Lattice-switch Monte Carlo method

A. D. Bruce, A. N. Jackson, G. J. Ackland, and N. B. Wilding

Department of Physics and Astronomy, The University of Edinburgh, Edinburgh EH9 3JZ, Scotland, United Kingdom

We present a Monte Carlo method for the direct evaluation of the difference between the free energies of two crystal structures. The method is built on a lattice-switch transformation that maps a configuration of one structure onto a candidate configuration of the other by ''switching'' one set of lattice vectors for the other, while keeping the displacements with respect to the lattice sites constant. The sampling of the displacement configurations is biased, multicanonically, to favor paths leading to *gateway* arrangements for which the Monte Carlo switch to the candidate configuration will be accepted. The configurations of both structures can then be efficiently sampled in a single process, and the difference between their free energies evaluated from their measured probabilities. We explore and exploit the method in the context of extensive studies of systems of hard spheres. We show that the efficiency of the method is controlled by the extent to which the switch conserves correlated microstructure. We also show how, microscopically, the procedure works: the system finds gateway arrangements which fulfill the sampling bias intelligently. We establish, with high precision, the differences between the free energies of the two close packed structures (fcc and hcp) in both the constant density and the constant pressure ensembles.

PACS number(s): 05.10.Ln, 65.50.+m, 64.70.Kb

I. INTRODUCTION

Let us pose the problem. We are presented with a material whose chemical composition is known; we are provided with a model of the interatomic interactions; and we have identified two candidate crystalline structures. How should we proceed to determine which structure will be favored under given conditions? Of course, equilibrium statistical mechanics tells us what we must do, in principle: the favored structure will be that which has the *greater a priori probability* or *configurational weight* or, in equivalent thermodynamic parlance, *lower free energy*. Thus the task is to *compare* the configurational weights of (determine the *difference* between the free energies of) the candidate structures.

A variety of approximate strategies exist for addressing this problem [1]. But it is clear that if one desires a technique that is both generally applicable and reliable (that is, has quantifiable uncertainties) one must look to the Monte Carlo (MC) method [2], the standard computational tool for dealing with many-body systems [3].

The application of MC methods to the study of phase-behavior presents a generic problem [4,5]: the free energy of a phase cannot be expressed (in a practically useful form) as a canonical average over the associated configurations; free-energy-estimation inevitably entails simulations that visit a substantially wider spectrum of configurations, which together form a *path* through configuration space [6]. The strategic choices to be made concern the path itself—ultimately, the physical character of the additional configurations sampled—and the sampling procedure.

An acceptable path will fall into one or other of two categories—we will call them reference-state and interphase paths. A reference-state path links (comprises sets of configurations that interpolate between) the configuration space associated with each phase to the configuration space associated with some reference system [7] whose free energy is known. An interphase path links the configuration space of

one phase to that of the other. Both categories of path embrace many further subcategories. Thus, a reference-state-path may run through a space of thermodynamic coordinates or through a space of model parameters. An interphase path may be ''physically motivated''—modeling authentically the configurations through which a system actually passes in the course of a phase transformation or it may be ''computationally motivated'' (''nonphysical'')—designed, pragmatically, to deliver a result.

The sampling procedures used to explore the chosen path also fall, broadly, into one or other of two categories – we will call them multistage and single stage. The multistage approach entails a *series* of independent simulations each of which explores a different point on the path; the simulations may determine simply the derivative of the free energy at each point [the integration method (IM)] or the difference between the free energies of adjacent points (the overlap method). The single-stage approach involves, in essence *one* simulation exploring the *entire* path.

There are very many ways in which one can respond to these strategic choices. Many of them are represented in the large literature devoted to this problem [8–14]. But all of them, in our view, lack one or more of the characteristics (generality, transparency, precision) of a definitively satisfactory solution to such a fundamental and simply posed problem.

In seeking that solution it seems to us there are good *a priori* grounds for favoring an inter phase path, explored by single-stage sampling. The prejudice on the choice of path reflects the fact that, in using a reference state path, one has to determine, separately, the absolute free energies of each phase. These absolute free energies are typically very large—*arbitrarily* so in the vicinity of a phase boundary—compared to the quantity (their difference) which is actually of interest. In contrast, using an interphase path allows one to focus directly on this quantity. The *a priori* preference for a single-stage sampling rests on the transparency with which
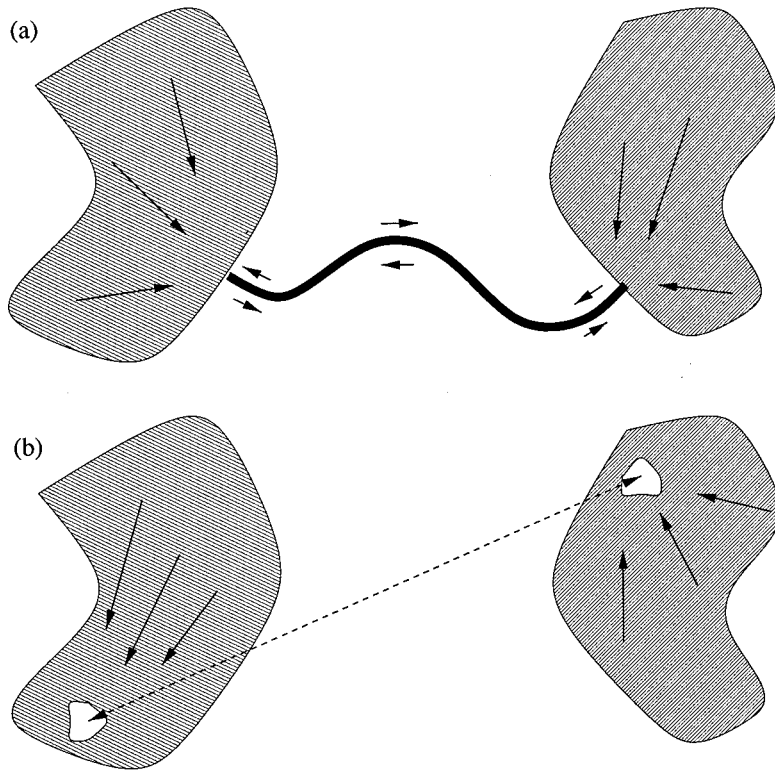
(a)

(b)

FIG. 1. Schematic representations of the different ways in which multicanonical sampling methods can be used to achieve interphase crossing. In the conventional approach (a) the sampling algorithm is biased so as to enhance the probability of the *mixed* phase states lying along a path (the heavy dark line) linking the two regions of configuration space. In the lattice-switch method (b) the bias is constructed so as to enhance the probability of the subsets of states (the white islands), *within* the single-phase regions, from which the switch operation (the large dashed arrow) will be accepted.

the associated uncertainties (error bounds) are prescribed. We shall return to these points in Sec. V. With these strategic choices made, one is left with two tasks—one conceptual (designing the interphase path) and the other practical (formulating the sampling algorithm).

The practical issue is relatively easily addressed. In recent years, the Monte Carlo toolkit has been significantly enhanced to provide a range of extended sampling techniques—multicanonical [15], expanded ensemble [16], and simulated tempering [17]. These methods (whose origins can be traced back to much earlier pioneering work [18]) allow one to construct a MC procedure that will traverse virtually any desired path through configuration space. Here we adopt the multicanonical framework. In this framework, the desired path is represented as a discrete series of macrostates, defined by some chosen macroscopic property [6]; in the multicanonical sampling procedure each macrostate is visited with a probability that is enhanced, or diminished, with respect to its canonical value, by an amount that is controlled by a *multicanonical weight*; the set of weights is constructed so that, while the canonical probabilities vary vastly over the path, the multicanonical probabilities are essentially constant, allowing the whole path to be negotiated in one simulation.

The core issue is, then, the *design* of the interphase path—at heart, the choice of an appropriate order parameter [19]. The choice is important: it determines, implicitly [20], the nature of the configurations sampled during the interphase traverse, the MC time required for that traverse, and thence the statistical quality of the final results.

Outside the context of *structural*-phase behavior—in the case of liquid-gas phase behavior, for example—the choice is clear and a multicanonical strategy is securely in place. The order parameter is identified with that—the density—associated with the accompanying critical point. The result-

ing interphase configurations are then generically *inhomogeneous*, comprising two coexisting regions (one of each phase), separated by an interface. On this path, it is the free energy cost of this interface that provides the ergodic barrier which has to be surmounted by multicanonical weighting [21]. The passage along the path (the motion of the interface) involves processes which differ only in scale from those already represented in the microscopic dynamics of a single phase. This approach is illustrated schematically in Fig. 1(a). It has been successfully used in studies of phase behavior in ferromagnets [22], fluids [23], and lattice gauge theories [24].

In the context of *structural* phase behavior it is clear that this kind of strategy will seldom be fruitful [25,26]. In such systems a traverse through an inhomogeneous two-phase (necessarily noncrystalline) region will involve substantial, physically slow, restructuring—vulnerable to further ergodic traps, and compounding the intrinsic slowness of the multicanonical sampling process. To the two general *a priori* preferences expressed above we thus add a third, specific to the structural context: the interphase path should comprise macrostates that are single phase and crystalline. This paper shows how to identify, build, and exploit a path of this kind.

