



Sweatman, M.B. and Quirke, N. (2004) Simulating fluid-solid equilibrium with the Gibbs ensemble. *Molecular Simulation*, 30 (1). pp. 23-28. ISSN 08927022

<http://strathprints.strath.ac.uk/13476/>

Strathprints is designed to allow users to access the research output of the University of Strathclyde. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. You may not engage in further distribution of the material for any profitmaking activities or any commercial gain. You may freely distribute both the url (<http://strathprints.strath.ac.uk>) and the content of this paper for research or study, educational, or not-for-profit purposes without prior permission or charge. You may freely distribute the url (<http://strathprints.strath.ac.uk>) of the Strathprints website.

Any correspondence concerning this service should be sent to The Strathprints Administrator: eprints@cis.strath.ac.uk

Self-referential Monte Carlo method for calculating the free energy of crystalline solids

M. B. Sweatman*

Department of Chemical and Process Engineering, University of Strathclyde, Glasgow, G1 1XJ, United Kingdom

(Received 2 March 2005; published 20 July 2005)

A self-referential Monte Carlo method is described for calculating the free energy of crystalline solids. All Monte Carlo methods for the free energy of classical crystalline solids calculate the free-energy difference between a state whose free energy can be calculated relatively easily and the state of interest. Previously published methods employ either a simple model crystal, such as the Einstein crystal, or a fluid as the reference state. The self-referential method employs a radically different reference state; it is the crystalline solid of interest but with a different number of unit cells. So it calculates the free-energy difference between two crystals, differing only in their size. The aim of this work is to demonstrate this approach by application to some simple systems, namely, the face centered cubic hard sphere and Lennard-Jones crystals. However, it can potentially be applied to arbitrary crystals in both bulk and confined environments, and ultimately it could also be very efficient.

I. INTRODUCTION

There are many Monte Carlo methods [1–11] in the literature for calculating the absolute free energy of classical crystalline solids but none of them are entirely satisfactory. Here, I describe a different approach for such systems, a “self-referential” (SR) Monte Carlo technique for calculating the free energy of classical crystalline solids, both in the bulk and in confined spaces. The aim of this work is to demonstrate the feasibility of this approach by application to some simple crystalline systems, namely, the face centered cubic crystals of hard sphere and Lennard-Jones particles. Although this general approach has been considered before by Barnes and Kofke [12] in the context of a “hard rod on a line” system, and there are also some strong parallels with the work of Mon and colleagues [13,14] who studied lattice systems, the method provided here is quite different and it can potentially be applied to arbitrary crystals whether in bulk or confined environments. The current implementation of this technique is rather inefficient, but in the summary I outline some ideas that should allow this method to be very efficient.

Phase transitions involving crystalline solids are often strongly first order and associated with significant hysteresis. So knowledge of the free energy of crystalline solids is important, particularly in computer simulations where, because of relatively small system sizes, nucleation events can be rare. At this stage it is also important to note a significant technical difference between simulation studies of crystals in bulk and confined environments. For a bulk system the experimental pressure and temperature can be imposed by simulating in the isothermal-isobaric (*NPT*) ensemble. This means that only free-energy differences between competing phases need be determined as a function of pressure to decide the location of an equilibrium transition. However, when simulating a confined system we are usually limited to

small systems that do not include the confined system–bulk system interface. So, in general, we cannot impose or measure the bulk (i.e., experimental) pressure corresponding to the confined state. Instead, it is essential that the absolute chemical potential is known, because this quantity is the same in the confined and bulk systems at equilibrium. Unfortunately, we cannot impose the chemical potential on simulated crystals (i.e., we cannot simulate crystals in the grand-canonical ensemble) because of difficulties associated with periodic boundary conditions and particle number fluctuations (this is explained in detail later). Rather, simulations *should* be performed with a fixed number of particles and the absolute chemical potential should be measured. For a pure system the chemical potential is simply the Gibbs free energy per particle of the confined isothermal-isobaric ensemble. So for simulation studies of confined crystals calculation of the *absolute* Gibbs free energy is essential if conditions inside the pore are to be related to experimental (or bulk) conditions, regardless of whether phase behavior is of interest or not.

Given the undoubted importance of calculating the free energy in Monte Carlo simulations it is no surprise that a spectrum of various approaches already exists for this purpose. Every method calculates the free-energy difference between the state of interest and another state whose absolute free energy can be calculated relatively easily. So they can all be classified according to (1) the reference state for which the absolute free energy is calculated, and (2) the approach used to calculate the free-energy difference. Several options are available for each. Consider the choice of reference states first. Many methods employ a simple model crystal, such as the Einstein [3] or acoustic [15] crystals, as a reference state. All the others [4–6] use a (meta)stable fluid state. These choices are not entirely satisfactory because these reference states are very different from the crystalline solid state of interest. This is particularly important for molecular and confined systems where the complexity of the path, or the “path length,” between the reference state and the state of interest can be long and sometimes difficult to define. The key and radical contribution of the work described here is that the

reference state is exactly the same state as the final state (hence the term self-referential, coined by Barnes and Kofke [12]), except that the reference state has a different number of particles (or crystal unit cells to be precise). This choice has several distinct advantages that will potentially allow short paths between these two states to be easily defined. They are as follows: (1) a direct path between initial and final states is traversed avoiding first-order phase transitions (longer paths designed to avoid these transitions are unnecessary), (2) the path needs to traverse only relevant degrees of freedom and need not traverse other degrees of freedom or parameters, i.e., the molecular model (both molecule-molecule and molecule-wall interactions) is constant along the path, and (3) together with the retiling algorithm (see later) only the absolute minimum number of particles need travel this path, i.e., only one crystal unit cell for a bulk system (in all other methods all particles in the ensemble travel the path). These aspects are discussed in more detail in the comparison section below and in the Summary.

