Lattice-switch Monte Carlo method: Application to soft potentials

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

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

The lattice-switch Monte Carlo method, recently introduced and applied in the context of hard spheres, is extended to particles interacting through a soft potential. The method utilizes a transformation that switches between configurations of two different crystalline structures, allowing the phase space of both structures to be explored in a single simulation and the *difference* between their free energies to be determined directly. We apply the method to determine the fcc-hcp crystalline phase behavior of the classical Lennard-Jones solid.

I. INTRODUCTION

In recent papers [1,2], we introduced a new technique for tackling an old problem. The technique is called lattice-switch Monte Carlo (LSMC); the problem is the determination of the relative stability of two crystalline structures. The LSMC method is built on a 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 multicanonically biased [3] 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 are sampled within a *single* simulation, and the difference between their free energies evaluated from the ratio of their measured probabilities.

In [1,2] we set out the LSMC method in some detail; we argued that it offers significant advances with respect to existing strategies [4], preeminently integration methods [5,6]; and we explored its operation through extensive studies of the relative stability of the crystalline (fcc, hcp) phases of hard spheres. More recently [7], we have shown that LSMC can be set in a somewhat more general framework that allows exploration of *freezing* (crystal-*liquid* phase coexistence); that work also was set in the context of hard spheres.

In this paper we describe the application of the method to model crystalline solids with a *soft* interaction potential of Lennard-Jones (LJ) form. The generalization from "hard" to "soft" potentials has a number of consequences. Temperature and pressure become independent control parameters [8]. There is the possibility of real structural phase behavior (i.e., crystal-to-crystal transitions induced by changes of temperature and pressure) [9]; and harmonic crystal dynamics and its refinements [10] provide an alternative, approximate, strategy against which to benchmark the LSMC calculations. The LJ model is also of rather wider physical interest: it provides an excellent account of many features of the behavior of rare gas solids (RGS) [11]; one might hope it would also account for the observed *phase* behavior of RGS, at least under "normal" conditions [12]. But the observed phase behavior (the favored RGS crystal structure appears to be fcc, almost everywhere [13]) has proved surprisingly hard to elucidate [14]. The difference between the ground-state energies of hcp and fcc actually favors the hcp structure; the difference is small and its sign can be changed both by small

deformations of the potential [15] and by truncating its range [16–18]. In such circumstances fluctuation effects are clearly crucial. Such effects *may* have either an anharmonic or a quantum character (sometimes both) and *will* necessarily do so, at high enough or low enough temperatures, respectively. A satisfactory, general, computational strategy must deal coherently with both anharmonic and quantum effects. The LSMC method can be developed to do so. But in this work (in keeping with the majority of existing simulation approaches to this problem) we shall neglect quantum effects: we study the structural phase behavior of the (one might say *a*, given the sensitivity to the potential) *classical* LJ model.

The layout of the paper is as follows. Section II defines the model, and explores the sensitivity of the ground state to the range at which the potential is truncated. In Sec. III we summarize the statistical mechanics relevant to the phase stability problem. Section IV describes the standard approach to this problem, based on the small-amplitude (quasiharmonic) treatment of the fluctuation spectrum, and again explores the consequences of potential truncation. Section V formulates the LSMC method in the context of soft potentials; Sec. VII presents the results it delivers. They show that the LSMC method captures the results of the harmonic theory in the regime (high enough density, low enough temperature) where one would expect the harmonic approximation to be satisfactory; and that it extends naturally beyond this regime, revealing anharmonic corrections that grow with reducing density, in an intelligible way. Since the method focuses directly on the *difference* of the free energies between the phases it is able to measure it with precision, despite its intrinsic smallness. We thus have confidence that the implied phase diagram is truly representative of the classical LJ model. Finally in Sec. VIII, we review what we have learned from this study about the LS methodology and its value in relation to the many other strategies that have been used to attack the problem of free-energy measurement; and about LJ systems themselves.

II. MODEL AND GROUND STATE

We consider a system of N particles of spatial coordinates $\{\vec{r}\}$, confined within a volume V and subject to periodic boundary conditions. Each particle is associated with a

site of a lattice [19] defined by a set of vectors $\{\vec{R}\}_\alpha$ where α indexes the lattice type [20].

The particles interact through some approximation to the LJ potential

$$\phi(r_{ij}) = 4\epsilon \left[\left(\frac{\sigma}{r_{ij}} \right)^{12} - \left(\frac{\sigma}{r_{ij}} \right)^6 \right], \quad (1)$$

where $r_{ij} = |\vec{r}_i - \vec{r}_j|$. The customary approximation is some form of truncation, which is required, in principle, to avoid self-interactions in a finite periodic system and is also a practical expedient to limit computational cost. One may define a truncated potential in a number of different ways. Thus, one may truncate the potential at a *fixed* cutoff $r_x = m\sigma$ (for some chosen m), defining

$$\phi_a(r_{ij}) = \begin{cases} \phi(r_{ij}) - \phi(r_x) & r_{ij} \leq r_x \\ 0 & r_{ij} > r_x \end{cases}. \quad (2)$$

In this case the potential is also "shifted" so as to suppress the discontinuous change occurring when particle separations pass through the cutoff value, which they do, both stochastically in the course of the crystal dynamics and systematically with changes in control parameters such as density or pressure.

Alternatively one may choose a truncation scheme in which pairs of particles are identified at the outset as interacting or not, according to the separation of the *lattice sites* with which they are associated. Thus, setting $r_x = mr_0$ where r_0 is the equilibrium nearest-neighbor separation, we define

$$\phi_b(r_{ij}) = \begin{cases} \phi(r_{ij}) & R_{ij} \leq r_x \\ 0 & R_{ij} > r_x \end{cases}, \quad (3)$$

where $R_{ij} \equiv |\vec{R}_i - \vec{R}_j|$. The fact that the cutoff r_x now scales appropriately with density ensures that the list of interacting particles prescribed by Eq. (3) is indeed fixed (independent of density) so that a potential *shift* is redundant in this case.

The significance of the choice of truncation is revealed in the (classical) ground-state energy

$$E^0 = E(\{\vec{R}\}) = \sum_{\langle ij \rangle} \phi(R_{ij}), \quad (4)$$

where the sum counts (once) the contribution of each *pair* of interacting particles. Figure 1 shows the *difference* [21] between the ground-state energies of the two structures (per particle)

$$\Delta \epsilon^0_{\alpha,\tilde{\alpha}} \equiv \frac{E(\{\vec{R}\}_\alpha) - E(\{\vec{R}\}_{\tilde{\alpha}})}{N}, \quad (5)$$

as a function of density, with $\alpha \equiv$ fcc and $\tilde{\alpha} \equiv$ hcp. In addition to the exact result following from implementing Eq. (5) without approximation [22–24], we show the results that follow if the potential is approximated by the truncated forms ϕ_a [Eq. (2)] and ϕ_b [Eq. (3)].
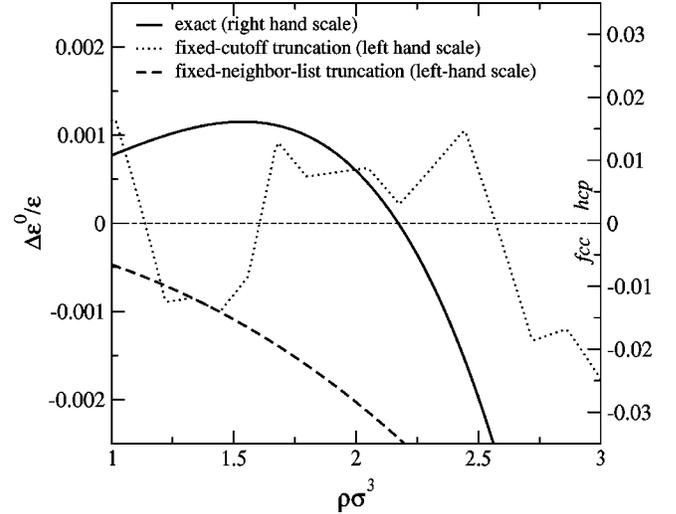
FIG. 1. The difference between the ground-state energies per particle for fcc and hcp structures, as function of density. The solid line shows the result with no potential truncation. The dotted and dashed lines show the results based on the fixed-cutoff truncation scheme [Eq. (2) with $r_x = 2.5\sigma$] and the fixed-neighbor-list truncation scheme [Eq. (3) with $r_x = 2.5r_0$] respectively. Note the order-of-magnitude difference between the scales on the left and right hand axes that are used for "exact" and "truncated" results respectively.