The key ideas are simple. In any crystalline configuration each atomic position coordinate may be expressed as the sum of a lattice vector and a displacement vector. The configuration space associated with each structure, individually, may be explored by standard MC procedures which stochastically update the displacement vectors while keeping the lattice vectors constant. In *principle* the passage from one phase to the other may be accomplished by a lattice switch (LS) in which one entire set of lattice vectors is replaced with the other, while the displacement vectors are held fixed. Formally this LS can be incorporated into the MC procedure simply by treating the lattice type as an additional stochastic

variable. In practice *this* interphase ''path'' (blind leap) will not work. Implemented this way the LS will map a ''typical'' configuration of one structure onto an ''untypical'' (high-energy) conjugate configuration [27]; the associated MC step will generally be rejected. To make it work the LS needs to be extended to include two segments of ''path'' (each lying entirely within one phase) which connect the sets of equilibrium configurations with the special configurations (we will call them gateway configurations) from which a successful LS can be initiated [28]. These path-segments may be labeled by an ''order parameter'' which measures the mismatch between the energies of the configurations linked by LS. This order parameter has a high value for the equilibrium configurations, lying at one end of a path segment: these configurations are not energy matched [29] to their conjugates. It has a low value for the gateway configurations at the other end: gateway configurations (whatever other attributes they may have) are necessarily energy matched to their conjugates. Multicanonical weights are attached to the macrostates of this order parameter, so that the multicanonical sampling procedure explores both path segments evenly, surmounting the probabilistic barrier which, in this case, reflects the *smallness* of the statistical weight of the gateway configurations. Together, the multicanonical sampling and the lattice switch provide a configuration space ''look and leap'' [Fig. 1(b)] which visits both phases while remaining at all times crystalline.

The LS method was introduced by us and described in outline form in an earlier brief communication [30]. Since that time it has been applied by two other groups [14,31]. The present paper has three principal objectives.

The first objective (with which we have already engaged in the preceding discussion) is to explain the core idea more fully: the ''idea'' (biased sampling to facilitate a global coordinate change) represents, we believe, a significantly new form of extended sampling, which merits further exposure.

Our second objective is to achieve a deeper understanding of how the process works—in particular the implications of the *form* chosen for the LS operation adopted (it is not unique) and the microscopic character of the gateway configurations which the system locates in response to the multicanonical weighting, tailored to support that operation. We show that the efficiency of the LS operation depends significantly on the extent to which it conserves correlated microstructure. And we find that the gateway configurations have features which reflect the specific nature of the lattice-switch transformation we adopt, in a microscopically intelligible (even intelligent) way.

Our third objective is to extend our study of the phase behavior of hard spheres. This problem is of enduring interest, displaying a richness that belies the simplicity of the model itself [32]. The relative stability of the two closed-packed (fcc and hcp) structures is particularly finely balanced: the entropy difference [33] is so small (smaller than the entropy change at freezing by of order 10^{-3}) that it can easily be lost in statistical uncertainties. Discrepancies (large, in relative terms) between a recent IM study [10] and its predecessors [9] provided the motivation for our development of the LS method. In this paper we present results in the constant-density ensemble, both near the melting density and at the close-packed limit. In so doing we resolve the

discrepancy between the results, near melting, reported in our initial study [30] and those—also using LS—reported recently by Pronk and Frenkel [31]: the fault was ours, stemming from a failure to recognize the consequences of center-of-mass drift. We also show that the method can be extended straightforwardly—in this case at least—to the constant-pressure ensemble.

The paper is structured as follows. Section II sets out the theoretical framework. We define the model, the competing structures, and the associated configurational weights: in the case of hard sphere systems the latter are purely entropic. We identify an appropriate form of lattice-switch transformation: here, it is designed to capitalize on the close-packed layers common to both structures. To bias the displacement sampling we need to define an appropriate measure of the ''energy cost'' of the lattice switch; we will see that the number of pairs of overlapping spheres created by the transformation fulfills this role simply and effectively. The efficiency of the method also potentially depends on the choice of representation of both the lattice-to-lattice mapping and the particle displacements: we discuss the principles involved in the choice of representation. Section III provides computational implementation details, including the procedures used to evolve an appropriate multicanonical sampling distribution. Section IV contains our results. Finally, in Sec. V, we offer our conclusions in relation to both the hard sphere system and the lattice-switch method.

II. FORMULATION

A. The model system

We consider a system of N particles, of spatial coordinates $\{\vec{r}\}$, confined within a volume V, and subject to periodic boundary conditions. The interactions are those of hard spheres of diameter D; the configurational energy is of the form

$$E(\{\vec{r}\}) = \begin{cases} 0 & \text{if } r_{ij} \geq D \ \forall i,j, \\ \infty & \text{otherwise,} \end{cases} \quad (1)$$

where $r_{ij} = |\vec{r}_i - \vec{r}_j|$. The total configurational weight associated with this system is

$$\Omega(N,V) = \prod_i \left[\int_V d\vec{r}_i \right] \prod_{\langle ij \rangle} \Theta(r_{ij} - D), \quad (2)$$

where $\Theta(x) \equiv 1(0)$ for $x \geq 0 (<0)$, and the product on $\langle ij \rangle$ extends over all particle pairs. The associated entropy density is

$$s(N,V) \equiv \frac{1}{N} \ln \Omega(N,V). \quad (3)$$

We are concerned with the entropy of specific phases (the two familiar crystalline close-packed structures) of this system. In general, the entropy of a phase measures the weight of the configurations satisfying some constraint that is characteristic of that phase. It is therefore necessary in principle (although in practice the issue is typically skirted) to formulate a constraint that identifies a configuration as ''belonging to'' a given crystalline phase. One can do so—very naturally,

and in the traditions of lattice dynamics [34]—by decomposing the particle position coordinates into a sum of ''lattice'' and ''displacement'' vectors:

$$\vec{r}_i = \vec{R}_i^\alpha + \vec{u}_i. \tag{4}$$

Here $\{\vec{R}\}_\alpha \equiv \vec{R}_i^\alpha, i = 1, \ldots, N$ is a set of fixed (configuration-independent) vectors associated with the crystalline structure labeled α. We will refer to them as ''lattice vectors.'' But we use this term a little loosely: more precisely, we mean the set of vectors identified by the orthodox crystallographic lattice, convolved with the orthodox basis [35,36]. The other vectors $\{\vec{u}\}$ represent displacements with respect to the ''lattice'' sites; the symmetry of the structures of interest here ensures that these displacements have zero ensemble average. This framework provides us with a number of ways of identifying the configurations to be associated with structure α. First, one might adopt the criterion that all particle displacements, with respect to the associated lattice sites, lie within some nominated *spatial* cutoff, chosen to be sufficiently large that the results are independent of its specific value. This criterion has the merit that it does not stray beyond the concepts of equilibrium statistical mechanics. Alternatively one might identify the relevant configurations as the set that are accessible from *some* member of the set (the perfect crystalline state, for example) within some nominated *temporal* cutoff. The merit of *this* choice is that it is a quasiformal expression of what, in practice, computer simulation attacks on this problem actually *do*, albeit implicitly: the free energy assigned to a phase (in, for example, IM-based studies) represents the weight of the configuration space sampled on the time scale of the simulation. The result should be independent of that time scale provided it (the scale) is long enough that the configuration space of each structure is effectively sampled, but still short compared to interphase crossing times. Whichever view one takes (in practice we adopt the latter: see Sec. III A) one may write, for the configurational weight associated with structure α

$$\Omega(N, V, \alpha) = \prod_i \left[\int_\alpha d\vec{u}_i \right] \prod_{\langle ij \rangle} \Theta(r_{ij} - D), \tag{5}$$

where $\int_\alpha$ signifies integration subject to the chosen configurational constraint.

In the thermodynamic ($N \to \infty$) limit, the associated entropy density

$$s(N, V, \alpha) \equiv \frac{1}{N} \ln \Omega(N, V, \alpha) \tag{6}$$

depends only on the particle number density, which we write in the dimensionless form

$$\tilde{\rho} \equiv \frac{\rho}{\rho_{CP}} \equiv \frac{N/V}{\sqrt{2}/D^3}, \tag{7}$$

where ρ_{CP}, the number density at close packing, provides the natural scale. The range of interest to us here extends from the melting density $\tilde{\rho} \approx 0.736$ [37] through to the close-packed limit $\tilde{\rho} = 1$.

In the close-packed limit the configurational integral [Eq. (5)] may be rewritten [38] as the product of two terms:

$$\Omega(N, V, \alpha) = \Omega_0(N, V) \Omega_\alpha \tag{8}$$

The first term here is defined by

$$\Omega_0(N, V) = \left[\frac{D\epsilon}{1 - \epsilon} \right]^{3N} \tag{9a}$$

with

$$\epsilon \equiv 1 - \tilde{\rho}^{1/3}. \tag{9b}$$

The associated contribution to the entropy is logarithmically divergent in the close-packed limit [39], but independent of the phase. The second contribution to the configurational integral is defined by

$$\Omega_\alpha = \prod_i \left[\int_\alpha d\vec{u}_i \right] \prod_{\langle ij \rangle}^{nn} \Theta(u_{ij}^{\|} + 1) [1 + O(\epsilon)], \tag{10a}$$

where [40]

$$\vec{u}_{ij} \equiv \vec{u}_i - \vec{u}_j \equiv u_{ij}^{\|} \hat{n}_{ij}^\alpha + \vec{u}_{ij}^\perp \tag{10b}$$

while $\hat{n}_{ij}^\alpha$ is a unit vector from lattice site j to nearest neighbor lattice site i. The associated contribution to the entropy is finite, but depends on the phase through the geometry of the nearest neighbor vectors. It may be visualized as that of a set of hard dodecahedra [41].

Now let us recall that the quantity of immediate interest is the *difference* between the entropy densities of the two phases. It may be written as

$$\Delta s_{\alpha\beta} \equiv s(N, V, \alpha) - s(N, V, \beta) = \frac{1}{N} \ln \mathcal{R}_{\alpha\beta}(N, V), \tag{11}$$

where

$$\mathcal{R}_{\alpha\beta}(N, V) \equiv \frac{\Omega(N, V, \alpha)}{\Omega(N, V, \beta)} = \frac{P(\alpha|N, V)}{P(\beta|N, V)}. \tag{12}$$

Here $P(\alpha|N, V)$ is the probability that a system, free to explore the joint configuration space of the two structures (and visiting configurations with the appropriate probabilities—all equal in this case) will be found in some configuration of structure α.