Many alternatives exist for calculating free-energy differences [2,8,9], including thermodynamic integration (TI), free-energy perturbation, and “parameter hopping” (or multicanonical simulation) as seen in the umbrella-sampling [2], Landau free energy [5], and phase-switch [4] approaches. It is possible to couple numerical integration (quadrature) with each of these techniques to produce an efficient method for traversing the required free-energy difference. In the implementation described below, parameter hopping is used without quadrature to traverse the path, i.e., the entire free-energy difference is traversed without any interpolation by a succession of connected hops, and consequently the SR approach described here is not particularly efficient. Some ideas that could dramatically improve its efficiency are discussed in the Summary.

The remainder of this paper is organized as follows. In the next two sections I discuss some issues relating to the motivation for this work. Then I describe the SR method as it is applied to the simple systems considered here. I avoid attempting a more general definition of the methodology that could be applied to arbitrary crystals, which would be rather complicated, in favor of clarity, which is more important considering that the aim of this work is to validate a new idea. This developed methodology is applied to calculate the chemical potential of pure fcc hard-sphere and Lennard-Jones bulk crystals. Results are presented at conditions of solid-fluid coexistence and compared with existing reference data. Finally, I discuss potential future work including ideas to improve efficiency.

A. The choice of ensemble

In a conventional simulation of a crystalline solid with periodic boundary conditions the number of lattice sites is fixed. Further, in a grand-canonical-ensemble simulation the volume is fixed, and it follows that for a space-filling crystal the lattice site density must also be fixed. If the lattice site density is fixed then, clearly, there is no way that the simulated crystal can attain the correct equilibrium lattice site density unless, somehow, it is initiated in this state. This is

the principal reason why space-filling crystals cannot be simulated in the grand-canonical ensemble. Of course, another problem has to do with the probability of acceptance of insertion and deletion moves, but this is only a practical problem that might be tackled with sufficient computing resources or clever algorithms. The problem concerning the lattice site density is more fundamental. Note that this is not a finite-size effect, i.e., it is not necessarily resolved by simulating a sufficiently large system because no matter how large the system one cannot guarantee that it will be initiated with the correct lattice site density. For two-dimensional (2D) crystals, for example, those that form as an adsorbed layer on a surface, the lattice site density is fixed if the simulation box area parallel to the layer is fixed. So equilibrium 2D crystals can only be simulated if the simulation box lengths parallel to the surface are allowed to fluctuate. This holds even if the 2D layer is in contact with a fluid phase.

Of course, simulations of pure (defect-free) crystalline solids can be carried out in the NPT ensemble for bulk crystals and the isothermal-isotension ensemble ($N\sigma T$) for confined systems. For a system confined in an ideal slit pore, for example, interfacial tension is fixed in a direction parallel to the slit. Now the lattice site density is free to fluctuate and attain the true equilibrium state. This approach is fine for bulk crystals where the simulated pressure corresponds to experimental pressure. But crystals confined within pores of fixed width (slit pores, for example) do not experience the bulk pressure directly, unless the confined system–bulk system interface is simulated. Instead, to establish the experimental or bulk conditions corresponding to a confined state one must know the chemical potential, because the chemical potential is constant throughout an equilibrium system regardless of inhomogeneity. So we would like to simulate confined crystals in the grand-canonical ensemble where the chemical potential is imposed, yet we cannot. Rather, we should simulate them in an $N\sigma T$ ensemble and *measure* the absolute chemical potential (and hence the Gibbs free energy) to establish experimental (or bulk) conditions. Despite this, there are several examples [16–19] in the literature where the grand-canonical ensemble has been used to simulate space-filling crystalline solids and 2D crystals with fixed box area. All this work should be considered carefully, because in every case these are not equilibrium simulations. Note also the simulations by Dominguez and colleagues [20] in which crystalline solid free energies for a slit-pore system are calculated in an ensemble in which the slit-pore area is fixed. We can expect that their free-energy calculations and hence their predicted phase transition points are dependent on their choice of slit-pore area, a fact admitted by them when referring to the strain on their simulated crystals.

The key point of this discussion is to emphasize the importance of methods for calculating the chemical potential (or equivalently the Gibbs free energy) of confined crystals. This is an important factor in the motivation for this current work.

B. Comparison with other approaches

The above comments concerning the advantages of the self-referential approach deserve more detailed justification.

Given the plethora of alternative methods this justification could become exorbitantly long, so this section will concentrate on comparing the self-referential approach with some of the more popular and robust methods in the literature, in particular the lattice-coupling-expansion TI method [2], the Landau free-energy approach [5], and the phase-switch method [4] (these last two methods were not originally designed to calculate absolute free energies, but provided the absolute free energy of the reference fluid phase can also be calculated they can clearly be used for this purpose). To compare these approaches let us consider a cubic phase of ice confined within a slit pore at fixed conditions of pore width, temperature, and interfacial tension (or equivalently fixed grand potential per unit area of slit). As we are using this system as a hypothetical application there is no need to worry about whether this phase does actually exist.

