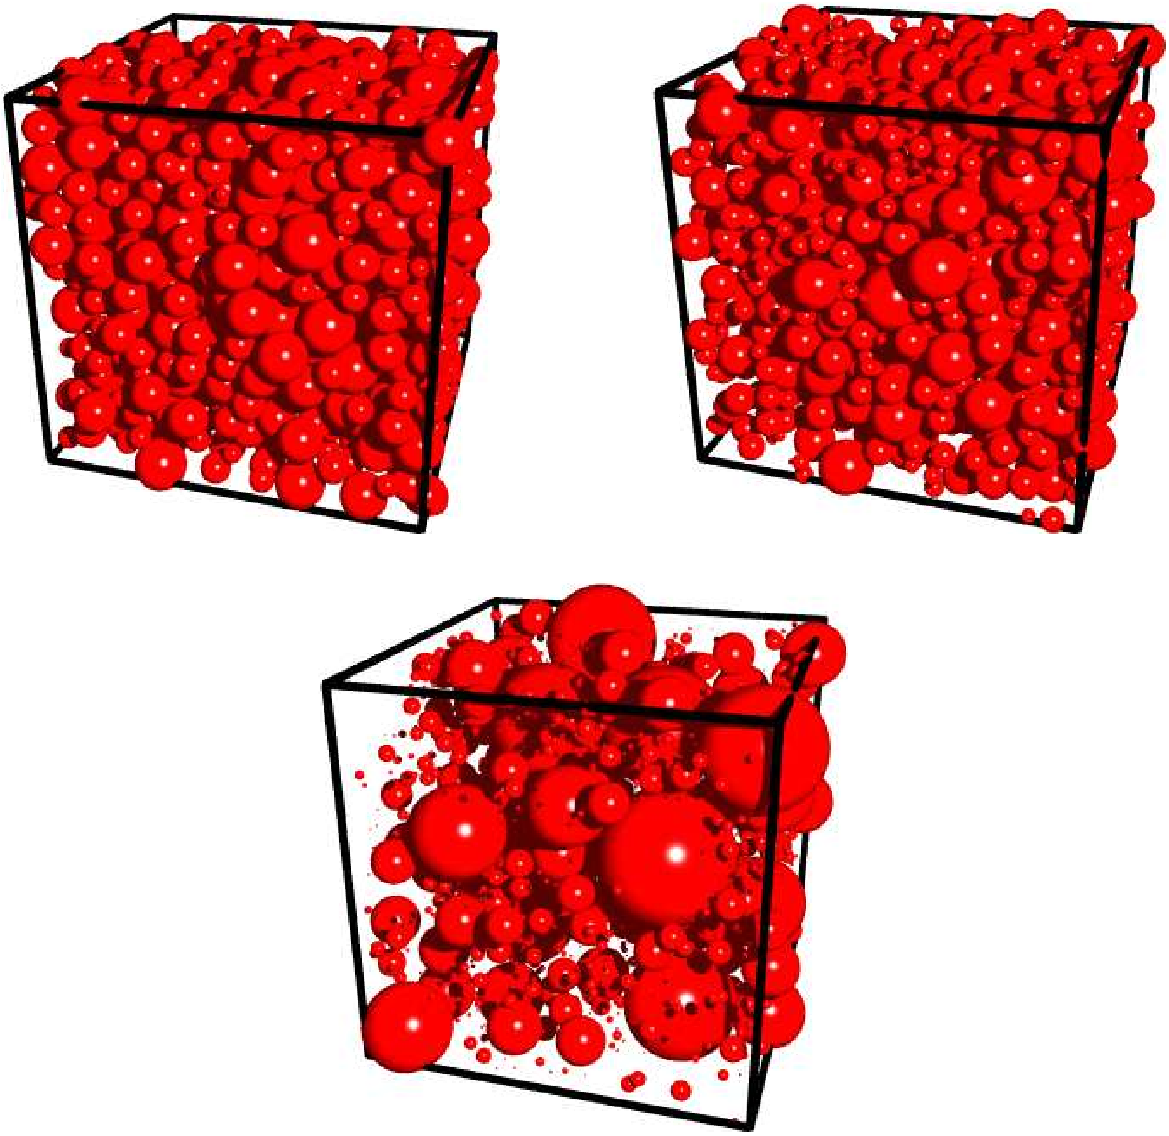

FIG. 1. Snapshots of configurations. **(a)** Schulz distribution ($z = 15$), $\eta = 0.43$, $L = 12$. **(b)** Schulz distribution ($z = 5$), $\eta = 0.38$, $L = 16$. **(c)** Log-normal distribution, $\eta = 0.29$, $L = 48$.