The exact result shows the well-known $T = 0$ phase behavior [11]: there is a single-phase boundary (at $\rho = \rho_0 = 2.1727\sigma^{-3}$) separating hcp ($\rho < \rho_o$) and fcc ($\rho > \rho_o$) regions. By contrast the fixed-cutoff approximation [Eq. (2) with $r_x = 2.5\sigma$] leads to a ground-state energy that varies wildly with density, as the number of particles within the interaction range evolves. The implied (seemingly rich) phase behavior [17] can be traced directly to the truncation. The fixed-interaction-list approximation [Eq. (3) with $r_x = 2.5r_0$] fares differently, but no better: the evolution is smooth, but shows no phase boundary at all.

The moral is clear: to capture phase behavior that is any trustworthy sense representative of the true LJ model one must handle the ground-state energies without recourse to truncation. While it is not our aim to give a definitive treatment of the LJ problem (our treatment is classical, after all) we shall respect this constraint. Thus, in what follows, we shall take for the *general* configurational energy the form

$$E(\{\vec{r}\}) = E^0 + \sum_{\langle ij \rangle} \phi_c(r_{ij}), \quad (6)$$

with

$$\phi_c(r_{ij}) = \begin{cases} \phi(r_{ij}) - \phi(R_{ij}) & R_{ij} \leq r_x \\ 0 & R_{ij} > r_x \end{cases}. \quad (7)$$

The ground-state energy E^0 is treated essentially exactly [25]. The effects of the truncation are restricted to the fluctuation spectrum; we shall discuss them (and the choice of r_x) in due course.

III. STATISTICAL MECHANICS

The fluctuation spectrum (the "significant" configurations $\{\vec{r}\}$) of a crystalline phase α is naturally described by the displacements with respect to the lattice sites $\{\vec{R}\}_\alpha$. Thus we make the familiar decomposition

$$\vec{r}_i = \vec{R}_i^\alpha + \vec{u}_i. \tag{8}$$

The partition function (configurational weight) of structure α, at temperature T, may then be written as

$$Z(N,T,V,\alpha) = \prod_i \left[\int_{[\alpha]} d\vec{u}_i \right] e^{-\beta E(\{\vec{r}\})}. \tag{9}$$

Here $\int_{[\alpha]}$ signifies integration subject to an appropriate configurational constraint that picks out (from the entire configuration space of N particles in volume V) those configurations that "belong" to crystalline phase α [26]. The partition function defines the Helmholtz free-energy density through

$$f(N,T,V,\alpha) = -\frac{1}{N\beta} \ln Z(N,T,V,\alpha). \tag{10}$$

The relative stability of the two phases reflects the *difference* between their free energies, and thus the *ratio* of their partition functions [27]

$$\Delta f_{\alpha\tilde\alpha} \equiv f(N,T,V,\alpha) - f(N,T,V,\tilde\alpha) = -\frac{1}{N\beta} \ln \mathcal{R}_{\alpha\tilde\alpha}(N,T,V), \tag{11}$$

where

$$\mathcal{R}_{\alpha\tilde\alpha}(N,T,V) \equiv \frac{Z(N,T,V,\alpha)}{Z(N,T,V,\tilde\alpha)}. \tag{12}$$

If the difference between the densities of the two phases is small, the phase boundary follows immediately as the locus of points for which $\Delta f_{\alpha\tilde\alpha}$ vanishes (the two phases have equal statistical weights). In general, however, the problem is more conveniently formulated in the NTP ensemble. The relevant partition functions are then

$$\mathcal{Z}(N,T,P,\alpha) = \int dV Z(N,T,V,\alpha) e^{-\beta PV}. \tag{13}$$

The difference between the Gibbs free-energy densities follows as

$$\Delta g_{\alpha\tilde\alpha} \equiv g(N,T,P,\alpha) - g(N,T,P,\tilde\alpha)$$
$$= -\frac{1}{N\beta} \ln \mathcal{R}_{\alpha\tilde\alpha}(N,T,P), \tag{14}$$

where now

$$\mathcal{R}_{\alpha\tilde\alpha}(N,T,P) \equiv \frac{Z(N,T,P,\alpha)}{Z(N,T,P,\tilde\alpha)}. \tag{15}$$

The phase boundary is defined by the condition that

$$\mathcal{R}_{\alpha\tilde\alpha}(N,T,P) = 1, \tag{16}$$

in the thermodynamic ($N\to\infty$) limit.

IV. HARMONIC APPROXIMATION

The foundation stone of analytical approaches to the problem of free-energy determination for crystalline solids is the harmonic approximation. The methodology is long established [10]; many studies of LJ systems within the harmonic approximation (and its refinements) exist [28]. Here we use the harmonic framework to provide a benchmark of the LSMC work to be presented in following sections. We describe it in outline only.

In the harmonic approximation we expand the configurational energy [Eq. (6)] to second order in the displacements $\{\vec{u}\}$ [29]

$$E(\{\vec{r}\}) = E(\{\vec{R}\}) + \frac{1}{2} \sum_{\langle ij \rangle, \mu, \nu} K_{ij}^{\mu\nu} u_i^\mu u_j^\nu + \cdots. \tag{17}$$

Here u_i^μ and u_j^ν denote cartesian components (μ, ν) of the displacements $\vec{u}$ of the particles associated with sites i, j. The "dynamical" matrix K is defined by [30]

$$K_{ij}^{\mu\nu} = \begin{cases} \left[\dfrac{\phi_c'(R_{ij})}{R_{ij}} - \phi_c''(R_{ij}) \right] \hat{n}_{ij}^\mu \hat{n}_{ij}^\nu - \dfrac{\phi_c'(R_{ij})}{R_{ij}} \delta_{\mu\nu} & (j \neq i) \\ -\sum_{k \neq i} K_{ik}^{\mu\nu} & (j = i) \end{cases} \tag{18}$$

with

$$\hat{n}_{ij} \equiv \frac{\vec{R}_{ij}}{R_{ij}}. \tag{19}$$

As a result of the truncation scheme adopted in Eq. (6) the ground-state energy in Eq. (17) is essentially exact (in the limit of large enough N), while the fluctuation term contains contributions only from interactions falling within the (density-dependent) cutoff.