Let us consider the lattice-coupling-expansion TI method first. With this method the (canonical) Helmholtz free energy is calculated at fixed slit-pore area. The first step in the TI scheme is to introduce a center-of-mass constraint so that subsequent paths can be traversed. The details of this step are not trivial [2] and, indeed, recent work [21] has corrected previous errors associated with this step. Next, we traverse a path that gradually couples the water molecules to their lattice sites with strong harmonic springs. Note that orientational [22] as well as translational coupling of the water molecules might be useful for this material in order to avoid traversing first-order phase transitions along subsequent paths. If orientational coupling is used then it might be advantageous to gradually turn off any electrostatic dipoles on the water molecular model—this could help with the subsequent expansion stage. Next, the area of the slit pore can be gradually increased until interactions between water molecules are insignificant. Note that the area of the final expanded state will depend on whether the electrostatic dipole interactions have previously been turned off [because permanent dipole–permanent dipole interactions decay more slowly (r^{-3}) than induced dipole–induced dipole (or dispersion) interactions (r^{-5})]. Next, any orientational coupling will need to be gradually turned off. Finally, the water molecule–pore wall interactions will need to be gradually turned off to arrive at the noninteracting Einstein crystal. Note that a further step involving calculating the free energy of the ideal gas (with orientational degrees of freedom) might be needed. Finally, this entire calculation should be repeated for a range of system sizes to ascertain the significance of finite-size effects induced by simulating in an ensemble with fixed slit-pore area.

Next, consider the Landau free-energy approach applied in an $N\sigma T$ ensemble consisting of fixed interfacial tension, number of water molecules, and temperature, where a path from the cubic crystalline phase to fluid water is to be traversed. The free energy of this ensemble is simply the water chemical potential multiplied by the number of water molecules. First, an order parameter that can effectively distinguish different solid and fluid phases must be formulated. Once this is achieved the simulation proceeds by comparing the statistical likelihood of macrostates constrained to narrow ranges of this order parameter—stepping continuously through the order parameter from the crystal to the fluid (or

vice versa) generates the required free-energy difference, avoiding any first-order phase transitions. To be successful the resulting fluid phase must be (meta)stable. If no stable liquidlike phase exists at the required conditions then a path in thermodynamic space (temperature and interfacial tension in this case) might be traversed to ensure the existence of such a fluid state. So some understanding of the phase behavior of the fluid state is required in advance. Finally, the free energy of the fluid water is obtained by standard methods. The efficiency of this approach depends on the free-energy barrier between initial and final states. The height of this barrier is not known *a priori*—it will depend on the order parameter, but it will essentially be proportional to the number of particles in the simulation.

Now let us consider the phase-switch method applied, again, in an $N\sigma T$ ensemble. This is essentially a parameter-hopping approach that measures the free-energy difference between two states by sampling both states in the same simulation. Specialized Monte Carlo (MC) moves allow the system to switch from one state to the other (by performing a global coordinate transformation) provided the system is sufficiently close to special microstates where this switch has a reasonable probability of being accepted. The system is biased to ensure that it approaches these special microstates with sufficient regularity. This is not actually so different from the Landau method above; in this case the bias directs the fluid subsystem toward special “frozen” or fixed fluid configurations that can then be rearranged (via the global coordinate transformation) to generate the crystal. So the same remarks pertain to this approach as to the Landau approach above, except that a different order parameter is needed.

Finally, consider the SR approach. Once again simulations are performed in the $N\sigma T$ ensemble. Two steps are performed (and are explained in more detail later). The first is a single parameter-hopping step that doubles the number of unit cells in the simulation—this is achieved by replication whereby a double-size system is created consisting of two nearly identical (to within a tolerance) single-size systems. The second step allows the coordinates of the newly created water molecules to relax fully from their initial highly constrained values.

Some of the advantages of the SR approach become clear in this example. Although the lattice-coupling-expansion TI approach is probably feasible, the path is arduous involving many different segments, including the center-of-mass adjustment, translational coupling, possible orientational coupling, electrostatic dipole elimination, expansion, and finally elimination of the pore walls. Great effort is made to define a path that avoids first-order phase transitions. If no such path can be found then additional parameter hops can be introduced to traverse any phase transitions. For very complex materials defining such a path could well be difficult and frustrating. In contrast, the phase-switch, Landau free-energy, and SR approaches offer a more concise path. However, the phase-switch and Landau free-energy approaches rely on the existence of a (meta)stable liquidlike phase whose free energy still needs to be calculated, and an order parameter must be defined in advance. So the advantages of the SR approach will become more important as the material

becomes more complex, for example, crystals of proteins and pharmaceuticals confined in structured pores. Enhancement of the SR approach in line with ideas outlined in the Summary could potentially lead to it also being a very efficient approach.

II. THEORY

Our aim is to calculate the Gibbs-free-energy difference between two crystalline solid systems with differing numbers of unit cells (but are otherwise identical). If one system has m unit cells and the other has $n > m$, then the free energy of this structure is $x\Delta G/(n-m)$ where x is the number of unit cells and ΔG is the Gibbs-free-energy difference between these two systems. For a pure crystal the chemical potential is then $\Delta G/N_c(n-m)$ where N_c is the number of particles per unit cell. So, to be clear, the SR approach is simultaneously both a free-energy difference and an absolute method. It calculates the Gibbs-free-energy difference between two systems that differ only in their size, yet because the Gibbs free energy scales linearly with size the absolute Gibbs free energy (and chemical potential for a pure system) is obtained automatically. This idea was first expounded by Barnes and Kofke in a slightly restricted sense, i.e., they formulated and applied this SR idea in the canonical ensemble where it is exact only in the thermodynamic limit (although their model system of hard rods was found to be sufficiently large that finite-size effects were unimportant). Note also the parallels with the work of Mon. In this work the free-energy difference between lattice systems (once again in the canonical ensemble) that differ only in their size is calculated, but linearity of the free energy with system size is explicitly not assumed. So this is not a SR approach in the fullest sense defined here.

For a mixture we must also know the difference in chemical potential between each species to obtain the absolute chemical potentials. But this is not the focus of this work, and is a problem with all methods that calculate Gibbs free energies. For alloys, it might be resolved using a semi-grand-ensemble simulation [2].