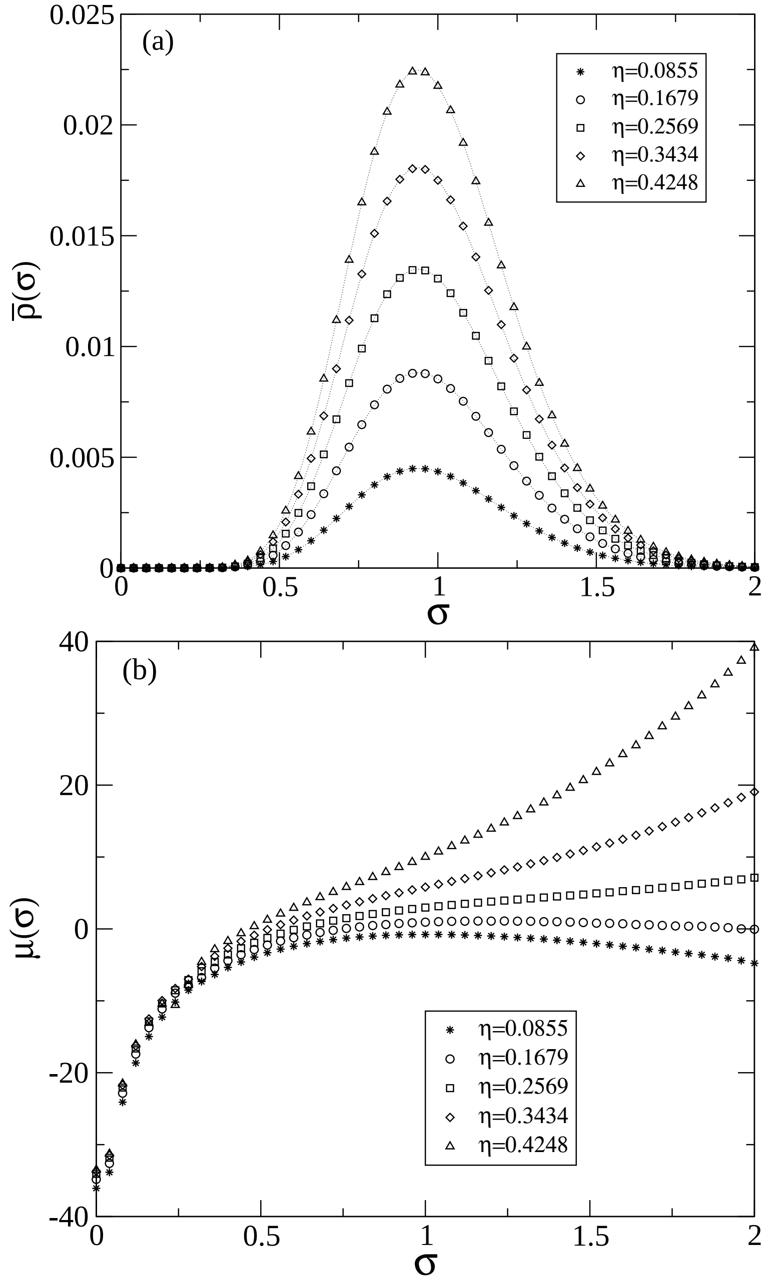

FIG. 2. Dilution line properties of hard spheres having the Schulz density distribution ($z = 15$, $\sigma_c = 3.0$) for system size $L = 9$. **(a)** Data points show the measured density distribution $\overline{\rho}(\sigma)$ at a selection of values of volume fraction η along the dilution line; dotted lines correspond to the target distribution $\rho_t(\sigma)$ **(b)** The corresponding chemical potential distribution $\mu(\sigma)$. Statistical errors do not exceed the symbol sizes.