In the constant density ensemble, then, the computational task is to determine the ratio defined by Eq. (12). In the constant pressure ensemble we require the ratio $\mathcal{R}_{\alpha\beta}(N, P^\star)$ of the partition functions

$$\mathcal{Z}(N, P^\star, \alpha) = \int dV \Omega(N, V, \alpha) e^{-P^\star V}, \tag{13}$$

where $P^\star$ is a measure of the pressure [42]. The associated thermodynamic potential is the Gibbs free energy density defined by

$$g(N, P^*, \alpha) \equiv -\frac{1}{N} \ln \mathcal{Z}(N, P^\star, \alpha) \tag{14}$$
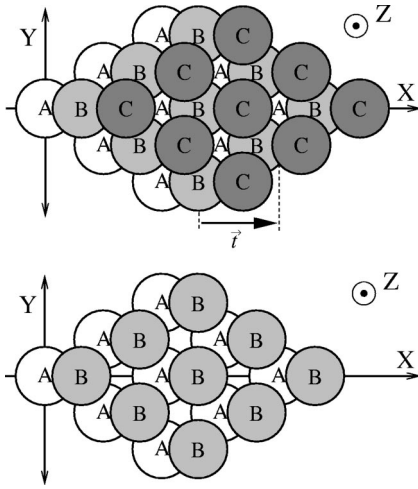
FIG. 2. Schematic representations of the two close-packed structures. The structures differ only in regard to the stacking pattern of the close-packed (x-y) planes which are of the form $ABCABC\cdots$ for fcc (upper) and $ABAB\cdots$ for hcp (lower). The vector labeled $\vec{t}$ is instrumental in defining the LS operation, shown in Fig. 3.

so that, in analogy with Eq. (11),

$$\Delta g_{\alpha\beta} \equiv g(N,P^\star,\alpha) - g(N,P^\star,\beta) = -\frac{1}{N}\ln \mathcal{R}_{\alpha\beta}(N,P^\star). \tag{15}$$

B. The lattice-switch method

The two close-packed structures of interest here are shown schematically in Fig. 2. In principle there are many transformations which will map one set of lattice vectors into the other; we shall consider the criteria guiding the choice in Sec. II C. The mapping used in most of the work reported here is shown schematically in Fig. 3. This scheme exploits the fact that the two structures differ only in respect of the stacking pattern of the close-packed planes. A suitable transformation can then be constructed that entails, simply, *translating* appropriate close-packed planes. By ''translate'' we mean, more precisely, ''relocate at a position defined by an
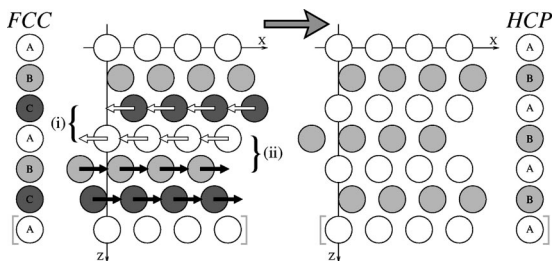

FIG. 3. The LS transformation applied to the perfect-crystal configuration. The diagram shows 6 close-packed (x-y) layers. [The additional bracketed layer at the bottom is the periodic image of the layer at the top.] The circles show the boundaries of hard spheres located at the sites of the two close-packed structures. In this realization of the fcc$\rightarrow$hcp lattice switch, the top pair of planes are left unaltered, while the other pairs of planes are relocated by translations, specified by the vectors $-\vec{t}$ (white arrows) and $\vec{t}$ (black arrows). The vector $\vec{t}$ is identified in Fig. 2.
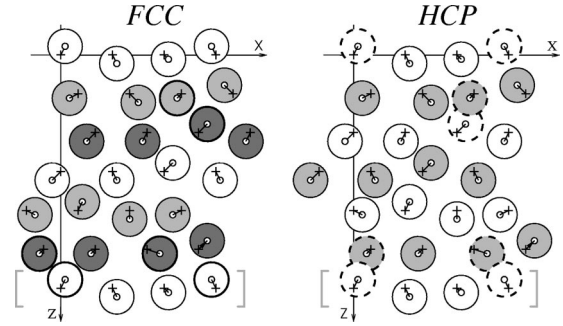

FIG. 4. The LS transformation applied to a ''typical'' configuration. The crosses identify the ''lattice sites''; the small circles locate the sphere centers in this configuration of displacements $\{\vec{u}\}$. This configuration is realizable (gives no overlaps) in the fcc structure; under the LS transformation it is mapped onto an (unrealizable) hcp configuration with three overlapping pairs of hard spheres (shown with dashed boundaries). Thus, for this configuration, the overlap order parameter $\mathcal{M}(\{\vec{u}\}) = 3$ [Eq. (16)].

appropriate translation vector'': one should not think of the planes as ''sliding through'' the intermediate states.

Figure 3 shows the application only to the *perfect*-crystal configurations where energy-matching is guaranteed [43]. In general (that is, for ''typical'' configurations: see Fig. 4 for an example) the two configurations related by the LS operation will not be energy matched: since adjacent planes are translated differently, the translations may—indeed, with overwhelming probability *will*—map a realizable configuration (of one structure), in which there are no overlapping spheres, onto an unrealizable configuration (of the other) in which there are overlaps. A MC lattice switch ''move'' will be rejected for *most* configurations. But not quite all: gateway configurations (configurations that are energy matched [29] to their conjugates) must exist, in significant measure. In particular, it is clear on grounds of continuity that configurations ''close enough'' to perfect-crystal form must fall into this category. One might therefore *choose* these ''small-displacement'' configurations to act as the gateway states, and define a multicanonical weighting procedure accordingly. However, one can avoid having to make this explicit choice, and, instead, let the system *find* gateway configurations itself. To do so we must define a measure of the mismatch between the energies of the configurations linked by the transformation.

In the present context that mismatch is quantified by the number of pairs of overlapping spheres created by the transformation. To that end let $M(\{\vec{u}\},\alpha)$ denote the number of overlapping pairs associated with the displacements $\{\vec{u}\}$ within the structure α. And define [44]

$$\mathcal{M}(\{\vec{u}\}) \equiv M(\{\vec{u}\},\text{hcp}) - M(\{\vec{u}\},\text{fcc}). \tag{16}$$

Since $M(\{\vec{u}\},\alpha)$ will necessarily be zero for any realizable set of displacements of structure α, the *overlap order parameter* $\mathcal{M}$ is necessarily $\geq 0 (\leq 0)$ for realizable configurations of the fcc (hcp) structure. Figure 4 provides a concrete example. The gateway configurations may then be identified (*without* prejudging their microscopic character) as the set of configurations for which $\mathcal{M} = 0$: a displacement pattern $\{\vec{u}\}$

for which $\mathcal{M}=0$ is realizable in *both* structures (energy matched). A LS MC step initiated from an $\mathcal{M}=0$ configuration *will* be accepted; if initiated from outside this set of configurations it will be rejected.

The sampling algorithm must thus be multicanonically customized so as to enhance the probability along a notional line in $\mathcal{M}$ space, extending from the ''equilibrium'' $\mathcal{M}$ values (reflecting the number of overlaps created by a LS acting on a *typical* configuration) through to the $\mathcal{M}=0$ gateway configurations. This aim is realized by augmenting the system energy function Eq. (1):

$$E(\{\vec{r}\}) \rightarrow E(\{\vec{r}\}) + \eta(\mathcal{M}(\{\vec{u}\})) \equiv \tilde{E}(\{\vec{r}\}) \qquad (17)$$

where $\eta(\mathcal{M}), \mathcal{M}=0, \pm 1, \pm 2 \cdots$ constitute a set of multicanonical weights [15]. These weights need to be chosen so as to allow the system to access the $\mathcal{M}=0$ gateway configurations, and thence (through the LS) the full joint configuration space of the two structures. The *desired* ratio of configurational weights, which reflects the *canonical* distribution $P(\mathcal{M}|N,V)$ [Eq. (12)] may then be estimated from the *measured multicanonical* distribution, $P(\mathcal{M}|N,V,\{\eta(\mathcal{M})\})$ with the identification

$$\mathcal{R}_{\text{fcc,hcp}}(N,V) = \frac{\displaystyle\sum_{\mathcal{M}>0} P(\mathcal{M}|N,V)}{\displaystyle\sum_{\mathcal{M}<0} P(\mathcal{M}|N,V)}$$

$$= \frac{\displaystyle\sum_{\mathcal{M}>0} P(\mathcal{M}|N,V,\{\eta(\mathcal{M})\}) e^{\eta(\mathcal{M})}}{\displaystyle\sum_{\mathcal{M}<0} P(\mathcal{M}|N,V,\{\eta(\mathcal{M})\}) e^{\eta(\mathcal{M})}}.$$

$$(18)$$

Here the exponential reweighting of the multicanonical distribution folds out the bias associated with the weights, whose residual effects are then simply as desired—the removal of the ergodic barrier between the two branches of the distribution.

C. Representations: tuning the lattice switch

We have presented the LS method in its simplest realization—the one we have used for most of the studies reported here. We now outline two important respects [45] in which some degree of generalization is possible, and may be desirable, in subsequent applications. Both involve the choice of *representation* of the LS transformation.

We have already alluded to the first point: there are many forms of lattice-to-lattice mapping. It is clear that the efficiency of the method will depend significantly upon the mapping chosen. Evidently the choice should be made so as to match up, as closely as possible, the *energy* of the two configurations it links. In the context of hard spheres this aim is realized by choosing the mapping which gives the smallest equilibrium overlap count (mean $|\mathcal{M}|$ value), which gives a measure of the entropic barrier that has to be negotiated by the multicanonical procedure. The smaller this barrier, the shorter is the path to the gateway configurations from which

a successful LS may be launched. Since the multicanonical simulations traverse this path only slowly (essentially diffusively, at best) the gains here are potentially substantial. It is intuitively clear that the scheme described above will fulfill this criterion well: in this representation, the LS translates close-packed planes bodily, so it can create overlaps only between spheres associated with *different* planes. But it is useful to explore other schemes—partly to check that there is no significantly better alternative, but principally to understand the different factors that control the efficiency. We have done so; the results are to be found in Sec. IV A.