The systems analyzed here are pure bulk crystalline solids with cubic unit cells where each particle only has translational degrees of freedom (denoted by \mathbf{r}) and the discussion below reflects this. Modifications that arise when considering more complex systems, such as confined or molecular crystals, will be discussed briefly in the Summary. In this implementation of the SR approach we choose for convenience $n=2m$ and ΔG is split into two contributions. The first measures the free-energy difference between the small (single-sized) system at pressure P_s and temperature T_s , and the large (double-sized) system at P_{d1} and T_{d1} , where one-half of the double-size system is a replica, to within a tolerance, of the other half. The second contribution measures the free-energy difference between this constrained double-size system and the unconstrained double-size system at P_s and T_s . MC schemes to measure each of these contributions are described below separately. Alternative schemes with different choices for n are possible and discussed in the Summary.

Each scheme is based on parameter hopping, i.e. two neighboring states, labeled a and b , are simulated simulta-

neously in a multicanonical ensemble. The free-energy difference between these two states, ΔG_{ab} , is simply

$$\Delta G_{ab} = -k_B T \ln(p_b/p_a) \quad (1)$$

where $p_a = 1 - p_b$ is the unbiased probability that the simulation is in macrostate a . This free-energy difference is also the bias needed in a non-Boltzmann sampling regime to ensure $p_a = p_b$.

The Gibbs free energy is related to the corresponding partition function Ψ via the fundamental relation

$$G = -k_B T \ln(\Psi). \quad (2)$$

Following the arguments put forward by Wilding [23] the isothermal-isobaric partition function for a classical crystalline solid is presented as

$$\Psi(N, P, T) = \int_0^\infty dV e^{-\beta PV} \frac{\Lambda^{-N}}{N} \int_{V'} d\mathbf{r}^N e^{-\beta H(\mathbf{r}^N)} \quad (3)$$

where H is the configurational contribution to the Hamiltonian, $\beta = 1/k_B T$ is the inverse temperature, V is the volume, and Λ is obtained by integrating momentum degrees of freedom. The factor N^{-1} deserves some explanation. Normally this factor, which accounts for double-counting states that have identical quantum wave functions, would be $N!$. But in Eq. (3) the particles are not allowed to permute their lattice positions (signaled by the constraint V'). So the number of identical states, allowing for periodic boundary conditions, is now only N . Of course, this is not a true partition function because it is not dimensionless. Work by Koper and Reiss suggests that the proper partition function for the NPT ensemble is

$$\begin{aligned} \Psi(N, P, T) &= \int_0^\infty dV \beta P^*(V) e^{-\beta PV} \frac{\Lambda^{-N}}{N} \int_{V'} d\mathbf{r}^N e^{-\beta H(\mathbf{r}^N)} \\ &= \int_0^\infty dV \beta P^*(V) e^{-\beta PV} Q_N \end{aligned} \quad (4)$$

where

$$P^* = \frac{\partial \ln(Q_N)}{\beta \partial V} \quad (5)$$

is the volume-dependent internal pressure, which is distinct from the pressure P imposed by the external reservoir. The deviation of P^* from P increases as V decreases, and is thought to be significant only for small systems. In this work this factor is omitted, potentially leading to another source of finite-size error in addition to that induced by periodic boundaries.

The partition function can be reexpressed [23] in a slightly simpler form by ‘‘clamping’’ the position of particle 1,

$$\Psi(N, P, T) = \int_0^\infty dV e^{-\beta PV} \frac{V}{N} \int_{V'} d\mathbf{r}^{N-1} e^{-\beta H(\mathbf{r}^N)} \quad (6)$$

where all factors of Λ are dropped as they play no role in this study, recognizing that calculated free energies are configu-

rational contributions only. This is the final form of the partition function used in this work.

A. Doubling and halving

MC moves attempt to double the system (if it is single sized) or halve the system (if it is double sized). These attempts are chosen with equal probability. A doubling move doubles the volume and number of particles. Particles are created with coordinates determined by their labels, i.e.,

$$\mathbf{r}_j = \mathbf{r}_{j-N} + \mathbf{L}_x + \Delta \mathbf{r}_j \quad (7)$$

for $j=N+1, \dots, 2N$, where \mathbf{L}_x is a vector, with length L_x , equivalent to one of the primary single-size simulation box lengths. $\Delta \mathbf{r}_j$ is a random displacement, chosen with uniform probability from within a sphere of radius $\alpha_1 L_x$, where $\alpha_1 \ll 1$ is the initial tolerance. A halving move halves the volume and deletes particles $N+1, \dots, 2N$.

In addition to these moves, other MC moves ensure temperature and pressure equilibration. For temperature equilibration the position coordinate displacement moves are specialized to improve sampling in the double-size system. In this case two particles, with labels j and $j+N$, are both moved simultaneously (j is chosen at random) by the same amount $\delta \mathbf{r}_j$, which is chosen as per the usual displacement selection criteria for particles in the small system. Then one of the particles (j or $j+N$, chosen randomly with equal probability) is displaced by a further smaller amount $\delta \mathbf{r}'$. If the particles exceed their tolerance then the move is rejected, i.e., if

$$|\mathbf{r}_j - (\mathbf{r}_{j+N} - \mathbf{L}_x)| > \alpha_i L_x \quad (8)$$

for $j=1, \dots, N$ (and where $i=1$ in this case). If this multiple-move approach were not adopted then only individual moves within the tolerance criteria would be allowed, and such moves would be very small and inefficient when α_i is small. Of course, the usual individual particle displacement moves are performed instead in the single-size system, and the position of particle 1 is never moved because it is clamped. To equilibrate pressure the usual volume scaling moves are attempted, i.e., the position coordinates of each molecule (including the clamped molecule) are scaled by randomly choosing a step in the volume with uniform probability on the range $\pm L_V$.