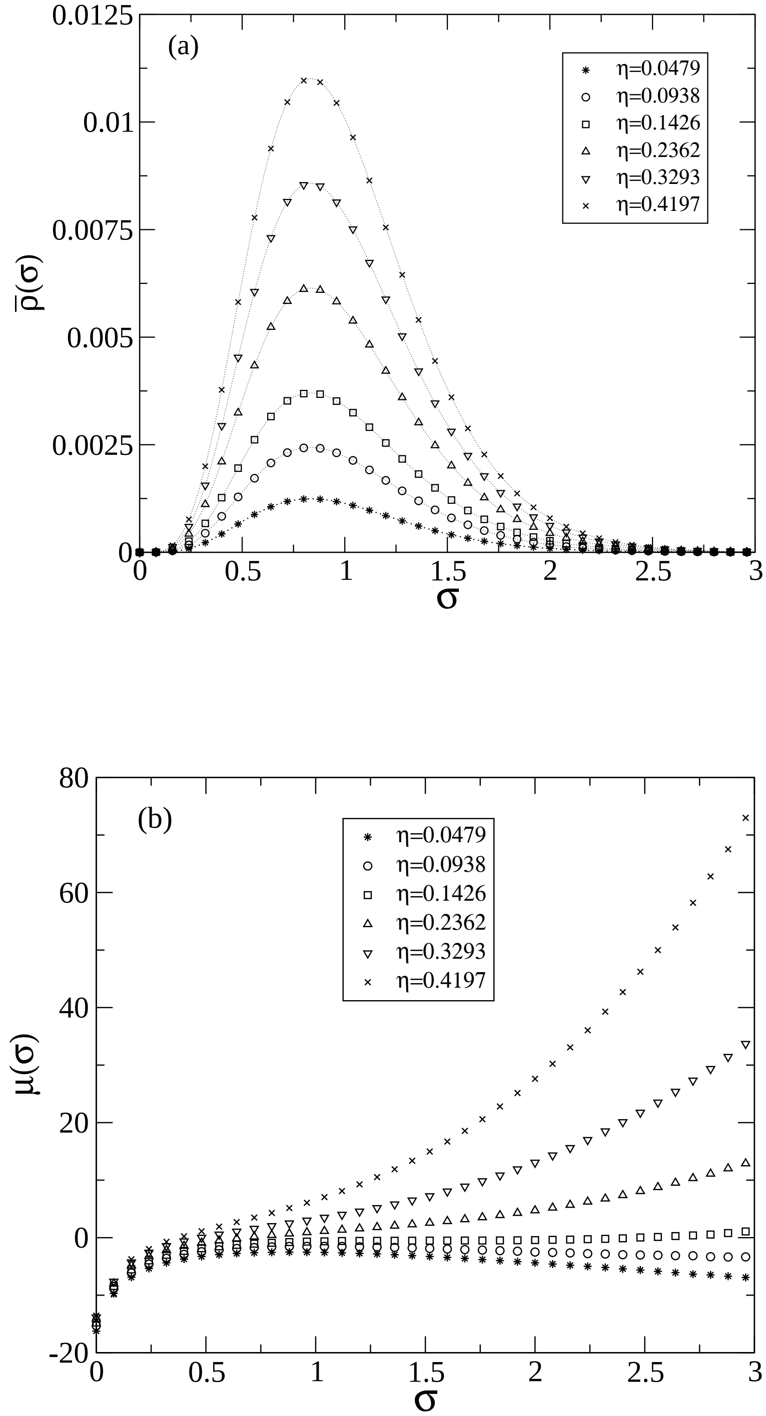

FIG. 3. Dilution line properties of hard spheres having the Schulz density distribution ($z = 5$, $\sigma_c = 4$), for system size $L = 12$. **(a)** Data points show the measured density distribution $\overline{\rho}(\sigma)$ at a selection of values of volume fraction η along the dilution line; dotted lines correspond to the target distribution $\rho_t(\sigma)$. **(b)** The corresponding chemical potential distribution $\mu(\sigma)$.