There is a second—less obvious—generalization of the framework. In the simple realization, the particle positions are written in the ''lattice plus displacement'' representation provided by Eq. (4). The LS operation then maps a configuration of one structure onto a configuration of the other with the *same* set of displacements. This is unnecessarily restrictive. More generally we are at liberty to write, in place of Eq. (4),

$$\vec{r} = \vec{R}^{\alpha} + \mathbf{T}^{\alpha} \cdot \vec{u}, \qquad (19)$$

where $\vec{r}, \vec{R}^{\alpha}$, and $\vec{u}$ are now column vectors with $3N$ elements and $\mathbf{T}^{\alpha}$ is a $3N \times 3N$ non-singular matrix, whose form (possibly $\{u\}$-dependent) is at our disposal. Equation (5) is then replaced by

$$\Omega(N,V,\alpha) = \prod_i \left[\int_{\alpha} d\vec{u}_i \right] \cdot \det \mathbf{T}^{\alpha} \prod_{\langle ij \rangle} \Theta(r_{ij} - D). \qquad (20)$$

From the standpoint of the (standard) single-phase part of the MC procedure, this change in representation is equivalent to changing the form of the configurational energy:

$$E(\{\vec{r}\}) \rightarrow E(\{\vec{r}\}) - \ln[\det \mathbf{T}^{\alpha}], \qquad (21)$$

This change introduces some computational overheads, which could be substantial if the $\mathbf{T}$ transformation is not local. The potential pay off lies in the LS part of the MC procedure. One might hope to be able to tune the form of the $\mathbf{T}$ matrix so that ''typical'' configurations of the one structure are mapped (by LS) into ''typical'' configurations of the other. In the case of the hard sphere problem, however, our results (Sec. IV A) suggest that there is little to be gained here by this kind of tuning.

III. IMPLEMENTATION

A. Monte Carlo procedures

First we consider the procedure for MC sampling of the particle displacements, for a given structure (set of lattice vectors). As discussed in Sec. II A this sampling should, in principle, satisfy some appropriate configurational constraint [46]. In our original studies [30] we chose to implement this constraint *explicitly*, through our sampling distribution: candidate displacements were drawn from a flat (''top hat'') distribution. This procedure can be made to work. But the constraint explicitly breaks the translational invariance; and one must deal with the consequences. In particular the configurational integral effectively being evaluated then *depends* upon the location of the center of mass and thence upon the

top-hat cutoff; this dependence sets in when the displacement acquired by the center of mass, in the course of its slow diffusive motion, becomes comparable with the top-hat cutoff. One can avoid this problem simply by fixing the center of mass. Our failure to do so in Ref. [30] led to results which differ significantly from those we present here. In the studies reported here we have chosen the ''implicit'' realization of the configurational constraint (practically, but not conceptually equivalent to ignoring it) which rests (Sec. II A) on time scales. Spheres were chosen at random, and trial *changes* to the current displacement drawn from a uniform distribution. The displacement update is accepted according to the Metropolis prescription [3]

$$p_a(\{\vec{u}\} \rightarrow \{\vec{u}'\}) = \min\{1, \exp[-\Delta\widetilde{E}(\{\vec{r}\})]\} \qquad (22)$$

where $\widetilde{E}(\{\vec{r}\})$ is defined in Eq. (17) [47]. In addition to the constraint that the update should yield a realizable configuration of the current phase, this acceptance probability reflects the chosen weights which are defined (Sec. III B explains how) on the space of the overlap order parameter $\mathcal{M}$ [Eq. (16)]. To minimize the computational time spent determining how a proposed move affects the value of $\mathcal{M}$ we used a local overlap array, holding information on which neighbors of a given sphere currently overlap with that sphere in the conjugate configuration generated by a LS.

The representation of the close-packed limit provided by Eq. (10a) can be handled with only minor amendments: the constraint $r_{ij} > D$ identifying realizable configurations is replaced by a constraint on the scaled displacement-difference coordinates $u_{ij}^{\parallel} > -1$. The overlap order parameter (measuring the number of times the hard sphere constraint is violated in the conjugate configuration) is redefined accordingly. In this limit particle ''interactions'' (encounters) may occur only between immediate neighbors. At other densities we allowed for the possibility of encounters between nominal second neighbors. We found, however, that although the number of such encounters grows rapidly with the approach to the melting density, the consequences for the relative entropy of the two structures is insignificant under the conditions studied here [48].

In addition to particle moves the constant-pressure simulations require updates of the simulation-cell parameters. In such an update (implemented on average once per sweep) a trial set of cell parameters are selected, and accepted with probability [49]

$$p_a(V \rightarrow V') = \min\{1, \exp[-\Delta\widetilde{E}(\{\vec{r}\}) - P^{\star}\Delta V \\ + N\ln(V'/V)]\}, \qquad (23)$$

where V' is the volume associated with the trial parameters. Note that this kind of update—a dilation—changes $\widetilde{E}(\{\vec{r}\})$ both trivially (so as to forbid moves causing ''real'' overlaps) and more subtly through changes in the count of the overlaps in the conjugate configuration. A volume update thus requires recalculation of the entire local overlap array.

Now consider the lattice switch. The switch may be viewed as an updating of the ''lattice'' type α, regarded as a stochastic variable. The prescription for such an update is quite simple. After every particle update the value of $\mathcal{M}$ is
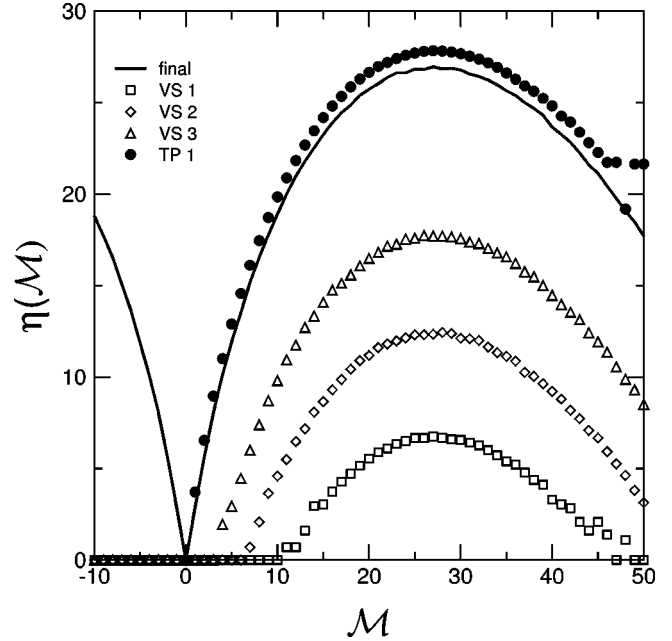

FIG. 5. Illustration of the weight-generation process, for a system of $N = 216$ hard spheres. The points marked VS are the results of the first 3 iterations of the visited-states algorithm, initiated from an fcc equilibrium state. The points marked TP emerge from one application of the transition probability method. The solid line shows a refined (usable) set of weights.

checked (it is already known). The LS is performed if (and only if) the gateway condition $\mathcal{M} = 0$ is satisfied.

B. Calculating the weights

The determination of an acceptable set of multicanonical weights [15] can be accomplished in a number of ways—none, seemingly, entirely systematic. We describe briefly the techniques we have used in the present study. Figure 5 provides some illustration. For further details and references to other work the reader is referred to Refs. [15,26,50,51].

The simplest method is the visited-state (VS) technique [50]. In this approach a suitable set of weights is evolved through an iterative process (Fig. 5), the next set of weights depending upon the distribution of the (overlap) order parameter over the macrostates visited using the current set of weights. This process is repeated until the weights yield a distribution $P(\mathcal{M}|N, V, \{\eta(\mathcal{M})\})$ that is effectively flat. This method proved quite adequate for our smallest system.

For larger systems, however, we found it more efficient to appeal to the transition probability (TP) method [50]. In the simplest realization of this method the simulation is initiated from a ''cold'' (zero displacement) configuration (a member of the $\mathcal{M} = 0$ macrostate) for one structure. In the course of its subsequent evolution towards equilibrium for that structure the numbers of transitions between different $\mathcal{M}$ macrostates are recorded, and subsequently used to construct an estimator of the macrostate-transition-probability matrix. This TP matrix can be used to estimate the macrostate probability distribution and thence to provide an estimate for a set of weights (Fig. 5), which can in turn be refined via VS. For our intermediate size system this method worked well.

In the case of our largest system we found it necessary to modify the method somewhat, so as to limit the rate at which

the simulation passes through $\mathcal{M}$ space. One way of doing this is to constrain the system to macrostates with overlap order parameters below some ''barrier'' value, which is gradually incremented (moved ''out'' in $\mathcal{M}$ space), at intervals of the order of the equilibration time.

By *fiat* the two structures have the same weights for $\mathcal{M}$ $=0$. In principle, the weights associated with the two structures for nonzero $|\mathcal{M}|$ are different [i.e., $\eta(\mathcal{M}) \neq \eta(-\mathcal{M})$], and have to be evolved separately. In practice the weights of the two structures are very similar—a reflection of the similarity of the entropies of the two phases. Consequently, one set of weights provides an excellent first approximation to the other, for refinement by VS.

C. Simulation details

The specific form of the LS operation we have chosen (Fig. 3) imposes restrictions on the geometry of the system simulated: with normal periodic boundary conditions the system must comprise integral multiples of 6 close-packed planes. It is possible to avoid this restriction by using more elaborate boundary conditions [14], but we chose to avoid this complication and simulate systems comprising $6^3 =216$, $12^3 =1728$, and $18^3 =5832$ spheres. Simulations were performed at two densities, namely, [see Eq. (7)] $\tilde{\rho} =0.7778$ [47], and $\tilde{\rho} =1$, the close-packed limit.