Within this approximation the NVT partition function [Eq. (9)] can be expressed in the form

$$Z(N,T,V,\alpha) = [2\pi/\beta]^{3N/2} \exp[-\beta E(\{\vec{R}\}_\alpha)][\det \mathbf{K}(\alpha)]^{-1/2}, \tag{20}$$

where $\mathbf{K}(\alpha)$ is the $3N \times 3N$ matrix whose elements are defined by Eq. (18) (applied to phase α). Then the fluctuation contribution to the free-energy density of phase α in the *harmonic approximation* is measured [31] by

$$f_\alpha^h \equiv \frac{1}{2N\beta} \ln[\det \mathbf{K}(\alpha)] = \frac{1}{2N\beta} \sum_{j=1}^{3(N-1)} \ln \lambda_j^\alpha, \tag{21}$$
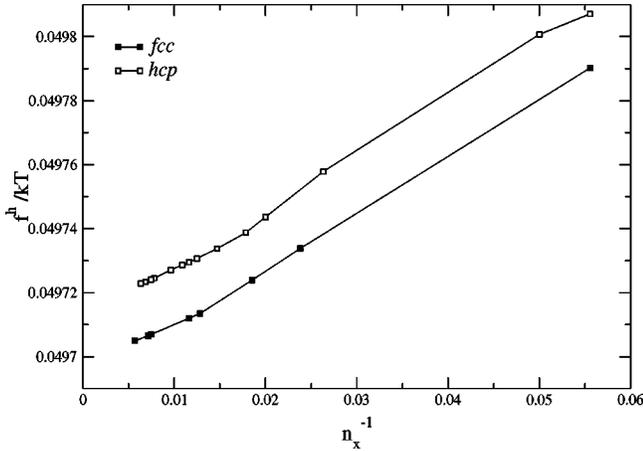
FIG. 2. The harmonic contributions to the free-energy density [Eq. (21); note [32]] for fcc and hcp structures with $N=216$, and $\rho\sigma^3=\sqrt{2}$. The results are expressed as functions of the number of interacting neighbors n_x [33].

where λ_j^α $[j=1,\ldots,3(N-1)]$ denote the nonzero eigenvalues of the matrix $\mathbf{K}(\alpha)$. The *difference* between the harmonic contributions to the free-energy densities of the two phases follows as

$$\Delta f^h_{\alpha\tilde{\alpha}}=f^h_\alpha-f^h_{\tilde{\alpha}}=\frac{1}{2N\beta}\ln\left[\frac{\det \mathbf{K}(\alpha)}{\det \mathbf{K}(\tilde{\alpha})}\right]. \qquad (22)$$

These expressions can be evaluated numerically with relative ease (for modest N values), allowing investigation of the influence of the potential truncation. Figure 2 shows f^h, as defined by Eq. (21) [32], for the two structures in systems of $N=216$ particles, at a density $\rho\sigma^3=\sqrt{2}$, for different choices of cutoff. We have chosen [33] to parameterize the cutoff by the inverse of the number n_x of interacting neighbors (the number of particles falling within the cutoff range). The free-energy "scale" on which to assess the significance of these results is set by the *difference* between the values for the two structures. On this scale the values of f^h for the two structures each evolve significantly between the two extremes. But the *difference* between these free energies varies relatively little: the values for the two extremes differ by only 2%. Figure 3 shows that this statement remains true over a wide range of densities, and for a substantially larger system size N. The conclusion we draw is that, in contrast to the *ground-state energy* (Sec. II, Fig. 1), the *fluctuation* contribution to the free-energy *difference* of interest in the *harmonic approximation* can be accurately captured by a model incorporating only relatively short-range interactions. It might seem reasonable to suppose that this state of affairs extends beyond the harmonic limit and we have indeed proceeded on this basis: for the most part our LSMC studies have focused on the truncated model defined in Eq. (7), with $r_x=1.5r_0$ and $n_x=18$. But the results presented below include a retrospective check on the consistency of this assumption.

Figure 3 also shows that the dependence of the free-energy difference upon system size N (finite-size effects) *is*
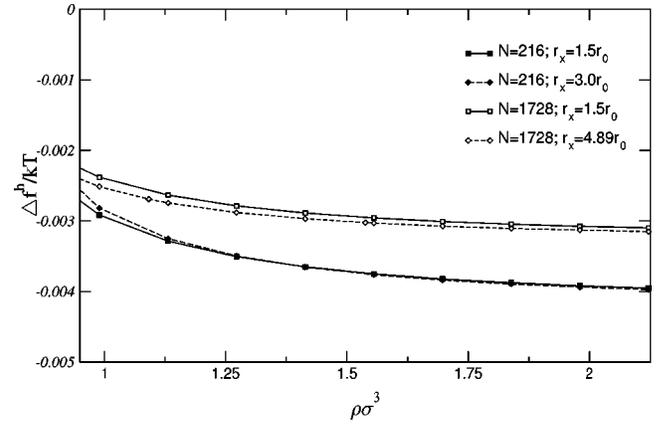
FIG. 3. The *difference* between the harmonic contributions to the free-energy density [Eq. (21)] of the two structures for system sizes $N=216$ and $N=1728$ as function of density. The solid and dashed lines correspond to different fixed-neighbor-list truncation schemes.

significant on the relevant scale: the harmonic free-energy difference changes by some 25% between $N=6^3=216$ and $N=12^3=1728$. Our direct matrix-diagonalization treatment is not readily extended to larger N values. But our results for the harmonic free-energy difference between the two structures for $N=12^3$ lie within 3% of the harmonic thermodynamic limit reported in [34]; and we find that a $1/N$ extrapolation of our $N=6^3$ and $N=12^3$ results is consistent with that limit. Accordingly our LSMC studies have also been largely restricted to these two system sizes.

V. LATTICE-SWITCH MC METHOD

The problem of phase behavior is succinctly expressed in Eqs. (12) or (15): the relative stability of (free-energy difference between) two phases is determined by the ratio of the associated partition functions. It is helpful (both conceptually and practically) to recognize that ratio for what it is: the ratio of the equilibrium probabilities with which a system, free to explore *both* phases, will visit each one. The LSMC method provides the strategy needed to captitalize on this identification. We have described the method fully elsewhere [2], in the context of a system of hard spheres. The description here will, therefore, be brief and will focus on the adaptations needed to accommodate soft interactions.

In principle, LSMC represents a simple and intuitive extension of the MC framework [6] conventionally used for exploration of a single (structural) phase. The "single-phase" MC method updates the particle displacements for fixed lattice vectors; LSMC augments this procedure by updates of the lattice vectors with the displacements held fixed. By "update" of the lattice vectors we mean replacing the entire set of lattice vectors describing the current structure with a set that describes the other structure. The current particle configuration, which "belongs" to the phase space associated with the current structure, is thus mapped onto a particle configuration that "belongs" to the phase space of the other structure.

Without further refinement this simple procedure will not

in general work: an equilibrium configuration of one structure will typically have as its conjugate (i.e., will *map into*) a relatively high-energy configuration of the other structure. The lattice switch MC step will typically be rejected. To make it work the sampling of the configurations of each structure must be multicanonically biased [3] to enhance the probabilities of those states (we call them gateway states) from which a switch can be attempted with a reasonable likelihood of success. The bias needs to reflect the *difference* between the excitation energies [35] of conjugate pairs. To that end we define an order parameter [36]

$$\mathcal{M} \equiv \mathcal{M}(\{\vec{u}\}) = \mathcal{E}(\{\vec{u}\}, \{\vec{R}\}_{\text{hcp}}) - \mathcal{E}(\{\vec{u}\}, \{\vec{R}\}_{\text{fcc}}), \quad (23)$$

where

$$\mathcal{E}(\{\vec{u}\}, \{\vec{R}\}_\alpha) = \beta[E(\{\vec{r}\}) - E(\{\vec{R}\}_\alpha)], \quad (24)$$

measures the excitation energy of configuration $\{\vec{r}\} \equiv \{\vec{u}\}, \alpha$ (the amount by which it exceeds the ground-state energy of structure α) in units of $kT = \beta^{-1}$. The energy difference defined in Eq. (23) is the analog of the "overlap order parameter" utilized in our studies of hard spheres [1,2].