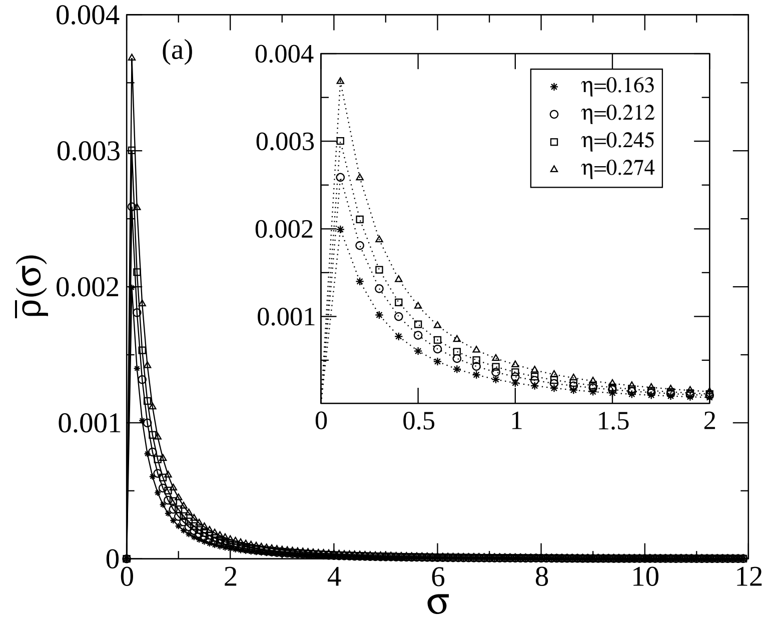

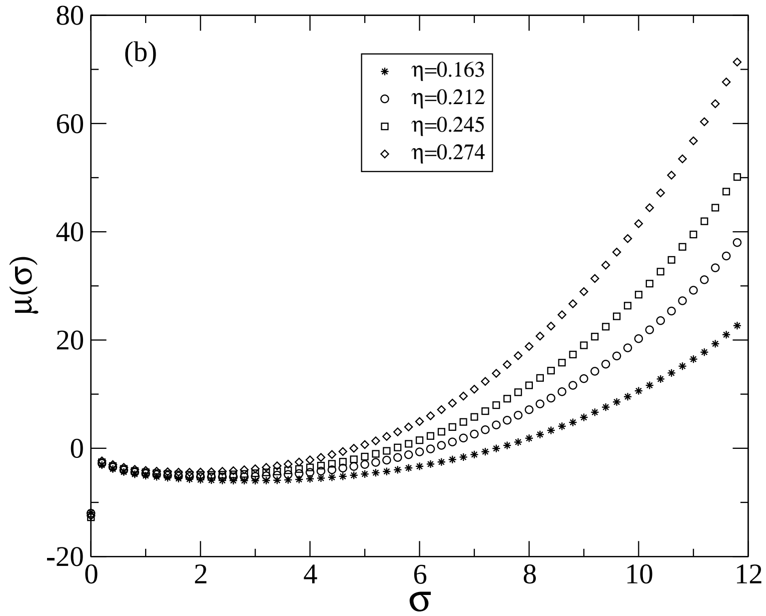

FIG. 4. Dilution line properties of hard spheres having the log-normal density distribution ($W = 2.5$, $\sigma_c = 12$) for system size $L = 36$. **(a)** Data points show the measured density distribution $\overline{\rho}(\sigma)$ at a selection of values of volume fraction η along the dilution line; dotted lines correspond to the target distribution $\rho_t(\sigma)$. The inset shows the region of small σ. Lines are merely guides to the eye. **(b)** The corresponding chemical potential distribution $\mu(\sigma)$. Statistical errors do not exceed the symbol sizes.

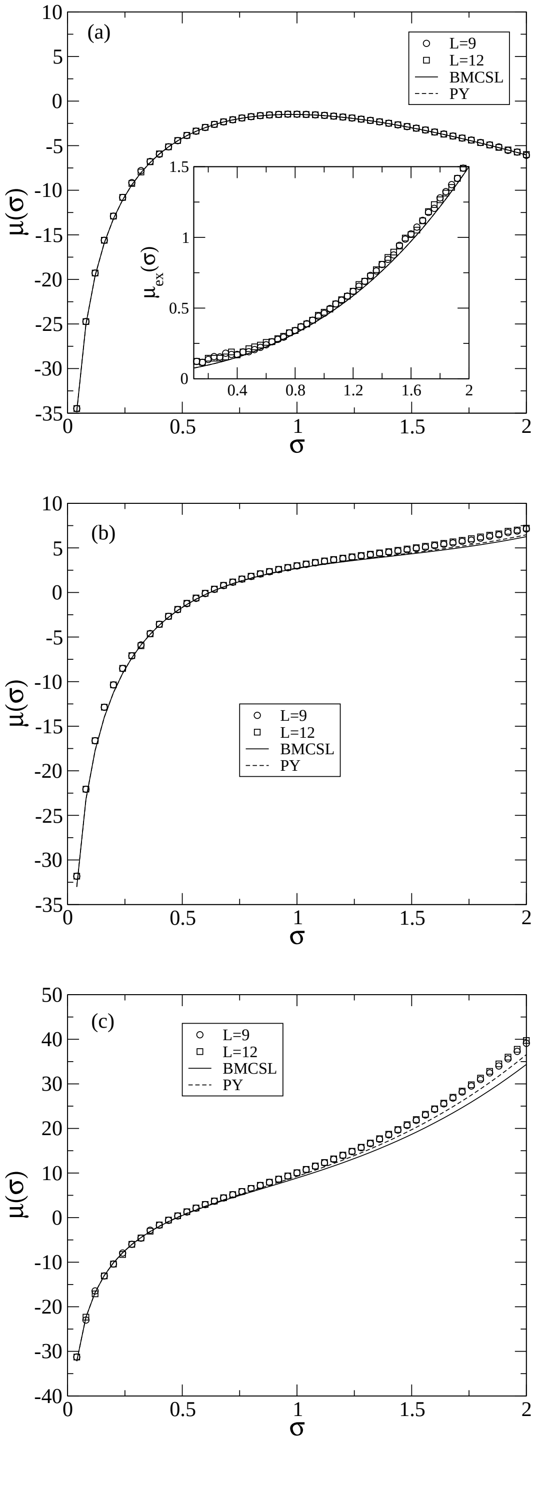

FIG. 5. Chemical potential distribution $\mu(\sigma)$ for the Schulz density distribution ($z = 15$, $\sigma_c = 3$), for system sizes $L = 9$ and $L = 12$. **(a)** $\eta = 0.056$. Here the inset shows the excess chemical potential given by eq. 3.5; **(b)** $\eta = 0.257$; **(c)** $\eta = 0.426$. Also shown in each case are the predictions of the BMCSL and PY equations of state. Statistical errors do not exceed the symbol sizes.