The maximum step size for displacement updating was chosen so as to minimize the autocorrelation time of the overlap order parameter [Eq. 16]. We found a maximum step size of $0.13D$ produced the best results at $\tilde{\rho} =0.7778$, while a value close to unity was found to be appropriate in the close-packed limit, in the representation (and scaled units) given in Eq. (10a).

A significant proportion of our simulation time was devoted to the process of weight-determination. For our largest system we used 10^6 Monte Carlo sweeps (MCS) to generate a first (TP) estimate of the weights, with a further 5×10^6 MCS devoted to weight-refinement using VS.

The free-energy differences of interest were then determined by further simulations in the multicanonically weighted ensemble. For each system (density, and size) we performed a series of runs each long on the scale of the autocorrelation time of the overlap order parameter. Each of these runs then provides an independent estimate of the (logarithm of the) probability ratio required [Eq. (18)]. The standard deviation of these estimates provides a basis for assigning an associated statistical uncertainty. Implementing this stage required simulation times ranging from $\sim 2.5 \times 10^8$ MCS for $N=216$ to $\sim 4 \times 10^7$ MCS for $N=5832$.

IV. RESULTS

A. The effects of the representation

As discussed in Sec. II C the LS operation can be implemented with different choices of representation of the lattice mapping or the particle displacements [45]. The efficiency of a lattice mapping is measured (inversely) by the equilibrium overlap count. Table I shows results for a variety of mappings, chosen to expose the different factors that control the mapping efficiency. Mapping number 1 is the one described

TABLE I. The efficiency of different lattice mappings (for $N =1728$ and $\tilde{\rho} =0.7778$), as measured by the number of overlaps (per sphere) that they generate. Refer to the text for details.

mapping	description	effect	$m=\mathcal{M}/N$
1	$(0,-\vec{t},+\vec{t})$	fcc→hcp	0.150(1)
2	$(0,2\vec{t},-2\vec{t})$	fcc→hcp	0.183(1)
3	$(0,3\vec{t},-3\vec{t})$	fcc→fcc	0.194(1)
4	random-plane	fcc→hcp	0.373(2)
5	random-site	fcc→hcp	0.820(3)

in Fig. 3, and used throughout this work: the notation $(0, -\vec{t},+\vec{t})$ signifies that the three pairs of planes counting from the top of Fig. 3 are translated respectively by $0,-\vec{t}$, and $+\vec{t}$. A similar convention is used to label mappings 2 and 3. In mapping 4 (''random-plane'') an hcp configuration is generated by taking an fcc configuration and restacking its close-packed planes in a random order, in an hcp pattern. In mapping 5 (''random-site'') an hcp configuration is generated by mapping the particle displacements in an fcc configuration randomly on to the sites of an hcp lattice.

The random-site mapping (number 5) shows the largest overlap count. One can account for its value, rather well, by regarding the particle displacements as isotropic, Gaussian, and independent of structure [52], and estimating the probability that two particles associated with nearest-neighbor sites, and with displacements drawn *randomly* from this distribution, will overlap.

Using the random-plane mapping (number 4) cuts the overlap count by a factor of (a little more than) 2 with respect to random site. This efficiency gain simply reflects the fact that of the $6N$ potential overlaps between near-neighbors, only the $3N$ associated with neighbors in different (but adjacent) planes can now contribute.

Mapping 3 simply generates one fcc configuration from another (it is useful only because it is informative): its overlap count is cut by a further factor of 2. This reflects the fact that this mapping (similar to mappings 1 and 2) moves close-packed planes in *pairs*, thus guaranteeing no overlaps between the two members of each pair.

Mappings 2 and 1 show further—smaller but still practically useful—cuts in the overlap count. The origin of these gains is more interesting. It is clear that they must reflect the size of the translation vector used: mappings 1 through 3 differ only in this respect. This vector controls the extent of the shear which the mapping introduces between successive pairs of planes. The following interpretation seems reasonable. The displacement patterns in adjacent planes will be correlated to some extent, with undulations in one surface (the z components of the displacements) matched to undulations in its neighbor. The smaller the shear, the more closely these undulations will *remain* matched to one another (in the conjugate configuration), and the smaller the overlap count. With increasing shear, this advantage is lost and the behavior should (and indeed does) approach the limit (one quarter of the overlap count for mapping 5) one would expect in the absence of such correlations. The fact that this ''approach'' is already apparent in the performance of mapping 2 is consistent with the fact that the measured correlation length of
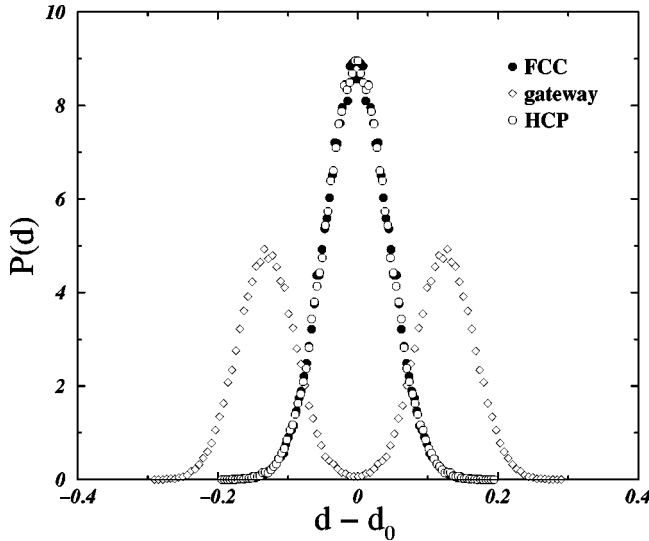
FIG. 6. Distribution of the separation d between adjacent close-packed planes in a system of 216 spheres at $\tilde{\rho}=0.7778$, in the equilibrium hcp and fcc macrostates, and in the gateway ($\mathcal{M}=0$) macrostate. The separation is measured with respect to the equilibrium separation d_0 and is expressed in units of the sphere separation δ [57].

the surface undulations at the density concerned is found to be close to the magnitude of the translation vector $\vec{t}$.

These results help to clarify the factors which control the overlap count of the mapping (number 1) we have actually used. It is tempting to attribute the overlaps to the fact that the LS (fcc→hcp, say) maps each particle from an environment in which adjacent close-packed planes have different stacking labels (A and C, say) to one in which they have the same label (C, say). The results for mappings 1–3 show that it would be misleading to think this way. The overlaps simply reflect the numbers of particles that ''see'' a new adjacent close packed plane (irrespective of its label), and the extent to which it is ''new.'' This is the reason for the simi-

larity between the overlap counts for the two structures (Sec. IV C). It shows, moreover, that any simple [53] tuning of the displacement representation (the choice of **T** matrix) is likely to be of no advantage here [54].

B. How it works: the gateway configurations

A LS operation will work (be accepted as a MC move) only when launched from a small subset of the configurations actually visited: these, by definition, are the ''gateway configurations.'' As noted earlier, one could identify *a priori* configurations (those characterized by ''small enough'' displacements) which fall into this set. But we have elected, rather, to let the system (the algorithm) identify them on the basis of their defining characteristic—that they have zero overlap order parameter $\mathcal{M}$ [28]. It is then interesting to investigate the *microscopic* characteristics of the configurations picked out by this constraint. Figure 6 shows the distribution of the separation d between adjacent close-packed (x-y) planes [55], for $\mathcal{M}$ macrostates corresponding to equilibrium fcc, equilibrium hcp, and gateway ($\mathcal{M}=0$) regions. The macrostates corresponding to the equilibrium crystal structures have similar, near-Gaussian, d distributions. In contrast, for the gateway macrostate the distribution is bimodal: in this macrostate, some planes are systematically moved closer to one another, while (in equal measure) others are shifted apart. On closer examination one finds that it is the planes which are *translated together* by the LS [e.g., the pair of planes marked (i) in Fig. 3] that fall into the first category, while the planes that are *translated differently* by the LS [e.g., the pair of planes marked (ii) in Fig. 3] fall into the second. The evolution, with $\mathcal{M}$, of the mean plane separation (for both categories) is shown in Fig. 7(a). The behavior thus unearthed is entirely reasonable. The LS operation can *only* create overlaps between neighboring planes which are translated by different amounts (sheared with respect to one another). The algorithm resolves the task set by the bias towards $\mathcal{M}=0$ by moving these pairs of planes (the ones vulnerable to overlaps) further apart, at the expense of a
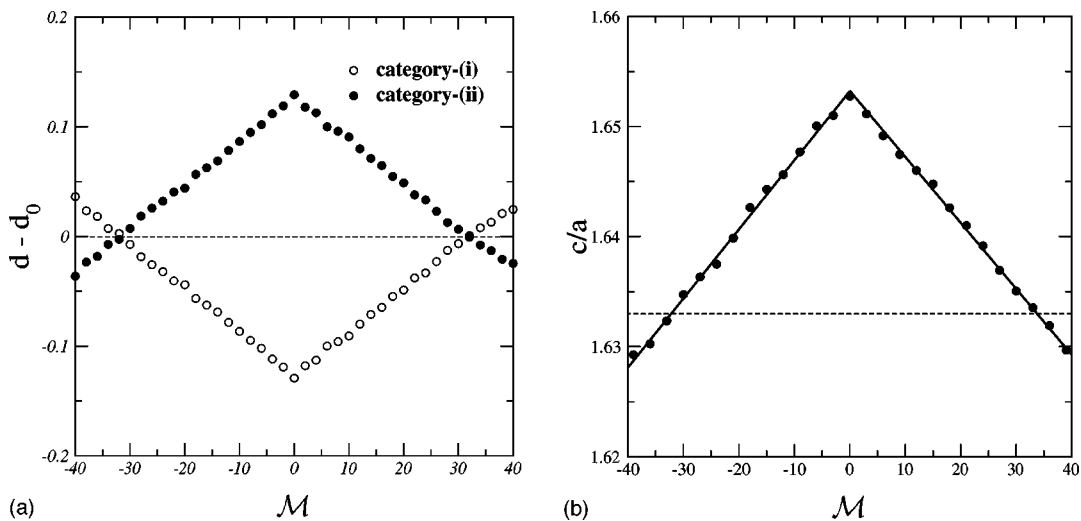