Given the relationship between the energy of a configuration and the energy of its conjugate one sees that displacement patterns $\{\vec{u}\}$ drawn from equilibrium sampling of the fcc (hcp) configuration space will give predominantly positive (negative) values of $\{\mathcal{M}\}$. The gateway states lie in between [37]; in the thermodynamic limit they have vanishingly small *equilibrium* probability. We thus require to engineer a multicanonical sampling algorithm to enhance the probability along a notional line in $\mathcal{M}$ space, extending from the "equilibrium" $\mathcal{M}$ values (whose *magnitude* reflects the energy cost of a LS acting on a *typical* configuration) through to the small-$\mathcal{M}$ gateway configurations. This aim is realized by augmenting the system energy function Eq. (6) appropriately,

$$\beta E(\{\vec{r}\}) \rightarrow \beta E(\{\vec{r}\}) + \eta[\mathcal{M}] \equiv \tilde{E}(\{\vec{r}\}). \quad (25)$$

Here $\eta[\mathcal{M}]$ is a multicanonical weight function [3] that must be chosen so as to allow the system to access the gateway configurations. For this purpose it is sufficient (though not necessarily optimal) to construct the weights such that the multicanonical distribution $P(\mathcal{M}|N,T,\{\eta\}\dots)$ [38] is roughly flat. The desired ratio of partition functions [Eq. (12) or (15)], is identified with the ratio of the aggregate *canonical* probabilities of the two phases

$$\mathcal{R}_{\alpha\tilde{\alpha}}(N,T,\dots) \equiv \frac{P(\alpha|N,T,\dots)}{P(\tilde{\alpha}|N,T,\dots)} = \frac{\displaystyle\sum_{\{\mathcal{M}\}} P(\mathcal{M},\alpha|N,T,\dots)}{\displaystyle\sum_{\{\mathcal{M}\}} P(\mathcal{M},\tilde{\alpha}|N,T,\dots)} \quad (26)$$

and can be obtained from the corresponding (measured) *multicanonical* distributions by reweighting (folding out the weights)
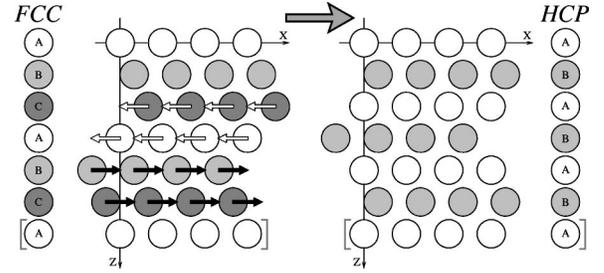
FIG. 4. The LS transformation applied to the ground-state configuration. The diagram shows six close-packed (x-y) layers. (The additional bracketed layer at the bottom is the periodic image of the layer at the top.) The circles show the boundaries of particles located at the sites of the two close-packed structures. In this realization of the fcc→hcp lattice switch, the top pair of planes are left unaltered, while the other pairs of planes are relocated by translations, specified by the black and white arrows.

$$\frac{\displaystyle\sum_{\{\mathcal{M}\}} P(\mathcal{M},\alpha|N,T,\dots)}{\displaystyle\sum_{\{\mathcal{M}\}} P(\mathcal{M},\tilde{\alpha}|N,T,\dots)}$$

$$= \frac{\displaystyle\sum_{\{\mathcal{M}\}} e^{\eta[\mathcal{M}]} P(\mathcal{M},\alpha|N,T,\{\eta\},\dots)}{\displaystyle\sum_{\{\mathcal{M}\}} e^{\eta[\mathcal{M}]} P(\mathcal{M},\tilde{\alpha}|N,T,\{\eta\},\dots)}. \quad (27)$$

The last step removes the *explicit* bias resulting from the weights, leaving only the desired legacy [39]—that both phases are visited with the canonical probabilities. Notice that, in contrast to the hard-sphere problem, where $\mathcal{M}$ serves as a unique identifier of the phase, the relevant sums in Eqs. (26) and (27) involve the phase labels explicitly.

VI. LATTICE-SWITCH MC IMPLEMENTATION DETAILS

Our MC procedure involves stochastic sampling of three classes of variable: the displacement vectors $\{\vec{u}\}$; the volume V; and the set of lattice vectors $\{\vec{R}\}_\alpha$ indexed, collectively, by the lattice label α. The sampling of $\{\vec{u}\}$ and V follows standard prescriptions [6], with the effective configurational energy (controlling metropolis acceptance probabilities) given by Eq. (25). Stochastic updating of the lattice label α (the lattice switch) is discussed extensively in [2]. We summarize the essentials.

There are many ways of defining the switch transformation (lattice-to-lattice mapping); for the fcc-hcp problem the prescription illustrated in Fig. 4 is simple and efficient.

In the cases in which the two phases related by the switch have significantly different volumes the switch operation can accommodate an appropriate dilation [40]. In this case we did not find that necessary: thus, as implemented here, the switch transformation preserves both V and $\{\vec{u}\}$ [41,42]. With the prescription adopted in Eq. (23) the switch *also* leaves $\mathcal{M}$ (and thus $\eta[\mathcal{M}]$) unchanged [43]. The switch
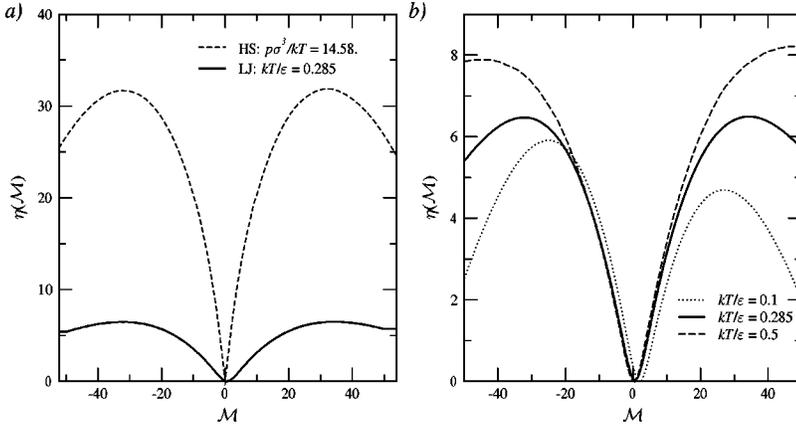
FIG. 5. Examples of ''best-estimate'' weight functions for systems of $N=6^3$ particles. (a) The Lennard-Jones (LJ) weight function (for $p=0,kT/\epsilon=0.285$) is compared with the weight function for hard spheres (HS; see [2]) at a pressure chosen such that the densities of the two systems are comparable. (b) Weight functions for the $p=0$ LJ system at three different temperatures.

acceptance probability is then simply

$$P_a(\alpha\rightarrow\tilde{\alpha})=\min[1,\exp(-\Delta\tilde{E})], \quad (28)$$

where [44]

$$\Delta\tilde{E}\equiv\beta[E(\{\vec{u}\},\{\vec{R}\}_{\tilde{\alpha}})-E(\{\vec{u}\},\{\vec{R}\}_{\alpha})], \quad (29)$$

is the energy cost of the switch.

The practical core of the whole procedure lies in the determination of the multicanonical weight function $\eta[\mathcal{M}]$. In generic terms the task is simple: the weight function is parameterized by a discrete set of weights each associated with an interval of $\mathcal{M}$ space of width m_0 (that we set to a constant value, $m_0=1$) and together extending across the relevant range of $\mathcal{M}$ space (the values appropriate for the equilibrium structures); the weights are refined iteratively until the sampling distribution is sufficiently close to the multicanonical ideal. Details of the techniques we used are described elsewhere [45].