The acceptance criteria for these moves are defined to achieve microscopic reversibility. This requires satisfaction of

$$\pi_o p_{on}^{ch,acc} = \pi_n p_{no}^{ch,acc} \quad (9)$$

where π_x is the probability density that the system resides in an infinitesimal element of coordinates corresponding to microstate x , p_{on}^{ch} is the probability of choosing the trial move such that microstate o (the old state) is transformed to microstate n (the new state) given that the system resides in microstate o , and p_{on}^{acc} is the probability of accepting this trial move. The acceptance probabilities are related by

$$p_{on}^{acc} = \min(1, q_{on}^{acc}), \quad (10)$$

$$q_{on}^{acc} = (q_{no}^{acc})^{-1}.$$

From the partition function (6) we have

$$\pi \propto \frac{V}{N} e^{-\beta PV} e^{-\beta H} dV d\mathbf{r}^{N-1} \quad (11)$$

and so given p_{on}^{ch} and p_{no}^{ch} the p_{on}^{acc} are uniquely defined.

Consider molecular coordinate moves first. For the single-size system q_{on}^{acc} is exactly the same as for any conventional isothermal simulation if $p_{on}^{ch} = p_{no}^{ch}$. For the double-size system we have

$$q_{on}^{acc} = \exp[-\beta_{di}(H_n - H_o)] \quad (12)$$

as well, where the subscript di denotes the double-size system corresponding to tolerance α_i . But now, depending on the replication tolerance, the value of the configurational energy of the double-size system is approximately double its value for one-half of this system. So for $i=1$ we can write

$$q_{on}^{acc} \approx \exp[-2\beta_{d1}(H_n^h - H_o^h)] \quad (13)$$

where the superscript h denotes the Hamiltonian of one-half of the double-size system. This indicates that the constrained double-size system will behave like an unconstrained system at a lower temperature. So, if we set $T_{d1} = T_s$, where subscript s denotes the single-size system, then in general the crystal will contract. This needs to be avoided so that system doubling and halving moves are performed efficiently (see below). Hence, we instead choose $T_{d1} > T_s$ to ensure that the double-size system does not significantly contract and so that the energy of the constrained double-size system with tolerance α_1 is sufficiently close to twice the energy of the single-size system.

Now consider volume moves where $p_{on}^{ch} = p_{no}^{ch}$. Because particle 1 is clamped we have

$$q_{on}^{acc} = \frac{\pi_n}{\pi_o} = \frac{\gamma V_o e^{-\beta H_n} e^{-\beta P \gamma V_o} dV (\gamma d\mathbf{r})^{N-1}}{V_o e^{-\beta H_o} e^{-\beta P V_o} dV d\mathbf{r}^{N-1}} = \gamma^N e^{-\beta(H_n - H_o)} e^{-\beta P V_o (\gamma - 1)} \quad (14)$$

where $\gamma = V_n / V_o$, i.e., precisely the same acceptance criteria as for conventional isothermal isobaric simulations [2]. Note that because the tolerance is defined to scale with system size these volume-scaling moves can never violate the tolerance constraints (8). However, there is an additional complication for the double-size system. If we consider a crystal of hard spheres and denote the probability that no particle overlaps are generated on attempted compression for the unconstrained double-size system by p_c , then for the constrained double-size system this probability is roughly $\sqrt{p_c}$, depending on the replication tolerance. Expansion moves are not affected in this way. So in general the constrained double-size system behaves like an unconstrained system at a higher pressure than P_{d1} . So, just as for particle displacement moves above, if we set $P_{d1} = P_s$ then in general the crystal will contract. Hence, we instead choose $P_{d1} < P_s$ to ensure that the volume of the constrained double-size system with tolerance α_1 is sufficiently close to twice the volume of the single-size system.

Finally, consider doubling and halving moves. These are compound moves that involve a change in volume and a change in particle number. Consider the volume-change submove first. This move is most easily defined as the limit of a more general kind of move. That is, consider the move consisting of expanding (or shrinking) an end slab of the simulation so that total volume is doubled (or halved). In this case, the end slab of length L'_x of the simulation box is expanded (or contracted) by a factor $\kappa=L_x/L'_x$. By taking the limit $L'_x \rightarrow 0$ it is possible to ensure that there are never any particles within this end slab. In the case of a doubling move the next submove consists of replicating the particles in the original volume according to the coordinate choices given in Eq. (7). In the case of a halving move this submove consists of deleting particles $N+1, \dots, 2N$. To ensure that the correct permutation of particle labels is preserved when performing these addition and deletion submoves we must be careful with calculating periodic images. First, before these submoves are performed, periodic images should be “unwound” so that particles that have swapped sides are moved back to their nonperiodic positions outside of the original simulation cell. Then the addition or deletion submoves are performed in accordance with Eq. (7). Finally, periodic images are re-applied in accordance with the new simulation cell boundaries. Note that this convention must also be followed when analyzing the positions of particles j and $j+N$ according to Eq. (8).

From Eqs. (6) and (9) we have for doubling moves

$$q_{sd}^{acc} = \frac{2V_s e^{-\beta_{d1}H_d} e^{-2\beta_{d1}P_{d1}V_s}}{2N_s} 2dV_s d\mathbf{r}^{2N_s-1} \\ = \frac{2N_s}{V_s e^{-\beta_s H_s} e^{-\beta_s P_s V_s}} dV_s d\mathbf{r}^{N_s-1} \left(\frac{d\mathbf{r}}{V_r} \right)^{N_s} \\ = 2V_r^{N_s} e^{-(\beta_{d1}H_d - \beta_s H_s)} e^{-V_s(2\beta_{d1}P_{d1} - \beta_s P_s)} \quad (15)$$

where $V_r = 4\pi(\alpha_i L_x)^3/3$. Because q_{ds}^{acc} is defined via Eq. (10) we find generally, depending on the replication tolerance, that one of the acceptance probabilities for these doubling and halving moves is very close to 1 while the other is very close to 0. So to improve statistics the double-size system is biased by an amount