FIG. 7. (a) The mean value of the separation d between adjacent close-packed planes in a system of 216 spheres at $\tilde{\rho}=0.7778$, for macrostates of different $\mathcal{M}$. The separation is measured with respect to the equilibrium separation d_0 in units of δ [57]. Category (i) planes (see Fig. 3) are translated together by LS; category (ii) planes are translated through different amounts by LS. (b) The evolution with $\mathcal{M}$ of the c/a ratio [58] in a constant-*pressure* ensemble [with the same parameters as (a)]. The horizontal line marks the ideal-close-packed value.
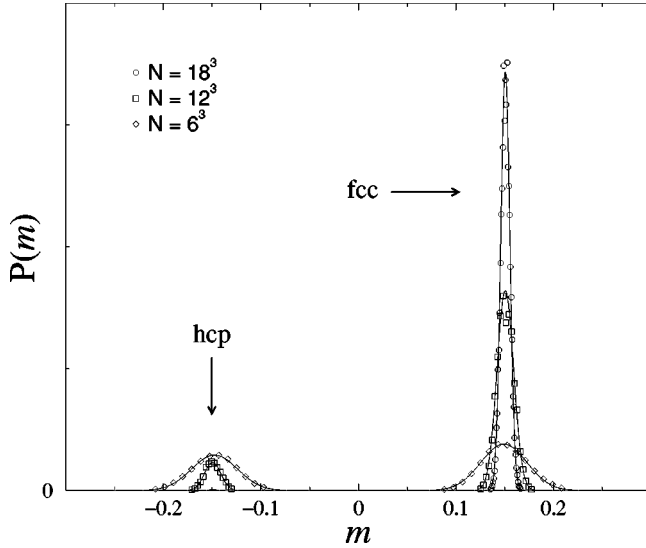
FIG. 8. The probability distribution of the overlap order parameter per particle, $m \equiv \mathcal{M}/N$, for systems of three different N values at $\tilde{\rho}=0.7778$. The lines provide Gaussian guides to the eye; the statistical uncertainties on the data points are smaller than the symbol size. The entropy difference is identified from the logarithm of the ratio of the integrated weights of the two peaks. The hcp peak for the largest system is not visible on this scale.

TABLE II. The difference in the entropy densities of the fcc and hcp structures, $\Delta s \equiv \Delta s_{\text{fcc,hcp}}$ [Eq. (11)]; the associated uncertainties are in parenthesis. The results of the present work (PW) supercede those of Ref. [30]. The results of Ref. [64] supercede those of Ref. [10]. IM stands for integration method; SM is the lattice shear method of Refs. [14,63].

ρ/ρ_{CP}	N	$\Delta s (10^{-5} k_B)$	Method	Ref.
0.731	512	85 (10)	SM	[14]
0.736	12000	230 (100)	IM	[64]
0.736	12096	87 (20)	IM	[62]
0.739	512	90 (4)	LS	[14]
0.7778	216	132 (4)	LS	[31]
0.7778	1728	112 (4)	LS	[31]
0.7778	1728	113 (4)	IM	[31]
0.7778	216	133 (3)	LS	PW
0.7778	1728	113 (3)	LS	PW
0.7778	5832	110 (3)	LS	PW
1.0	12000	260 (100)	IM	[64]
1.0	512	110 (20)	SM	[14]
1.0	64	91 (5)	LS	[14]
1.0	216	107 (4)	LS	[14]
1.0	512	119 (3)	LS	[14]
1.0	1000	113 (4)	LS	[14]
1.0	216	131 (3)	LS	PW
1.0	1728	125 (3)	LS	PW

compression of the others [56]. In simulations conducted at constant pressure this effect (still present) is supported by a second. Figure 7(b) shows that the algorithm now exploits the additional degrees of freedom (the *shape* of the simulation cell) to locate gateway states with values of the c/a ratio enhanced above the ideal close-packed value [58]. Again, the advantages with respect to the switch are clear.

It is tempting to say that the sampling is intelligent. In any event it is clear that the algorithm locates and utilizes configurations which it would be difficult to exploit explicitly in the design of the switch operation.

C. Entropies of crystalline structures

The essential output of a LS simulation is in the form of the normalized probability distribution of the overlap-order parameter, reweighted to remove the bias in the multicanonically weighted distribution actually measured. Figure 8 shows the results for this distribution (at $\tilde{\rho}=0.7778$) for three different N values. As one would expect the distributions each comprise two peaks (one associated with each phase) each of which is *nearly* Gaussian [59] and sharpens with increasing N [60]. Note the close correspondence between the equilibrium overlap counts for the two structures. This result is not *required* by definition, or any obvious symmetry. Rather it should be seen as a further manifestation (the smallness of the entropy difference between the phases is the prime one) of the similarity of the local particle environments in the two structures.

The relative weights of the two peaks is a direct measure of the difference between the entropies of the two structures [Eqs. (11), (12), (18)]. Since the entropies are extensive the ratio of the peak weights grows exponentially with N [61]; the fact that (in this case, at least for our smaller systems) the two peaks can even be displayed on the same scale is a

reflection of the exceptionally delicate balance between the two entropy densities.

Figure 8 allows one to *see* that fcc is the thermodynamically preferred structure. This conclusion is expressed quantitatively in the results gathered in Table II. Our results at $\tilde{\rho}=0.7778$ correct those of our earlier work [30], as explained in Sec. III A. They are in full accord with the results (both LS- and IM-based) reported by Pronk and Frenkel [31]. The close correspondence between the results for $N=1728$ and $N=5132$ confirms that the former system is already representative of the thermodynamic limit. Table II also shows the results of our studies at the close-packed limit, using the hard-dodecahedron representation [Eq. (10a)]. Our results seem at variance with the IM-based result of Woodcock [64], even allowing for the large uncertainty attached to that result. They are close to those (based on LS) reported by Mau and Huse [14]. But the differences (for the smaller systems, particularly) appear to be statistically significant [65]. Figure 9 gives an alternative view of these results. It utilizes the parametrization of the measured pressure difference between the two phases provided by Speedy [66] to determine the entropy difference, as a function of density, *given* the entropy difference at a chosen reference density; we have used the results of the present work at $\tilde{\rho}=0.7778$.

Table III shows the results of our studies in the constant pressure ensemble. The quantity of interest here is the difference between the Gibbs free energy densities at the chosen pressure, which follows from the relevant distribution with the aid of Eq. (15). In fact the Gibbs free energy density difference Δg for a given pressure, and the entropy density difference Δs at a physical density that is the thermodynamic
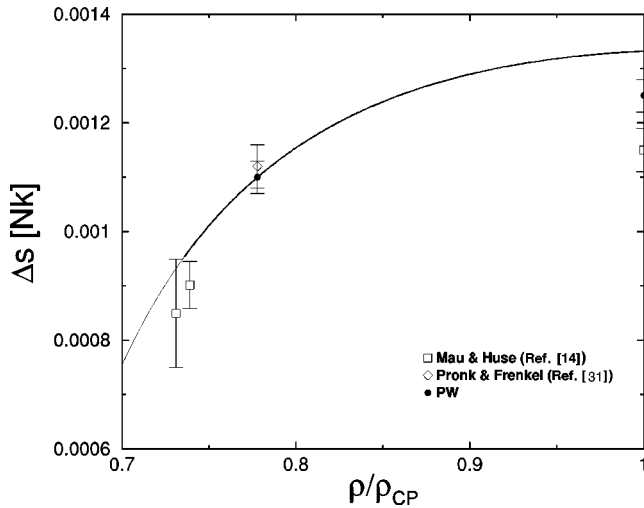
FIG. 9. The difference in the entropy densities of the fcc and hcp structures, $\Delta s \equiv \Delta s_{\text{fcc,hcp}}$ [Eq. (11)], as a function of reduced density $\tilde{\rho}$. The data points are as given in Table II. The solid line is the result of an integration of the pressures of the phases [66]. Note that this line passes through our result at $\tilde{\rho} = 0.7778$ by *construction*.

conjugate of that pressure for one of the phases, differ (in magnitude [67]) by terms that are *second order* in the pressure difference between the two phases. That pressure difference is extremely small [66], as is the difference between the measured densities of the two structures (Table III). In these circumstances one would expect the magnitude of Δg to fall on the Δs plot in Fig. 9; and indeed, within the residual uncertainties, it does.

V. DISCUSSION: REVIEW AND PROSPECTS

In the work described here we have been concerned both with a *system* of long-standing interest—the hard sphere crystal—and a *method*—lattice-switch Monte Carlo—with potentially wide applicability. We divide our concluding discussion accordingly.

The full agreement between the present work and that of Ref. [31] leaves little doubt that the equilibrium entropy difference between the two close-packed structures has finally been established securely and with high precision—at least at one density. Although a small discrepancy with respect to the results of Ref. [14] remains, the accord of our close-packed limit results with those established using pressure difference measurements [66] suggests that the curve in Fig. 9 provides a relatively complete and trustworthy picture of the density dependence.

Notwithstanding the simplicity of the model, these results do have implications for experimentally realizable systems. The immediate relevance to atomic systems is tenuous [68], but the model has been widely used to account for the behavior of assemblies of ''hard,'' ''spherical'' colloidal particles [32]. Since the predicted entropy-density difference is so small there are potentially many ways (residual interactions between the spheres; polydispersity) in which the applicability of the theory may be compromised. But, of these, it seems that the most significant issues to be addressed are to do with *scales*—length and time.

TABLE III. The difference in the gibbs free energy densities of the fcc and hcp structures $\Delta g \equiv \Delta g_{\text{fcc,hcp}}$ [Eq. (15)]; the associated uncertainties are in parenthesis. $P^{\star}$ gives the pressure [42] in units of $k_B T/D^3$.