VII. LATTICE-SWITCH MC RESULTS

A. Features of the LS operation

We begin by presenting a selection of results chosen to illustrate features of the LS operation itself; results for the quantities of more immediate physical interest appear in the subsections that follow.

Figure 5 provide examples of weight functions $\eta[\mathcal{M}]$ constructed such that the associated *weighted* distribution $P(\mathcal{M}|\eta[\mathcal{M}])$ is effectively ''flat'' (the multicanonical ideal [3]). To the extent that it *is* flat one can regard the weight function as a logarithmic measure of the *unweighted* (equilibrium) distribution $P(\mathcal{M})$. The peaks in the weight functions thus locate the ''equilibrium'' values of $\mathcal{M}$ for each phase, the peaks at positive (negative) $\mathcal{M}$ corresponding to fcc (hcp) phases respectively. The heights of these peaks measure, on a logarithmic scale, the size of the free-energy barrier to be negotiated for an escape from that phase, along the route offered by LS [46] .

Figure 5(a) compares the results for the soft LJ potential of primary interest here, with the results for a system of the same number of particles interacting as hard spheres [2]. The thermodynamic parameters were chosen so that the two sys-

tems have essentially the same densities. It is clear that ''softening'' the potential greatly reduces the barriers against switching [47].

Figure 5(b) shows the weights for the (zero-pressure) LJ system for a range of temperatures. The ''barrier heights'' grow with increasing temperature reflecting the growing energy mismatch between typical configurations and their conjugates.

The value of $\mathcal{M}$ for which $\eta[\mathcal{M}]$ assumes its minimum (zero [46]) value—corresponding to the macrostates of lowest probability in the equilibrium ensemble—identifies the gateway macrostates. This identification is apparent operationally in Fig. 6 that shows that the probability of a switch being accepted is maximal in the region where $\eta[\mathcal{M}]$ has its minimum. Though this identification is not self-evidently true, it is not surprising: the defining constraint on the ''gateway'' macrostates (namely that they should allow a switch) is sufficiently tight to ensure that they have the lowest probability of all the macrostates that it is *necessary* to sample.

Close scrutiny of Fig. 5(b) shows that the $\mathcal{M}$ value $\mathcal{M}_G$, locating the gateway states is nonzero and shifts with temperature. To understand this one should recall that the gate-
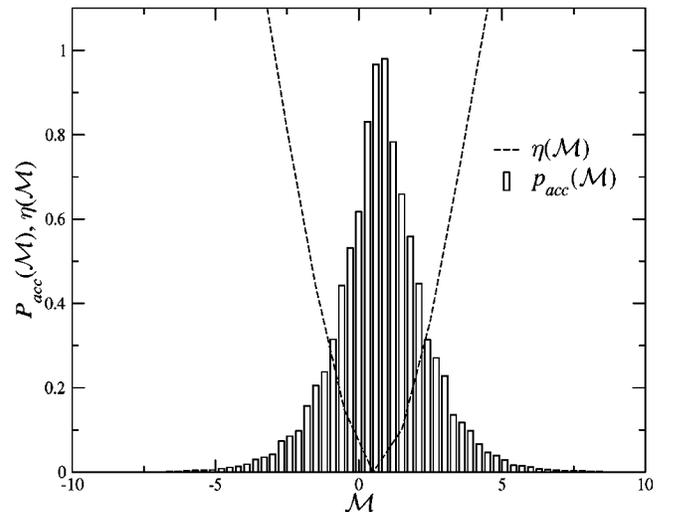
FIG. 6. The multicanonical weights η (solid line) and the probability of acceptance of a switch $p_{\rm acc}$ as a function of $\mathcal{M}$ in the region near $\mathcal{M}=0$. The system parameters are $N=6^3$, $p=0$, and $kT/\epsilon=0.285$.
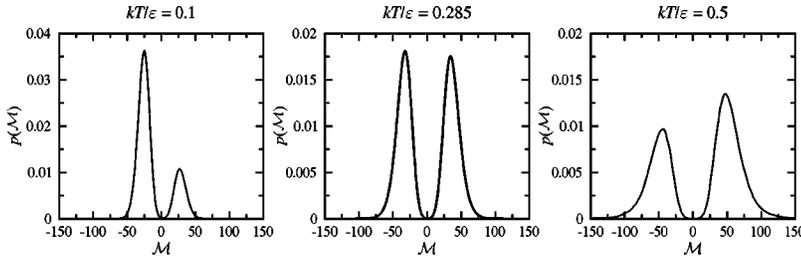
FIG. 7. Temperature dependence of the probability distribution $P(\mathcal{M})$ for a system of $N = 6^3$ Lennard-Jones particles at zero pressure.

way states are such that the switch energy cost $\Delta \widetilde{E}$ [Eq. (29)] is close to zero [48]. Recalling the definition of $\mathcal{M}$ we see that the condition $\Delta \widetilde{E} = 0$ implies $\mathcal{M} = \mathcal{M}_G = N\beta\Delta\epsilon^0_{\mathrm{fcc, hcp}}$: the behavior of $\mathcal{M}_G$ can thus be traced to the difference in ground-state energies.

Figure 7 shows the equilibrium $\mathcal{M}$ distributions for the zero-pressure LJ system at the three different temperatures featuring in Fig. 5(b): in each case we sample the multicanonical distributions defined by the appropriate set of weights (allowing *both* phases to be sampled) and then re-weight to obtain the equilibrium distribution. Equation (15) shows that the relative stability of the two phases is directly reflected in the ratio of the integrated areas of the peaks in the equilibrium $\mathcal{M}$ distribution. We thus see immediately and transparently from Fig. 7 that, while the hcp structure is favored at low enough temperature (as implied by the difference between the ground-state energies at this pressure, Fig. 1), the favored phase at high enough temperatures is fcc.

B. Benchmarking: LSMC and the harmonic approximation

A comparison between the results based on harmonic crystal dynamics (Sec. IV) and LSMC is worthwhile for two reasons. In the regime of low enough T, the harmonic theory is necessarily correct: comparison with LSMC results then provides a check on the overall consistency of the LS framework. In the regime of higher T, we must expect the harmonic theory to be growingly inaccurate: comparison with LSMC allows us to explore the extent of these inaccuracies.

Figure 8 shows the results of LSMC studies in the NVT ensemble, at two different densities, for a range of temperatures. In each case the solid line is the free-energy difference in the harmonic approximation, which combines the ground-state-energy difference [Eq. (5)] and the harmonic contribution to the free energy [Eq. (22)]:

$$\frac{\Delta f}{\epsilon} = \frac{\Delta \epsilon^0}{\epsilon} + \frac{\Delta f^h}{kT} \frac{kT}{\epsilon} \quad \text{(harmonic).} \quad (30)$$

The LS values follow from simulations determining the $\mathcal{M}$ distributions for the appropriate state point and applying Eq. (11). Evidently, the two methods agree at low enough temperatures, where the harmonic theory is trustworthy; this provides a reassuring consistency check. As we have already noted, the results of Ref. [34] provide consistency checks on the harmonic studies themselves.

The differences between the results of the two methods should, then simply reflect the anharmonic effects that are captured by the LS studies. The differences grow with increasing temperature and decrease with increasing density; this behavior is consistent with simple qualitative expectations [49].

At *both* densities the deviations from harmonic behavior are such as to favor hcp: we see no change in the sign of the "anharmonic" contributions, with increasing density, suggested by the perturbation-theory calculations in [18].