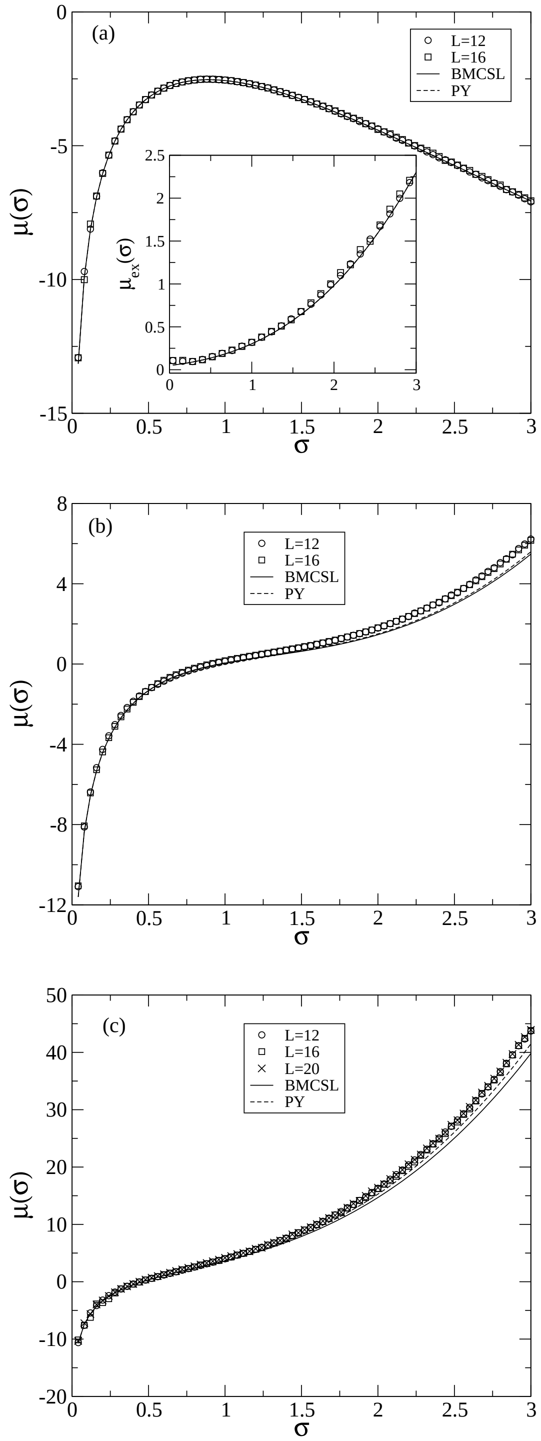

FIG. 6. Chemical potential distribution $\mu(\sigma)$ for the Schulz density distribution ($z = 5$, $\sigma_c = 4$). **(a)** $\eta = 0.048$, $L = 12, 16$, the inset shows the excess chemical potential given by eq. 3.5; **(b)** $\eta = 0.231$, $L = 12, 16$; **(c)** $\eta = 0.36$ $L = 12, 16, 20$. Also shown in each case are the predictions of the BMCSL and PY equations of state. Statistical errors do not exceed the symbol sizes.

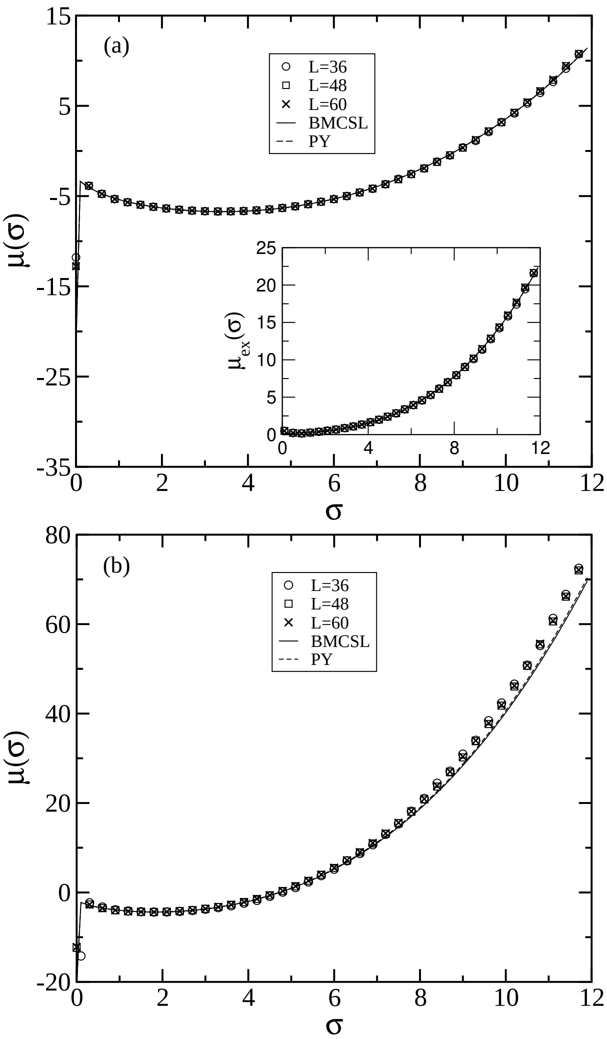

FIG. 7. Chemical potential distribution $\mu(\sigma)$ for the log-normal density distribution ($W = 2.5$, $\sigma_c = 12$) for system sizes $L = 36, 48, 60$. **(a)** $\eta = 0.126$, the inset shows the excess chemical potential given by eq. 3.5; **(b)** $\eta = 0.307$. Also shown in both cases are the predictions of the BMCSL and PY equations of state. Statistical errors do not exceed the symbol sizes.