$$\frac{\exp(\beta_s \mu_0 N_s)}{(\hat{V}_r)^{N_s}} \quad (16)$$

where $\hat{V}_r = 4\pi(\alpha_i \hat{L}_x)^3/3$ and \hat{L}_x is a fixed and typical value of L_x . This yields

$$\ln(q_{sd}^{acc}) = N_s[\beta_s \mu_0 + \ln(V_r/\hat{V}_r)] - (\beta_{d1}H_d - \beta_s H_s) \\ - V_s(2\beta_{d1}P_{d1} - \beta_s P_s) + \ln(2) \quad (17)$$

for doubling and halving moves. None of the other types of MC move are affected by this bias. The initial tolerance α_1 is chosen to ensure that the acceptance probability of a doubling move is not too small (i.e., $H_d \sim 2H_s$), while μ_0 is chosen to ensure equal sampling of single- and double-size systems.

B. Relaxation

Having transformed the single-size system to a constrained double-size system we now need to relax the constraint, measuring the free-energy difference between the constrained double-size system and the unconstrained double-size system. Potentially the most efficient way of achieving this is with quadrature, i.e., calculating the rate of change of free energy at many points distributed along the required path and then integrating along the path on the basis of an interpolation scheme. But in this work, which is focused on the feasibility of the SR approach, simple parameter hopping is employed instead, which calculates the free-energy difference between the two end points of a hop, and then connects very many hops together to traverse the required path.

For each hop a multicanonical simulation is performed that samples states consistent with α_i and α_{i+1} , i.e., each hop consists of a “dual-canonical” simulation. Relaxing and constraining moves are allowed that enable a system with N_d particles at pressure P_i , temperature T_i , and tolerance α_i to switch back and forth to a system at pressure P_{i+1} , temperature T_{i+1} , and tolerance α_{i+1} (note that the particles themselves are not moved during these switches). Relaxing moves are attempted if the double-size system is constrained while constraining moves are attempted (with equal probability) if the double-size system is relaxed. A path is traced thorough (P, T, α) space such that the initial point $(P_{d1}, T_{d1}, \alpha_1)$ is connected to the final point (P_s, T_s, α_m) via $m-2$ other points (corresponding to $m-1$ separate multicanonical simulations) where the final tolerance is sufficiently large that it no longer has any significant effect on the system, i.e., the free-energy difference between (P_s, T_s, α_{m-1}) and (P_s, T_s, α_m) is negligible. The path is easily defined; $\chi = \alpha_{i+1}/\alpha_i$ can be set to a constant, while a simple algorithm, described below, sets $\Delta P = P_{di+1} - P_{di}$ and $\Delta T = T_{di+1} - T_{di}$ to ensure that the average particle and energy densities of each system do not stray too far from those of the unconstrained system at P_s and T_s . In addition to these relaxing and constraining moves the same particle displacement and volume change moves as described above are also allowed.

The acceptance ratio for relaxing moves if $p_{i(i+1)}^{ch} = p_{i(i+1)i}^{ch}$ is

$$q_{i(i+1)}^{acc} = \left(\frac{V_{r,i}}{V_{r,i+1}} \right)^{N_s} e^{-H_d(\beta_{di+1} - \beta_i)} e^{-V_d(\beta_{di+1}P_{di+1} - \beta_i P_i)} \\ = \chi^{-3N_s} e^{-H_d(\beta_{di+1} - \beta_i)} e^{-V_d(\beta_{di+1}P_{di+1} - \beta_i P_i)}. \quad (18)$$

For constraining moves the acceptance ratio is $(q_{i(i+1)}^{acc})^{-1}$.

Finally, to improve statistics we bias the i th state by $\exp(\beta_{di} N_s \mu_i)$ so that relaxed and constrained states are sampled equally. The acceptance ratio for these moves is then

$$q_{i(i+1)}^{acc} = \chi^{-3N_s} e^{-H_d(\beta_{di+1} - \beta_i)} e^{-V_d(\beta_{di+1}P_{di+1} - \beta_i P_i)} e^{N_s(\beta_{di+1}\mu_{i+1} - \beta_{di}\mu_i)} \quad (19)$$

where $\mu_1 = 0$.

The total Gibbs-free-energy difference between the single- and unconstrained double-size systems is then

TABLE I. Simulation parameters and results for the SR method applied to the fcc hard sphere crystal. N_s is the number of particles in the single-size simulation and “Trials” is the total number of attempted MC moves—dividing by m gives the number for each relaxation simulation. χ is the tolerance ratio (see the text) in each relaxation simulation and $\mu = \mu_0 + \mu_m$ is the total chemical potential. The numbers in parentheses are statistical errors to one standard deviation.

Type	N_s	m	Trials (10^6)	χ	μ_0^*	μ_m^*	μ^*
Replication	108		50		1.2311(6)		
Relaxation	108	2500	10 000	1.00268		14.769(14)	16.000(14)
Replication	256		100		1.2367(7)		
Relaxation	256	5000	40 000	1.00134		14.811(12)	16.048(12)

$$\Delta G = N_s \{ \mu_0 + \mu_m - k_B T \ln[\hat{V}_r(\alpha_m)] \}. \quad (20)$$