At $kT/\epsilon = 0.4$ and $\rho\sigma^3 = 1.092$, we find $\Delta f/\epsilon = -0.000\,074(16)$, which is virtually indistinguishable from the predictions of the harmonic theory. In contrast Ref. [50] (which uses an integration method to determine the free energies of each phase separately) reports $\Delta f/\epsilon = -0.066(4)$ at $kT/\epsilon = 0.4$ and $\rho\sigma^3 = 1$; this seems irreconcilable with our results.

At the temperatures of most immediate physical interest (where $\Delta f = 0$) the discrepancies between the two methods are small, even for the zero-pressure density. We may thus expect that the harmonic theory of the fcc-hcp phase boundary will not prove far wrong; this is indeed the case.

C. The fcc-hcp phase boundary

We have established that LSMC provides us with a way of determining the difference between the free energies of the two phases at *some* point in the space of thermodynamic coordinates. We must now turn to the practical task of interest here: the determination of the phase behavior itself. This task may be divided into two—extrapolation and tracking. First ("extrapolation") we need to use the data accumulated
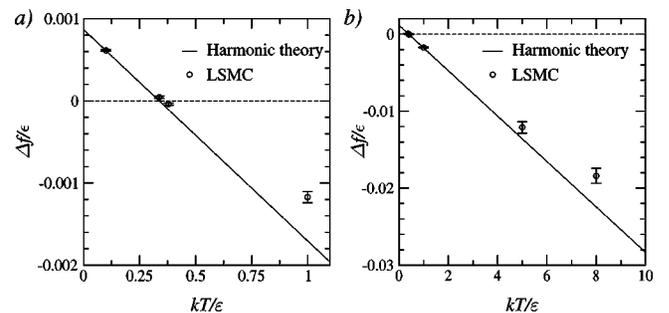
FIG. 8. The difference [21] between the Helmholtz free-energy densities of the two structures for a system of $N = 12^3$ particles at a density of (a) $\rho\sigma^3 = 1.092$ close to the zero-pressure density and (b) $\rho\sigma^3 = 1.538$ the density at which $\Delta\epsilon^0_{\mathrm{fcc, hcp}}$ is maximal. In each case the solid line represents the results of a harmonic calculation, which follows Eq. (30) with intercepts and gradients given by $\Delta\epsilon^0/\epsilon = 0.000\,869\,8$ [0.001 152] and $\Delta f^h/kT = -0.002\,574$ [$-0.002\,949$] for cases (a) [(b)] respectively. The data points are the results of LSMC studies utilizing Eq. (11).
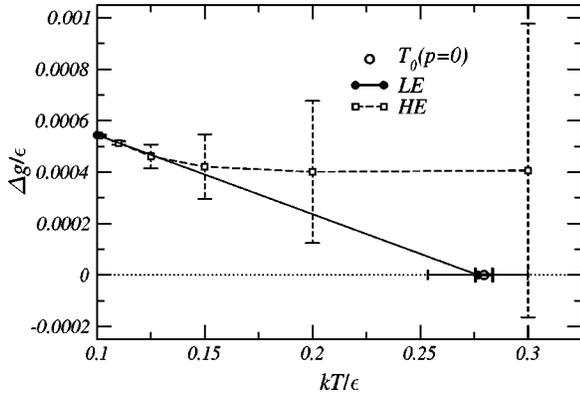
FIG. 9. Comparison of linear extrapolation (LE) and histogram extrapolation (HE) methods for determining the $p=0$ hcp-fcc phase boundary temperature $T_0(p=0)$ on the basis of simulation data accumulated at $kT/\epsilon=0.1$ and $p=0$ for a system of $N=6^3$ Lennard-Jones particles. Our best estimate of $T_0(p=0)$ is also shown.

at our initially chosen state point to infer the location of *some* point on the phase boundary (and refine it as necessary). Second ("tracking") we need to map out the phase boundary emanating from that point. There are many ways of addressing these tasks; we shall focus on the strategies we have found most useful, while noting other approaches that proved less satisfactory here.

Both extrapolation and tracking stages make use of the fact that the *derivatives* of the free-energy difference can be expressed in terms of the differences between single-phase averages of appropriate observables

$$\frac{\partial[\beta\Delta g]}{\partial\beta}=\frac{1}{N}\langle\Delta E+p\Delta V\rangle=\Delta h \qquad (31a)$$

and

$$\frac{\partial[\beta\Delta g]}{\partial p}=\frac{1}{N}\langle\Delta V\rangle=\Delta v, \qquad (31b)$$

where Δh and Δv represent the differences between the enthalpy per particle and mean volume per particle, for the two phases. Since the densities of fcc and hcp phases are very similar at coexistence [51] the p derivative [Eq. (31b)] is hard to determine accurately. But the β derivative [Eq. (31a)] is readily measurable to the required precision, using the appropriate single-phase averages already available in the LS simulations used to determine Δg itself. Given Δg and its β derivative a simple linear extrapolation (LE) scheme provides an estimate of the phase boundary temperature for the chosen pressure. Figure 9 provides an example. The extrapolation works well—a reflection of the fact that the behavior is close to harmonic over the region involved. Figure 9 also shows the results generated by a single histogram extrapolation (HE) [52] in which the $\mathcal{M}$ distribution measured through the histogram accumulated at the initial state point is used to estimate the distribution at other temperatures. In this case HE fares much worse than LE. The reason is that the HE predictions for some extrapolated temperature β explic-
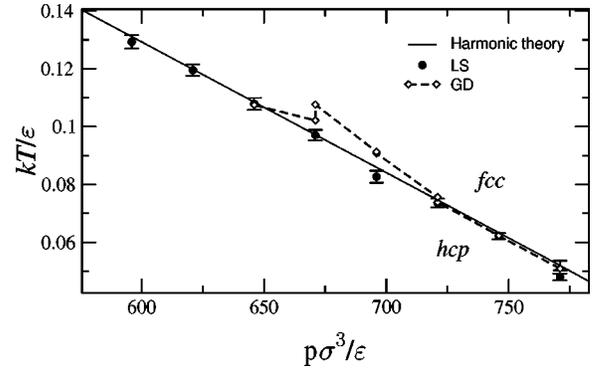
FIG. 10. Comparison of lattice-switch (LS) and Gibbs-Duhem (GD) methods for mapping out a high-pressure section of the phase boundary for a system of $N=6^3$ Lennard-Jones particles. The results of harmonic theory are also shown.

itly utilize the statistical information that the simulations (performed at β_0) give about the dominant configurations at β; this information typically resides in the poorly sampled wings of the measured distribution. By contrast LE implicitly *models* the measured distribution (in this case) as two Gaussians, each centered on the average for the phase, whose characteristics are then determined by the well-sampled regions of the measured distribution. LE does well when (as here) the near-Gaussian assumption is good.

Once a point on the phase boundary is established the rest of the boundary can, in principle, be mapped by numerical integration of the phase boundary derivative that follows from Eqs. (31a) and (31b). This strategy—Gibbs-Duhem (GD) integration—has been widely and effectively used [53]. In this context, however, it is undermined by the problem we have already noted with the measurement of the volume difference at coexistence. The resulting inaccuracies are compounded by the generic draw back of GD—the absence of any self-correction mechanism. Figure 10 shows the results of a low-order predict-evaluate-correct (PEC) GD integration: on its own GD is clearly not suited to the present problem. Instead we chose a strategy which, while making use of GD to the extent it is trustworthy, employs LS to correct it. Starting at some coexistence point p_1,β_1 say we choose a different pressure $p_2=p_1+\Delta p$. We use the GD derivative to predict the associated transition temperature β_2^e. If the difference between β_2^e and β_1 is not statistically significant (given the errors associated with the GD derivative) we reset $\beta_2^e=\beta_1$. The multicanonical weights used at β_1,p_1 provide estimates of the weights at β_2^e,p_2 that can be easily refined. The process is completed by a β extrapolation to the phase boundary (as discussed above). Sustaining switching between phases as the phase boundary is traced out delivers a running consistency check. In this case the computational cost is actually *less* than it would be to determine the GD derivative to the accuracy required to make GD itself viable. Figure 10 shows some representative results.