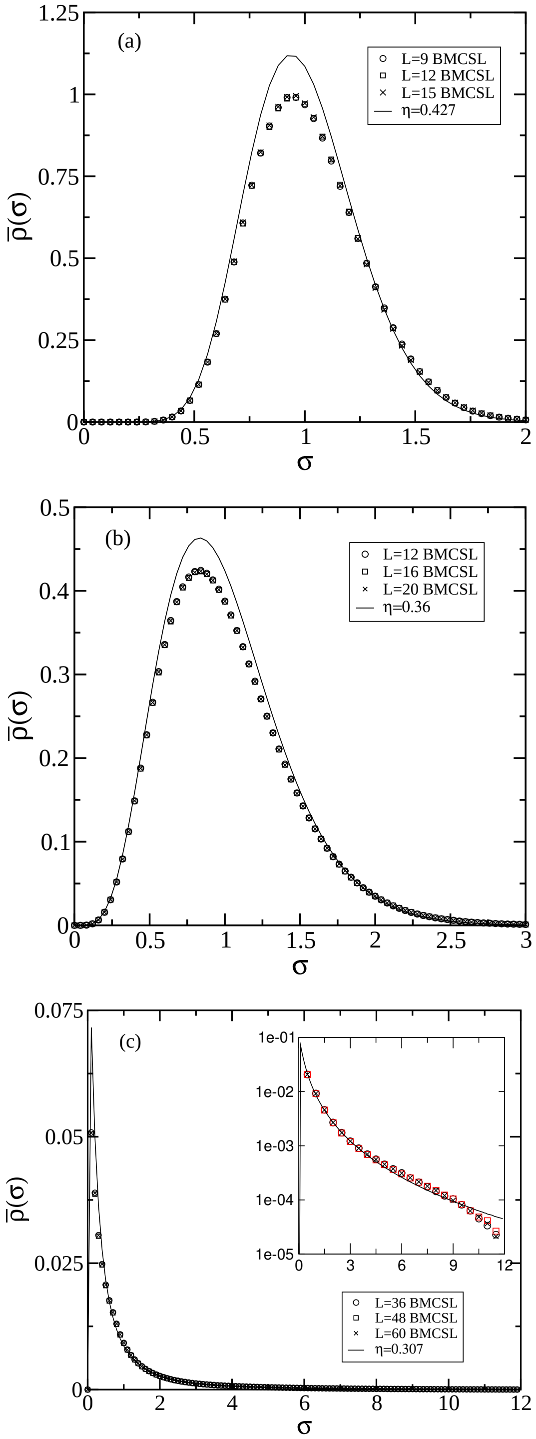

FIG. 8. The measured density distribution $\rho(\sigma)$ obtained using forms of $\mu(\sigma)$ predicted by the BMCSL equation of state at the stated volume fraction. **(a)** Schulz ($z = 15$), $\eta = 0.427$, $L = 9, 12, 15$; **(b)** Schulz ($z = 5$), $\eta = 0.36$, $L = 12, 16, 20$; **(c)** log-normal distribution, $\eta = 0.307$, $L = 36, 48, 60$, the inset shows the same data on a log scale. In each case the density distribution from which $\mu(\sigma)$ derives is shown as a solid line. Statistical errors do not exceed the symbol sizes.