There is no unique way of defining an algorithm for calculating the pressure and temperature steps ΔP and ΔT . In this work it is defined for both pressure and temperature as follows. First, target volumes and configurational energies are defined for the double-size system, $V_d = 2\langle V \rangle_s$ and $E_d = 2\langle E \rangle_s$, respectively, where the angular brackets denote an ensemble average. Then, P_{d1} and T_{d1} are chosen such that $\langle V_{d1} \rangle \cong V_d^t$ and $\langle E_{d1} \rangle \cong E_d^t$ and the indices k_V and k_E are set equal to 0. The average volume and energy are measured for each α_i . These indices are set to zero if $\langle V_{di} \rangle > V_d^t$ or $\langle E_{di} \rangle > E_d^t$, respectively, while they are increased by 1 (to a maximum of 3) if $\langle V_{di} \rangle < V_d^t$ or $\langle E_{di} \rangle < E_d^t$, respectively. For each α_{i+1} the pressure and temperature steps are defined as

$$\Delta P = \frac{(P_s - P_{di})}{m - i} \min(k_V, m - i),$$

$$\Delta T = \frac{(T_s - T_{di})}{m - i} \min(k_E, m - i). \quad (21)$$

III. RESULTS FOR SIMPLE SYSTEMS

The aim of this work is to test and validate the self-referential method described above. Simple model crystals are ideal for this purpose. Hard sphere and Lennard-Jones (LJ) face centered cubic crystals are chosen because of the availability of reference data. The hard sphere and shifted-force LJ potentials are

$$\phi_{HS}(r) = \begin{cases} \infty, & r < d, \\ 0, & r \geq d, \end{cases}$$

TABLE II. As for Table I except that the shifted-force LJ crystal is simulated.

Type	N_s	m	Trials (10^6)	χ	μ_0^*	μ_m^*	μ^*
Replication	256		100		-0.2939(3)		
Relaxation	256	2000	800	1.00225		-2.910(22)	-3.204(22)

$$\phi_{LJ}^{sf}(r) = \begin{cases} 4\epsilon(x^{-12} - x^{-6}) - \phi_{LJ}(r_c) - (r - r_c)\phi'_{LJ}(r_c), & r < r_c, \\ 0, & r \geq r_c, \end{cases} \quad (22)$$

respectively, where $x = r/\sigma$, r is the pair separation, the prime indicates the derivative with respect to separation, and $\phi_{LJ} = 4\epsilon(x^{-12} - x^{-6})$.

The perfect (defect-free) fcc hard sphere crystal is simulated at fluid-solid coexistence, i.e., at a reduced pressure [23] of $P_s^* = P_s d^3 / k_B T = 11.49$, where d is the hard sphere diameter (note that in Ref. [23] the reduced coexistence pressure is calculated with an uncertainty of ± 0.09). Temperature is arbitrary for this system. Simulations with $N_s = 108$ and 256 are performed to examine finite-size effects, which are expected to vary with system size with a leading-order contribution proportional to $1/N_d$. Statistical errors are calculated in the usual way [24] using block averages. Regarding the probabilities for choosing each trial move, i.e., $p_{on}^{ch} = p_{no}^{ch}$, they are in the ratio $1:N_s^{-1}:N_s^{-1}$ for displacement, volume, and doubling and halving moves during the replication stage, and $1:N_d^{-1}:1$ for displacement, volume, and relaxation or constraining moves during the relaxation stage.

Simulation parameters and results are given in Table I. I choose $P_{d1}^* = 6.375$ and $\alpha_1 = 0.001/\hat{L}_x$. The final result is $\mu^* = \mu/k_B T = 16.08 \pm 0.02$. These results confirm (within statistical error) the validity of the result in Ref. [23] for hard sphere fluid–perfect fcc solid coexistence because the (very accurate) equation of state of Kolafa *et al.* [25] for a hard sphere fluid gives $\mu^* = 15.99 \pm 0.10$ at $P^* = 11.49 \pm 0.09$.

For the perfect fcc shifted-force LJ crystal at its triple point, i.e., at a reduced pressure and temperature of $P^* = P\sigma^3/\epsilon = 0.00182$ and $T^* = k_B T/\epsilon = 0.56$, one simulation with $N_s = 256$ is performed. Simulation parameters and results are given in Table II. I choose $P_{d1}^* = 0$, $T_{d1}^* = 0.7616$, and $\alpha_1 = 0.01/\hat{L}_x$. The final result is $\mu^* = \mu/\epsilon = -3.20 \pm 0.02$, al-

though this does not include an analysis of finite-size effects because of the numerical expense involved. These results agree well with those of Errington and co-workers [26] and Sweatman and Quirke [27], which together give $\mu^* = -3.23 \pm 0.02$ (the gas phase result from a gas-liquid Gibbs simulation [27]).

IV. SUMMARY

The SR methodology described above is appropriate for simple atomic systems that form bulk crystalline solids with cubic unit cells. But this method is potentially very general, and is easily generalized to treat molecular crystals formed from noncubic unit cells. When treating molecular crystals it will be important to constrain internal molecular degrees of freedom in an appropriate manner during replication. For frozen internal molecular degrees of freedom (such as the orientation of water molecules in ice) a constraint is needed in much the same way as it is needed for molecular position coordinates. These types of constraint will increase the path length substantially. For nonfrozen molecular degrees of freedom (such as the orientation of nitrogen in its orientationally disordered crystalline β phase [28]) the constraint can be much weaker and should not contribute significantly to the numerical effort required to traverse the path. Constrained molecular degrees of freedom need to be treated in the same way as position coordinates, i.e., constrained on system doubling and then gradually relaxed. The same methodology can be employed to treat confined crystals with at least one degree of translational invariance, such as crystals confined in rigid slit pores. In this case pressure should be replaced by grand-potential density (or interfacial tension) and simulations are performed in the $N\sigma T$ ensemble by allowing fluctuations in simulation box vertices parallel to the translational invariance. Further work is required to investigate the possibility of including defects in the free-energy calculation.