Using these techniques we have mapped out the hcp-fcc phase boundary from the zero-pressure limit, through to the maximum pressure at which hcp is classically stable at $T=0$. The results are gathered together in Fig. 11. In the case
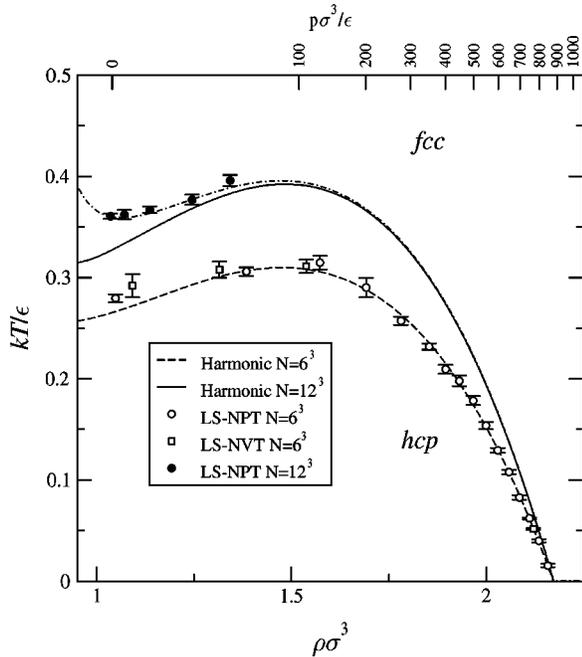
FIG. 11. A variety of approximations to the classical Lennard-Jones phase diagram. In all cases the difference between the ground-state energies of the two structures is essentially exact; the fluctuation contribution to the energy is truncated according to Eq. (7), with $r_x = 1.5 r_0$. The dashed and solid lines are the results of harmonic calculations (for the two system sizes). The dash-dotted line is a phenomenological parameterization of the anharmonic effects [54]. The scale at the top of the figure shows the pressures at selected points on the (LS $N = 12^3$, NPT) coexistence curve. Tie-line structure is unresolvable on the scale of the figure.

of the smaller ($N = 6^3$) system we gathered LS data along the entire phase boundary, demonstrating agreement with the results of harmonic theory except in a low-pressure region where the break-down of the harmonic approximation begins to be apparent. In the case of the larger ($N = 12^3$) system we focused our LS studies on this low pressure region, mapping out the phase boundary from $p = 0$ through to the pressure at which it becomes indistinguishable from the predictions of the harmonic theory. The dashed line through the LS points is based on a simple phenomenology of the growth of anharmonic effects with decreasing density [49,54]. While the deviations from the harmonic (NVT) theory are certainly resolvable, they are always small. And indeed the same general structure and scale of the phase diagram is apparent in pioneering (if semiquantitative) harmonic lattice dynamics studies of 50 years ago [24]. It is harder to make useful connections with more recent studies. Jackson and Swol [55] and Choi *et al.* [17] both appeal to perturbation theories using a truncated potential giving a complex phase behavior. In [18] the same method is applied to a system without truncation; but the perturbation theory presupposes hard-sphere results now known to be incorrect, and is in any case restricted to high temperatures ($kT/\epsilon \gtrsim 1$).

VIII. DISCUSSION

We divide this concluding section into three. First we shall take stock of where we stand on the *generic* problem

addressed here: the development of a general computational technique for the study of the phase behavior of crystalline matter. Then we shall summarize what we have learned about the *particular* problem we have chosen to address: the equilibrium phase behavior of the LJ model. Finally we shall consider the implications of the LJ model predictions for studies of the crystallization process itself.

1. The technique

To assess the LS technique we must now set it in context.

The task of determining the phase behavior of solids has traditionally been seen as requiring *separate* calculations of the *absolute* free energies of the competing phases, linking each to some reference system. This approach was pioneered by Frenkel and Ladd [5] and has been implemented in a variety of forms, differing according to the choice of reference system (Einstein solid, harmonic solid, ideal gas, . . .) and how the linking path is negotiated (integration, extended sampling) [50,56–59]. The problem facing this strategy is that the quantity of physical interest—the free-energy difference between the phases—is typically small (in the vicinity of a phase boundary necessarily so) compared to the absolute free energies of the separate phases; one is left with the task of determining the difference between two ''relatively large'' numbers [60] that are *themselves* of little or no physical interest. It is possible to make this strategy deliver results of comparable accuracy to that of LS, for comparable computational resource [61]. But the extent of the signal-to-noise problem provides strong motivation for developing a more direct strategy.

The essence of a ''direct strategy'' is that it should focus on the free-energy *difference* between (relative likelihood of) the two phases: this requires a ''path'' connecting the disjoint configuration spaces of the two phases. The pioneering study here is that of Rahman and Jacucci [62]. Subsequent studies in this vein can be differentiated according to whether (in the terminology of [62]) the path is ''physical'' or ''nonphysical'' and according to how the path is traversed [63–65]. Reference [65], devoted to the hcp-fcc phase behavior of hard spheres, provides the most recent instance of the use of a ''physical'' path between the two structures [66], explored with multicanonical sampling [3]. In that case the ''physical path'' was found to encounter some stubbornly physical barriers that (the authors report) made it less efficient than a LS strategy that they also explored.

The LS technique belongs to the class of ''direct'' methods that exploit a ''nonphysical'' path. One can see it [67] as a reformulation of the strategy used by Moody *et al.* [63] in early (though numerically inconclusive) studies of polymorphic crystalline behavior. The principal additional feature of LS is that we have used extended (''multicanonical'') sampling instead of multistage sampling; that change, combined with the availability of massively more computational resource produces a technique that—we have seen here and in [2]—delivers the precision one requires, with readily quantifiable uncertainties, in a physically transparent fashion.

There are two noteworthy ways in which the LS framework could be fruitfully developed. First, as noted in [2], one may fold into the switch operation any desired transforma-

tion of the displacement coordinates, provided one takes appropriate account (through the Jacobian) of the resulting warping of the configuration space [68]. In particular one can define a switch operation that preserves *collective* (normal mode) coordinates rather than *local* coordinates. In the *harmonic* region such an operation will work *without the need for any weighting*; the barrier (to be surmounted by multicanonical methods) is reduced to that associated with the anharmonic contributions. This approach is currently under investigation.

Second, the current framework is strictly classical. It is clearly desirable to extend the framework to include quantum effects; we plan to do so within a path integral formalism [69].

2. The system

We turn now to the particular system (the LJ model) we chose for this exploratory study. From some standpoints that choice proved ill advised; while LJ is regarded as the generic soft potential, its phase behavior is acutely sensitive to details (notably the range of interactions permitted) that are largely arbitrary adjuncts of the phenomenology. The diversity of LJ phase diagrams featuring in the literature is testimony to this sensitivity. Should one wish to take the LJ potential at face value and establish a definitive phase diagram one must certainly investigate the effects of potential-range truncation and (other) finite-size effects systematically. We have attempted to do so here: we have dealt exactly with the ground-state energies; we have seen that our truncation scheme for the excitation energy captures the thermodynamically limiting behavior well in the harmonic regime (Figs 2, 3); and we find no evidence to doubt it when one moves beyond the harmonic region [70]. We have some confidence then that Fig. 11 does indeed capture the classical LJ crystalline phase behavior.