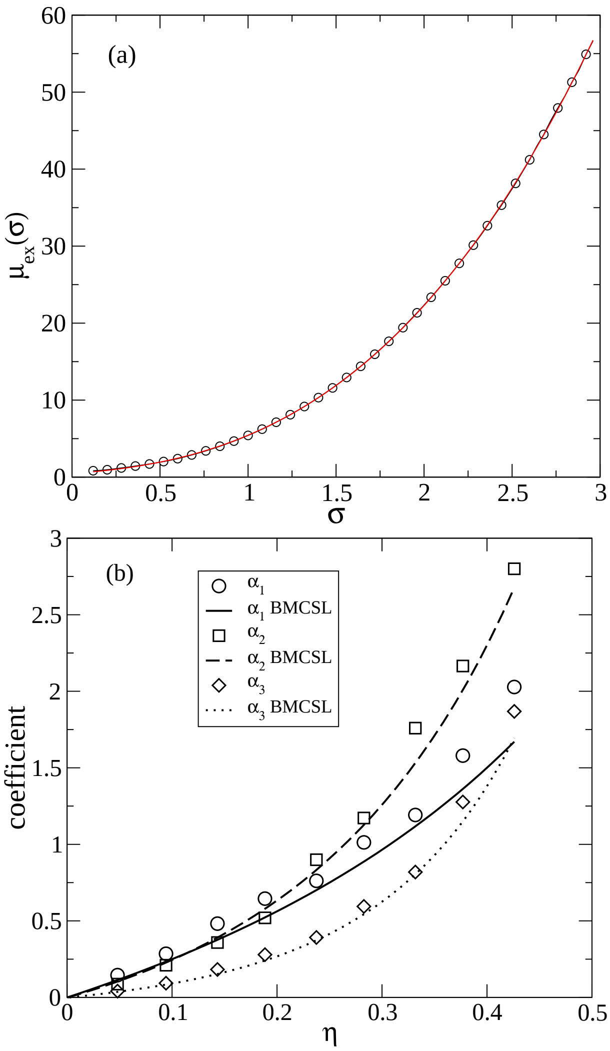

FIG. 9. **(a)** Data points show the measured excess chemical potential distribution $\mu_{\rm ex}(\sigma)$ for the Schulz density distribution ($z = 5$) at $\eta = 0.377$, for system size $L = 12$. The solid lines shows a fit to the data of the form given by eq. 3.6. **(b)** Values of the fit coefficients $\alpha_1, \alpha_2, \alpha_3$ for a selection of values of the volume fraction η. The solid lines show the predictions of the BMCSL equation. Statistical errors do not exceed the symbol sizes.

APPENDIX A: POLYDISPERSE HARD SPHERE EQUATIONS OF STATE

The equations of state that we quote here express the chemical potential as a function the sphere diameter σ. The expression due to Salacuse and Stell [20] is a generalization of the Percus-Yevick result for monodisperse hard spheres:

$$\mu_{\rm PY}(\sigma) = \ln[\rho_0 n(\sigma)] - \ln(1-\zeta^3) + \frac{\sigma^3\zeta_0 + 3\sigma^2\zeta_1 + 3\sigma\zeta_2}{1-\zeta_3} + \frac{3\sigma^3\zeta_1\zeta_2 + 9\sigma^2\zeta_2^2/2}{(1-\zeta_3)^2} + \frac{3\sigma^3\zeta_2^3}{(1-\zeta_3)^3} \tag{A1}$$

The BMCSL equation of state [18,19], on the other hand, generalizes the Carnahan-Starling expression, which for monodisperse hard spheres is more accurate than the Percus-Yevick result. It is given by

$$\begin{aligned}\mu_{\rm BMCSL}(\sigma) = {} & \ln[\rho_0 n(\sigma)] + \left(3\sigma^2\zeta_2^2/\zeta_3^2 - 2\sigma^3\zeta_2^3/\zeta_3^3\right)\ln(1-\zeta^3) \\ & + \frac{\sigma^3(\zeta_0 - \zeta_2^3/\zeta_3^2) + 3\sigma^2\zeta_1 + 3\sigma\zeta_2}{1-\zeta_3} \\ & + \frac{\sigma^3(3\zeta_1\zeta_2 - \zeta_2^3/\zeta_3^2) + 3\sigma^2\zeta_2^2/\zeta_3}{(1-\zeta_3)^2} + \frac{2\sigma^3\zeta_2^3}{\zeta_3(1-\zeta_3)^3}\end{aligned} \tag{A2}$$

where $\zeta_n = \pi\rho_0 m_n/6$ and m_n is the n-th moment of $\overline{\rho}(\sigma)$; note that $\zeta_3 = \eta$, the volume fraction of hard spheres.

Both equations of state include the ideal contribution to the chemical potential, $\mu_{\rm ideal}(\sigma) = \ln[\rho_0 n(\sigma)]$. For general temperature — remember that we set $\beta = 1$ for our hard sphere system — this contribution would read $\beta\mu_{\rm ideal}(\sigma) = \ln\bar{\rho}(\sigma)$. Because there has been some discussion in the literature regarding how the ideal chemical potential for polydisperse systems should be assigned (see e.g. [20]), it may be helpful to note that, within the grand canonical framework, $\mu_{\rm ideal}(\sigma)$ is unambiguously defined via the grand canonical partition function. Indeed, for an ideal system ($\mathcal{H}_N = 0$), one readily finds from eq. 2.1 that $\bar{\rho}(\sigma) = \exp(\beta\mu(\sigma))$, giving $\beta\mu_{\rm ideal}(\sigma) = \ln\bar{\rho}(\sigma)$ as before. One may worry about dimensions in the argument of the log here; but the definition in eq. 2.1 already implies that dimensionless units are used for lengths and for σ. Were this not the case, the integrals over $\vec{r}_i$ and σ_i would need to be normalized by a unit volume v_0 and a unit value σ_0 for the polydisperse attribute, and the ideal chemical potential would read $\beta\mu_{\rm ideal}(\sigma) = \ln[v_0\sigma_0\bar{\rho}(\sigma)]$, with the argument of the log now manifestly dimensionless. The normalization constant $v_0\sigma_0$ could in fact be made dependent on σ; this would just give a σ-dependent shift in the zero of the chemical potential scale.