The approach described above is not yet efficient. The overwhelming amount of CPU time is spent performing the relaxation stage. For example, for the double-size Lennard-Jones system the 5000 connected hops that traverse the relaxation path required about six days of computation on a standard desktop (3.0 GHz) PC. There are two reasons for this. First, the free-energy difference between the single- and double-sized systems is large, resulting in a long path length. Second, the parameter hopping approach, although robust, is not efficient. Future work will aim to improve this aspect and

make comparison with other methods. The most obvious idea for improving efficiency is to implement thermodynamic integration (perhaps along the lines described in Ref. [14]), rather than parameter hopping, as the technique used to relax the double-size system from its initially highly constrained state to its fully relaxed state. This could potentially reduce the number of individual relaxation simulations from several thousand to less than 100, consequently improving efficiency by perhaps 1 to 2 orders of magnitude. Another possibility is to attempt to adapt Tilwani and Wu's ingenious retiling algorithm [29], which has been developed for 2D disk systems, to general 3D crystalline solids. This would then allow the large system to be only slightly larger than the small system, say one unit cell larger for a bulk crystal rather than double size, resulting in a much shorter path between initial and final states. Combination of both these efficiency enhancements could perhaps yield a methodology several orders of magnitude faster than the current implementation for typical simulations.

The SR approach developed above can be considered a reinvention of the SR approach described by Barnes and Kofke [12], although there are some significant differences. In their work they considered for convenience a system of 1D hard rods on a line, which they argue can be considered a primitive form of crystal in a limited sense. They cast their canonical partition function in terms of vibrational modes, and use free-energy perturbation to traverse the path between single-size and double-size systems. In contrast, this work employs parameter hopping to traverse between the two subsystems—a technique that they also suggest as being worthy of investigation as it has the advantage of being reversible (see their work for a discussion of this). They also work in the canonical ensemble rather than the isothermal-isobaric ensemble used here, potentially leading to additional ensemble-induced finite-size effects (although their system was sufficiently large that these effects were not significant). They conclude with some suggestions for further work concerning the inclusion of temperature and higher dimensions so that more realistic 2D and 3D crystals can be treated. These concerns are effectively answered here. However, this work demonstrates that both the temperature and pressure need to be manipulated during the SR process to ensure that it is feasible.

ACKNOWLEDGMENTS

I sincerely thank Nick Quirke for advice with this work and his tireless support and encouragement over the years.

-
- [1] J. F. Lutsko, D. Wolf, and S. Yip, *J. Chem. Phys.* **88**, 6525 (1988).
[2] B. Smit and D. Frenkel, *Understanding Molecular Simulation: From Algorithms to Applications* (Academic, New York, 1996).
[3] D. Frenkel and A. J. C. Ladd, *J. Chem. Phys.* **81**, 3188 (1984).
[4] A. D. Bruce, N. B. Wilding, and G. J. Ackland, *Phys. Rev. Lett.* **79**, 3002 (1997).
[5] R. M. Lynden-Bell, J. S. van Duijneveldt, and D. Frenkel, *Mol. Phys.* **80**, 801 (1993).
[6] G. Grochola, *J. Chem. Phys.* **120**, 2122 (2004).
[7] A. P. Lyubartsev *et al.*, *J. Chem. Phys.* **96**, 1776 (1992).
[8] P. A. Monson and D. A. Kofke, *Adv. Chem. Phys.* **115**, 113 (2000).

- [9] J. M. Rickman and R. LeSar, *Annu. Rev. Mater. Res.* **32**, 195 (2002).
- [10] S. Y. Sheu, C. Y. Mou, and R. Lovett, *Phys. Rev. E* **51**, R3795 (1995).
- [11] N. B. Wilding and A. D. Bruce, *Phys. Rev. Lett.* **85**, 5138 (2000).
- [12] C. D. Barnes and D. A. Kofke, *Phys. Rev. E* **65**, 036709 (2002).
- [13] K. K. Mon, *Phys. Rev. Lett.* **54**, 2671 (1985).
- [14] K. K. Mon and K. Binder, *Phys. Rev. B* **42**, 675 (1990).
- [15] P. D. Beale, *Phys. Rev. E* **66**, 036132 (2002).
- [16] R. Radhakrishnan and K. E. Gubbins, *Mol. Phys.* **96**, 1249 (1999).
- [17] R. Radhakrishnan, K. E. Gubbins, and M. Sliwinska-Bartkowiak, *J. Chem. Phys.* **116**, 1147 (2002).
- [18] K. G. Ayappa and C. Ghatak, *J. Chem. Phys.* **117**, 5373 (2002).
- [19] M. Miyahara and K. E. Gubbins, *J. Chem. Phys.* **106**, 2865 (1997).
- [20] H. Dominguez, M. P. Allen, and R. Evans, *Mol. Phys.* **96**, 209 (1999).
- [21] J. M. Polson *et al.*, *J. Chem. Phys.* **112**, 5339 (2000).
- [22] L. A. Baez and P. Clancy, *Mol. Phys.* **86**, 385 (1995).
- [23] N. B. Wilding, *Comput. Phys. Commun.* **146**, 99 (2002).
- [24] M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids* (Clarendon Press, Oxford, 1987).
- [25] J. Kolafa, S. Labik, and A. Malijevsky, *Phys. Chem. Chem. Phys.* **6**, 2335 (2004).
- [26] J. R. Errington, P. G. Debenedetti, and S. Torquato, *J. Chem. Phys.* **118**, 2256 (2003).
- [27] M. B. Sweatman and N. Quirke, *Mol. Simul.* **30**, 23 (2004).
- [28] E. J. Meijer *et al.*, *J. Chem. Phys.* **92**, 7570 (1990).
- [29] P. Tilwani and D. Wu, *Direct Simulation of Phase Coexistence in Solids using the Gibbs Ensemble: Configuration Annealing Monte Carlo* (Colorado School of Mines, Golden, 2000).