$P^{\star}(D^{-3})$	$\tilde{\rho}_{\text{hcp}}$	$\tilde{\rho}_{\text{fcc}}$	N	$\Delta g(10^{-5}k_B T)$	
14.58	0.7776(1)	0.7775(1)	216	-113	(4)
14.58	0.7770(3)	0.7774(2)	1728	-112	(3)

First, the lengths. In the experiments reported in Ref. [69] the colloidal particles have diameters of order 10^{-7} m and the samples comprise crystallites with linear dimensions of order 10^{-5} m. The number of particles in such a crystallite ($N \sim 10^6$) is large compared to those in our simulation, which is (as we have seen) sufficient to allow us to *deduce* properties of the thermodynamic limit. But it is *not* large enough to guarantee that the behavior displayed will actually *be* that of the thermodynamic limit. To see this—and its principal implications—one needs to consider the stability of the perfect fcc crystal with respect to hcp-type stacking faults. Following Ref. [69] we may introduce a parameter α [70] measuring the probability that a chosen close-packed plane sits within an fcc environment as distinct from the hcp environment. A simple argument (Appendix A) using the pseudospin parametrization of stacking patterns provided in Ref. [14] then yields the result

$$\alpha = \frac{1}{2}\left(1 + \tanh\left[\frac{N_{\perp}\Delta s}{2}\right]\right), \quad (24)$$

where $N_{\perp}$ is the number of particles in a close packed layer and Δs (a function of $\tilde{\rho}$) is the fcc-hcp entropy difference per particle, as given in (and in the units of) Table II. The thermodynamic ideal ($\alpha = 1$) is thus realized only to the extent that $N_{\perp}\Delta s$ is large compared to unity. For the length scales given above, $N_{\perp}\Delta s \simeq 1$. The obvious implications are qualitatively consistent with the observations reported in Ref. [69] which show α values (deduced from Bragg scattering intensities) ranging from 0.5 [signaling essentially random-hexagonal-close packing, (rhcp)] through to $\alpha = 0.8$.

The observed spread in α values reflects, presumably, the issue of time scales. The smallness of the entropy difference (which supplies the kinetic driving force towards the equilibrium state) suggests that the equilibrium behavior will be observed only in samples which are grown sufficiently slowly and (or) given sufficient time for subsequent annealing [71]. The results of Ref. [69] do indeed suggest a correlation between observed α value and the slowness of the growth process. Experiments done in microgravity [72], where growth processes are greatly accelerated, yield essentially randomly close-packed crystals.

Now let us turn to the lattice-switch method. There are two questions here. One, does the method represent a significant advance with respect to existing methods? Two, is it generally applicable?

The main alternative method (the benchmark against which others need to be assessed) is probably integration along a reference path, of which the work reported in Ref. [31] represents, to our knowledge, the most refined example.

If one compares the two techniques (LS and IM) on the basis of precision-for-computational-buck there seems to be no clear winner in the hard-sphere studies to date: Ref. [31] reports calculations using both methods that achieve comparable levels of precision on the basis of comparable computational time. But one should note that the entropy difference ultimately determined is some *four orders of magnitude* smaller [73] than the separate entropies of the two phases, determined by IM. One can see this as a testimony to the care with which the recent IM studies have been carried out; or (as suggested in Sec. I) as a strong indicator that another approach using an interphase path is called for. There are also two other counts—both somewhat subjective—on which we suggest that the LS approach is superior. First, it seems to us relatively illuminating (by comparison with IM) to read-off the result for a free energy difference directly from a figure similar to Fig. 8 which *shows* what it *means*. Secondly it also seems to us that LS wins in regard to the transparency of the uncertainties to be attached to its results. The LS error bounds represent purely statistical uncertainties associated with the measurement of the relative weights of two distribution peaks. The IM error bounds have to aggregate the uncertainties associated with different stages of the integration process.

As regards the second question, we expect that the method will, with appropriate extensions, be widely applicable. The first extension must clearly be to accommodate soft potentials. The LS operation will then need gateway configurations in which the energies of the two structures (measured with respect to their ground-state energies—or indeed *any* fixed reference energy) are closely matched [43]. The "overlap order parameter" will need to be redefined accordingly. With no more than this degree of elaboration the method should be applicable immediately to investigate the widespread "competition" between fcc and hcp ordering in the phase behavior of the elements [74].

More generally, moving beyond the space of fcc-hcp structures, the choice of lattice-to-lattice mapping will require some thought. Mappings which preserve the relative positions of significant subsets of the particles (the analog of the close-packed planes) are likely to be optimal. The license to choose ones representation of the displacements (Sec. II C) may also prove useful. Simple transformations [53] will help if the mapping takes particles between environments in which the spectrum of single-particle displacements is significantly different. In such cases one might envisage using a MC-annealing procedure to refine the choice of representation. The use of normal coordinates has some advantages here—but possibly not enough to offset the fact that the interaction potential is nonlocal when expressed in Fourier space.

APPENDIX A: DISPLACEMENT ENTROPY VERSUS STACKING ENTROPY

Consider a system of N hard spheres arranged in $N_\parallel$ close-packed layers of $N_\perp$ particles. Following Ref. [14] one may conveniently index each of the close-packed layers with a pseudo-spin (Ising-like) variable σ, where $\sigma_i = 1$ signifies that layer i has an fcc environment (the two immediately adjacent layers are not aligned with one another) while $\sigma_i = -1$ implies an hcp environment (adjacent planes are aligned with one another). The probability of a particular stacking sequence $\{\sigma\}$ (if these variables may properly be regarded as annealed) then satisfies

$$\ln P(\{\sigma\}|N,V) = S(N,V,\{\sigma\}) + \text{const}, \qquad (A1)$$

where $S(N,V,\{\sigma\})$ measures the entropy associated with the configurations (displacements) consistent with the particular structure $\{\sigma\}$. Following Ref. [14] this entropy (we will refer to it here as "displacement entropy") can usefully be written in the form of an expansion:

$$S(N,V,\{\sigma\}) = N s_0 + N_\perp h \sum_i \sigma_i + N_\perp J \sum_{\langle ij \rangle} \sigma_i \sigma_j + \cdots. \qquad (A2)$$

The expansion is effectively ordered in the *range* of the entropic interlayer "interactions": the ellipsis represents contributions from interactions (microscopically, displacement-displacement correlation functions) extending over more than 4 layers. The analysis of Ref. [14] indicates that the series converges quickly, except close to melting. If we neglect the interaction terms altogether we may make the identification

$$h = \frac{1}{N_\perp N_\parallel}[S(N,V,\{\sigma = +1\}) - S(N,V,\{\sigma = -1\})]$$

$$= \frac{\Delta s_{\text{fcc,hcp}}}{2} \qquad (A3)$$

and, from Eq. (A1),

$$\langle \sigma \rangle = \frac{1}{N} \sum_{\{\sigma\},i} P(\{\sigma\}|N,V)\sigma_i = \tanh[N_\perp h]$$

$$= \tanh\left[\frac{N_\perp \Delta s_{\text{fcc,hcp}}}{2}\right] \qquad (A4)$$

from which Eq. (24) follows. The correspondence with a 1D paramagnet is clear. The familiar competition (between orientation energy and entropy) is played out here as a competition between displacement entropy and stacking entropy, with $N_\perp$ playing the role of an inverse temperature.

[1] G. Jacucci and N. Quirke, in *Complete Simulation in the Physics and Chemistry of Solids* edited by C.R.A. Catlow and W. C. Mackrodt, Vol. 166 of Lecture Notes in Physics, (Springer, Berlin, 1982), p. 38.

[2] Most, perhaps all, of the methods referred to here may also be realized within the framework of molecular dynamics.

[3] K. Binder, Rep. Prog. Phys. **60**, 487 (1997).

[4] D. Frenkel, in *Molecular-Dynamics Simulation of Statistical-Mechanical Systems*, edited by G. Ciccotti and W.G. Hoover (North-Holland, Amsterdam, 1986), p. 151.

[5] M.P. Allen, in *Monte Carlo and Molecular Dynamics of Condensed Matter Systems*, edited by K. Binder and G. Ciccotti (Italian Physical Society, Bologna, 1996), p. 255.

[6] A path is defined by a sequence of values of either some macroscopic observable or some parameter λ, which controls a thermodynamic field or a model parameter.

[7] In some instances one needs to make use of separate reference systems for each phase.

[8] A full survey is not appropriate here; we note only some examples. References [9] and [10] both use multistage (integration) methods involving reference paths—the former to an Einstein solid, and the latter to a single-occupancy-cell model of an ideal gas–to study the fcc-hcp phase behavior of hard spheres. Reference [11] describes a study of the fcc-bcc phase behavior of Lennard-Jones systems, using a multistage (overlap) method applied to a nonphysical interphase path, with the λ parameter [6] indexing a configurational energy that interpolates between those of the two structures. This method has been widely used (see, e.g., Ref. [12]). Reference [13] describes a physical interphase path linking NaCl and CsCl structures, but does not utilize the path for a free-energy-difference calculation. The ''lattice-shear'' method described in Ref. [14], and applied to hard spheres, is a natural refinement of Ref. [13] in which a single-stage sampling procedure is used to explore a path constructed out of a series of mutually sheared lattices that interpolate between the lattices of interest.

[9] D. Frenkel and A.J.C. Ladd, J. Chem. Phys. **81**, 3188 (1984).

[10] L.V. Woodcock, Nature (London) **384**, 141 (1997).

[11] A. Rahman and G. Jacucci, Nuovo Cimento D **4**, 357 (1984).

[12] M.C. Moody, J.R. Ray, and A. Rahman, J. Chem. Phys. **84**, 1795 (1985); B. Kuchta and R.D. Etters, Phys. Rev. **47**, 14 691 (1993).

[13] M. Parrinello and A. Rahman, J. Phys. (Paris) **42**, 511 (1981).

[14] S.-C. Mau and D.A. Huse, Phys. Rev. E **59**, 4396 (1999).