REFERENCES

[1] W.B. Russel, D.A. Saville and W.R. Schowalter, *Colloidal Dispersions*, (Cambridge U.P., Cambridge, 1989).

[2] M. Baus, L.F. Rull and J.-P Ryckaert (eds.), *Observations, Prediction and Simulation of Phase Transitions in Complex Fluids* (Kluwer, Dordrecht, 1995).

[3] P. Sollich and M.E. Cates, Phys. Rev. Lett. **80**, 1365 (1998); P.B. Warren, Phys. Rev. Lett. **80**, 1369 (1998); P. Sollich, P.B. Warren and M.E. Cates, Adv. Chem. Phys. **116**, 265 (2001); P. Sollich, J. Phys. Condens. Matter (in press).

[4] Superior to the CE is a constant-(NpT) ensemble, in which the overall density can fluctuate via changes in the system volume. However the problems associated with a fixed realization of $N(\sigma)$ remain.

[5] N.B. Wilding, Phys. Rev. **E52**, 602 (1995).

[6] A.Z. Panagiotopoulos; J. Phys. Condens. Matter **12**, R25 (2000).

[7] N.B. Wilding; Am. J. Phys. **69** 1147 (2001).

[8] Conceptually and operationally (cf. sec. II B) it is expedient to regard $\rho(\sigma)$ as a binned distribution (histogram) in the limit in which the bin size $\Delta\sigma$ goes to zero. The product in eq. 2.6 is to be understood to run over all bins.

[9] A.M. Ferrenberg and R.H. Swendsen, Phys. Rev. Lett. **61**, 2635 (1988); *ibid* **63**, 1195 (1989).

[10] D. Frenkel and B. Smit; *Understanding Molecular Simulation*, Academic Press, Boston (1996).

[11] There is some license with regards to the choice of values for σ_c and M. Clearly, where computationally practical, σ_c should be chosen to exceed the maximum value of σ for which significant weight occurs in $\rho(\sigma)$ at any point along the phase space path of interest. The choice of M should strike a suitable balance between data resolution and smoothness of the histogram.

[12] For anything other than hard particles, the energy u should also be written to the list, because it is required for histogram reweighting.

[13] Note that, by construction $\mu'(\sigma)$ is tuned so that $\rho'(\sigma)$ matches the intrinsically *smooth* target function. Thus, in contrast to most implementations of histogram reweighting, departure from the range of reliable extrapolation is *not* signaled by a histogram that appears ragged in its extremal regions.

[14] We have found good results using the DFP version of a quasi-Newton algorithm described in section 10.7 of *Numerical Recipes*; W.H. Press, S.A. Teulolsky, W.T. Vetterling, B.P. Flannery, Cambridge University Press (1992).

[15] Once the extrapolation has been made, and a new simulation performed, it is advisable to check the reliability of the extrapolation by comparing the measured $\rho(\sigma)$ with the target function.

[16] Owing to the lack of a temperature scale in hard sphere systems, we implicitly adopt the convention of assigning $\beta = 1$.

[17] For the case of hard spheres, the form for $\mu(\sigma)$ could equally well have been bootstrapped at low densities using the predictions of one of the analytical equations of state [18–20].

[18] T. Boublík, J. Chem. Phys. **53**, 471 (1970).

[19] G.A. Mansoori. N.F. Carnahan, K.E. Starling, T.W. Leland, J. Chem. Phys. **54**, 1523 (1971).

[20] J.J. Salacuse and G. Stell, J. Chem. Phys. **77**, 3714 (1982).
[21] P. Bartlett, J. Chem. Phys. **107**, 188 (1997).
[22] D. A. Kofke and E.D. Glandt, J. Chem. Phys. **87**, 4881 (1987).
[23] P.G. Bolhuis and D.A. Kofke, Phys. Rev. **E54**, 634 (1996); D.A. Kofke and P.G. Bolhuis, Phys. Rev. **E59**, 618 (1999).
[24] T. Kristóf and J Liszi, Mol. Phys. **99**, 167 (2001).
[25] P. Bartlett and P.B. Warren, Phys. Rev. Lett. **82**, 1979 (1999).
[26] P.B. Warren, Euro. Phys. Lett. **46**, 295 (1999).
[27] J.A. Cuesta, Eur. Phys. Lett. **46**, 197 (1999).
[28] R.P. Sear, Phys. Rev. Lett. **82**, 4244 (1999).
[29] This mirrors a similar finding for the BMCSL equation in the context of highly asymmetric binary hard sphere mixtures: T. Coussaert and M. Baus, J. Chem. Phys. **109**, 6012 (1998).
[30] Steps in this direction have recently been reported by C. Barrio and J.R. Solana, J. Chem. Phys. **113**, 10180 (2000).