[15] B.A. Berg and T. Neuhaus, Phys. Rev. Lett. **68**, 9 (1992); B.A. Berg, J. Stat. Phys. **82**, 323 (1996).

[16] A.P. Lyubartsev, A.A. Martsinovski, S.V. Shevkunov, and P.N. Vorontsov-Velyaminov, J. Chem. Phys. **96**, 1776 (1992).

[17] E. Marinari and G. Parisi, Europhys. Lett. **19**, 451 (1992).

[18] G.M. Torrie and J.P. Valleau, Chem. Phys. Lett. **28**, 578 (1974).

[19] Since, by definition, the end points of an interphase path lie in different phases, the associated macroscopic property may reasonably be described as an *order parameter*.

[20] The nature of the configurations visited along a path is not generally predictable *a priori* from the choice of order parameter. One may say only that the configurations sampled at a given point on the path will be those which have measurable *canonical* probabilities *conditioned* on that macrostate, and which are, moreover, accessible on the relevant time scales.

[21] The interfacial free energy emerges as a by-product—and in some cases [22] may actually be the principal focus of interest.

[22] B.A. Berg, U.H.E. Hansmann, and T. Neuhaus, Z. Phys. B: Condens. Matter **90**, 229 (1993).

[23] N.B. Wilding, Phys. Rev. E **52**, 602 (1995).

[24] B. Grossmann, M. Laursen, T. Trappenberg, and U.J. Wiese, Phys. Lett. B **293**, 175 (1992).

[25] It might work acceptably well if the dynamics of the interface between the two phases is favorable: systems with martensitic phase transitions may fall into this category: Z. Nishiyama, *Martensitic Transformations* (Academic, New York, 1978). Note also that the special case in which the structural phase transition involves no change of symmetry can be handled within the standard multicanonical framework: see Ref. [26].

[26] G.R. Smith and A.D. Bruce, Phys. Rev. E **53**, 6530 (1996).

[27] The *conjugate* of a given configuration is the configuration associated with the *same* set of displacements attached to the *other* set of lattice vectors.

[28] The *general* defining characteristic of the *gateway* configurations is that a LS operation, launched from within this set, will be accepted with a probability sufficient to make the attempt worthwhile.

[29] A configuration is *energy matched* (to its conjugate) if the difference between its energy and that of its conjugate is small, on the scale of $k_B T$.

[30] A.D. Bruce, N.B. Wilding, and G.J. Ackland, Phys. Rev. Lett. **79**, 3002 (1997).

[31] S. Pronk and D. Frenkel, J. Chem. Phys. **110**, 4589 (1999).

[32] P.N. Pusey, in *Liquids Freezing and the Glass Transition*, edited by J.P. Hansen, D. Leveesque, and J. Zinn-Justin (Elsevier, Amsterdam, 1991) p. 763.

[33] In its most general context the LS is the basis for estimating the differences between the *free energy* (Helmholtz or Gibbs) of two crystalline phases; in the case of hard sphere systems, at constant density, this free energy difference is purely *entropic*.

[34] M. Born and K. Huang, *Dynamical Theory of Crystal Lattices* (Clarendon, Oxford, 1968).

[35] The hcp ''lattice'' is not a ''Bravais lattice''; see e.g., Ref. [36], p. 79.

[36] N.W. Ashcroft and N.D. Mermin, *Solid State Physics* (Saunders, Philadelphia, 1976).

[37] W.G. Hoover and F.H. Ree, J. Chem. Phys. **49**, 3609 (1968).

[38] W.G. Rudd, Z.W. Salsburg, A.P. Yu, and F.H. Stillinger, J. Chem. Phys. **49**, 4857 (1968).

[39] The divergence is an artifact of the classical character of the model.

[40] As detailed in Ref. [38] the argument leading to Eq. (10a) actually entails a rescaling of the displacement coordinates by an ϵ-dependent factor.

[41] The ''visualization'' exercise has its limitations. The dodecahedra (and the lattice spacing) should be thought of as ''infinitely large'' compared to the mean separation of adjacent faces.

[42] We subsume a factor of $k_B T$ into the ''pressure'' $P^{\star}$.

[43] If the interparticle potential is not of the hard-sphere form [Eq. (1)] the ''perfect crystal'' configurations (classical ground states) of the two structures will generally have *different* energies. But one may handle the effects of this energy mismatch quite simply, by attaching different multicanonical weights to the gateway configurations of the two structures.

[44] The sign convention here has no deep significance. Defining $\mathcal{M}$ so that it has different signs in the two phases simply allows us to make a visually clear distinction (Fig. 8) between

the contributions which each phase makes to the $\mathcal{M}$ distribution.

[45] In fact, within the framework of periodic boundary conditions, the second generalization subsumes the first. Thus the various lattice-to-lattice mappings discussed in Sec. IV A can all be thought as a *single* mapping but with *different* hcp T matrices [Eq. (19)], chosen to interchange appropriate displacements.

[46] In the absence of *any* phase-defining constraint on the configurational integral specified in Eq. (5), the integral extends over *all* configurations compatible with the boundary conditions.

[47] This density value was chosen to coincide with one of those studied in Ref. [9].

[48] This conclusion is in qualitative accord with that of Ref. [14], although we find the effects of next-neighbor encounters to be somewhat smaller than is reported there.

[49] D. Frenkel, in *Computer Modelling of Fluids, Polymers and Solids*, edited by C.R.A. Catlow, C.S. Parker, and M.P. Allen (Kluwer Academic, Dordrecht, 1990), p. 83.

[50] G.R. Smith and A.D. Bruce, J. Phys. A , **28**, 6623 (1995).

[51] M. Fitzgerald, R.R. Picard, and R.N. Silver, Europhys. Lett. **46**, 282 (1999).

[52] These approximations are surprisingly good, except close to melting. See *inter alia* D.A. Young and B.J. Alder, J. Chem. Phys. **60**, 1254 (1974); J. Piasecki, and L. Peliti, J. Phys. A **26**, 4819 (1993); R. Ohnesorge, H. Löwen, and H. Wagner, Europhys. Lett. **22**, 245 (1993).

[53] By ''simple'' we mean a T matrix that preserves locality, merely acting to rotate or reflect individual displacements. Nonlocal transformations which mix the displacement vectors might reduce the overlap count further.

[54] The fact that the two phases have such similar entropies tells us that *in principle* there must be *some* representation of the LS in which virtually all realizable configurations of one phase transform into realizable configurations of the other: for this representation the white islands of Fig. 1(b) would fill the single phase regions; the effective length of the interphase path would be short, delivering results of essentially arbitrary precision. We think that this representation is not ''simple.''

[55] More precisely, the difference between the z components of the center of mass coordinates of these planes.

[56] The gateway configurations have further distinctive features (with respect to their equilibrium counterparts): the root-mean-square particle displacement, in the z direction, is reduced; and the in-plane correlation length of these displacements (the undulations of the close packed planes) is enhanced along the $\vec{t}$ direction.

[57] We define $\delta = D(\tilde{\rho}^{-1/3} - 1)$, giving the shortest distance between two sphere surfaces in the ideally close-packed perfect crystal configuration.

[58] We define a c/a ratio, for both hcp and fcc structures by c/a $= 2d_0/\tilde{D}$ with d_0 the mean separation between close-packed planes and $\tilde{D}$ the mean separation between neighboring sphere centers in close packed planes. In the hcp structure $c/a = \sqrt{\frac{8}{3}}$ $= 1.63299$ if the packing is *ideal* [36]; in the fcc structure it has this value by symmetry. In our constant-pressure simulations (for $N = 1728, P^\star = 14.58$) we found $[c/a]_{\text{fcc}}$ $= 1.6333(3)$ and $[c/a]_{\text{hcp}} = 1.6332(3)$. Evidently, whatever difference there is between the true values of these quantities is not easily resolved by single-phase averages.

[59] A Gaussian parametrization of the peaks provides a first approximation to a set of multicanonical weights—but one which significantly underestimates the entropic barrier against the LS.

[60] Note that in Fig. 8 the horizontal axis is the *intensive* variable $m \equiv \mathcal{M}/N$.

[61] The fact that the peak-weight difference diverges exponentially fast with N presents no *computational* problem: the multicanonical procedure is designed to cope with (and quantify) differences of this scale.

[62] P.G. Bolhuis, D. Frenkel S.-C. Mau, and D.A. Huse, Nature (London) **388**, 235 (1997).

[63] The ''shear'' implementation described by Mau and Huse [14] is (they report) substantially harder to implement than the switch method which we discuss here, and which they use in most of their work.

[64] L.V. Woodcock, Nature (London) **388**, 236 (1997).

[65] To provide one independent check of the close-packed limit code we determined the amplitude of the leading nontrivial term in the expansion of the pressure about the close-packed limit [38] in the fcc structure, finding agreement to 4 significant figures with results reported in Ref. [66].

[66] R.J. Speedy, J. Phys.: Condens. Matter **10**, 4387 (1998).

[67] But Δs and Δg are of opposite sign.

[68] Y. Choi, T. Ree, and F.H. Ree, J. Chem. Phys. **99**, 9917 (1993). show that predictions for the phase diagram of a Lennard Jones solid depend extremely sensitively on the fcc-hcp hard-sphere entropy difference.

[69] P.N. Pusey, W. van Megen, P. Bartlett, B.J. Anderson, J.G. Rarity, and S.M. Underwood, Phys. Rev. Lett. **63**, 2753 (1989).

[70] Note our double use of α as both a stacking-type probability and a structure label.

[71] Reference [31] estimates the rhcp-to-fcc annealing time for a hard-sphere-colloidal crystal of linear dimension 10^{-3} m to lie in the range ''months to years.''

[72] J. Zhu, M. Li, R. Rogers, W. Meyer, R.H. Ottewill, W.B. Russel, and P.M. Chaikin, Nature (London) **387**, 883 (1997).

[73] See the full accounts of the IM method provided in the original studies [9].

[74] D.A. Young, *Phase Diagrams of the Elements* (University of California, Berkeley, 1991).