3. The implications

If the calculations reported here are to represent more than an exercise in equilibrium statistical mechanics then their results need to be related to studies of the crystallization process itself. It is not easy to do so in any satisfying way.

If one tries to make contact with observations on RGS one is faced with a whole range of effects that are potentially important. We have already noted the need to extend the framework to include quantum effects—clearly important for the lighter RGS and essential in understanding the behavior of Helium. But the link with real systems is also complicated by two other manifestations of the fcc-hcp knife edge on which the LJ model sits. That knife edge is presumably reflected in the sensitivity of the observed structure to impurities [71]; it also means that effects that are missing from the two-body central force LJ phenomenology (and that are not crucial to other aspects of the RGS behavior) need to be addressed [14]. As a more limited objective one may try to make contact with *simulations* of the crystallization process in LJ systems. There have been many studies of this kind. Much of the focus has been on the initial stages of the process—in particular, the suggestion [72] that the route to the thermodynamically favored phase traverses a region with a bcc structure. This phenomenon (for which there is simulation support [73]) is a useful reminder of the importance of the kinetic pathways by which the equilibrium state is reached, in the long term. The evidence from the most detailed study known to us [74] is that a fcc structure emerges in that "long term" [75]. But this study utilized a LJ potential truncated at $r_x = 2.3\sigma$; and we have seen that such details matter as regards the "true" equilibrium behavior. In any case, given the LJ knife edge, it seems likely that the particular long term behavior that does emerge may well have more to do with the kinetics than the thermodynamics. That interplay is probably better explored in a system other than LJ. Recent work has begun to address this interesting area [76].

---

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

[2] A. D. Bruce, A. N. Jackson, G. J. Ackland, and N. B. Wilding, Phys. Rev. E **61**, 906 (2000).

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

[4] 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.

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

[6] D. Frenkel and B. Smit, *Understanding Molecular Simulation* (Academic, New York, 1996).

[7] N. B. Wilding and A. D. Bruce, Phys. Rev. Lett. **85**, 5138 (2000).

[8] In the case of hard spheres the temperature enters only trivially through the scale of the pressure.

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

[10] T. H. K. Barron and M. L. Klein, in *Dynamical Properties of Solids*, edited by G. K. Horton and A. A. Maradudin (North-Holland, Amsterdam, 1974), p. 391.

[11] J. A. Barker, in *Rare Gas Solids*, edited by M. L. Klein and J. A. Venables (Academic Press, London, 1976), Vol. 1, p. 212.

[12] One should not expect it to work at sufficiently extreme pressures and it does not: see for example, R. J. Hemley and N. W. Ashcroft, Phys. Today **51** (8), 26 (1998).

[13] Helium, of course, is different, see e.g., H. R. Glyde, *Rare Gas Solids* (Ref. [11]), p 382.

[14] See e.g., K. F. Niebel and J. A. Venables, in *Rare Gas Solids* (Ref. [11]), p. 558.

[15] B. J. Alder and P. H. Paulson, J. Chem. Phys. **43**, 4172 (1965).

[16] Compare Refs. [17,18]; see also Fig. 1.

[17] Y. Choi, T. Ree, and F. H. Ree, J. Chem. Phys. **95**, 7548 (1991).

[18] Y. Choi, T. Ree, and F. H. Ree, J. Chem. Phys. **99**, 9917 (1993).

[19] We use the term "lattice" rather loosely. The set of vectors $\{\vec{R}\}$

identifies the orthodox lattice convoluted with the orthodox basis. For the purposes of LSMC $\{\vec{R}\}_\alpha$ merely identifies one particular configuration of structure α. The particular simplicity (perfect periodicity) of this configuration is not essential to the LSMC procedure; the ''lattice switch'' is but a simple case of a more general strategy that switches between configurations of different phases [7].

[20] We shall frequently suppress the α label when no ambiguity results.

[21] Whenever we refer to the ''difference'' between the (free) energies of the two structures we shall mean, specifically, the value for fcc minus the value for hcp.

[22] The relevant lattice sums are defined in, e.g., [23]. The results generated and used here are in agreement with those tabulated in [24].

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

[24] T. H. K. Barron and C. Domb, Proc. R. Soc. London, Ser. A **227**, 447 (1955).

[25] The ground-state energies were computed as parameterized functions of density, in the limits of large n_x and N; these functions were then used in subsequent simulations.

[26] While this constraint deserves some thought [2] it is implemented naturally and automatically within a MC framework.

[27] In the present work (in contrast to our studies of hard spheres [1,2,7]) we shall display factors of $\beta = [kT]^{-1}$ explicitly.

[28] See, for example, Ref. [24]; D. J. Lacks and G. C. Rutledge, J. Chem. Phys. **101**, 9961 (1994); R. G. Della Valle and E. Venuti, Phys. Rev. B **58**, 206 (1998).

[29] For the purposes of the harmonic theory the ''lattice sites'' are assumed to be positions of classical equilibrium for the prevailing mechanical constraints.

[30] See e.g., A. A. Maradudin, in *Dynamical Properties of Solids*, edited by G. K. Horton and A. A. Maradudin (North-Holland, Amsterdam, 1974), p. 18.

[31] More specifically f^h is the harmonic contribution to the free-energy density to within terms that do not depend upon the phase label α and are thus irrelevant to the free-energy *difference* of interest.

[32] The origin of the ordinate scale in Fig. 2 is fixed by implementing Eq. (21) with the eigenvalues λ_j^α expressed in units of ϵ/σ^2.

[33] If one forms the abscissa for Fig. 2 from r_x rather than n_x each f^h function shows discontinuities at those r_x values where n_x changes; the difference between the two functions shows discontinuities where *either* n_x value change (such as those seen in the plot of the ground-state energy difference in Fig. 1). Parameterizing the cutoff by n_x, and interpolating between the values associated with the observable discrete values of this variable, reveals essentially smooth underlying behavior.

[34] Z. W. Salsburg and D. A. Huckaby, J. Comput. Phys. **7**, 489 (1971) evaluate the harmonic contribution to the free energies of each phase, at the densities that minimize their respective ground-state energies, using a potential including first- and second-neighbor interactions.

[35] In simulations conducted in the *NVT* ensemble it is convenient to factor out the difference between ground-state energies $\Delta E^0 \equiv N\Delta\epsilon^0$ (which is then *independent* of configuration) and define $\mathcal{M}$ in terms of the excitation energies alone. In the *NPT*

ensemble the ground-state energy difference is configuration dependent (varies with density) and there is less advantage to making this separation.

[36] The sign convention adopted here has no deep significance.

[37] They lie near (but not at) $\mathcal{M} = 0$, see Sec. VII A.

[38] Here . . . identifies the appropriate ensemble variable *V* or *P*.

[39] One may see the use of multicanonical weights (sampling with them; then folding them out) as a way of ''simply improving'' (albeit infinitely) the statistical precision of an unbiased estimator of the desired ratio.

[40] See, for example, the study of the freezing transition in [7].

[41] There are other possibilities to which we return in Sec. VIII.

[42] In our studies of hard spheres [2] we treated the ratio of the unit cell parameters *c* and *a* for the hcp structure as a simulation variable. We were unable to resolve any departure from the ideal close-packed limit $c/a = \sqrt{\frac{8}{3}}$. We did not reopen this issue in the present studies: the c/a ratio is fixed throughout at its ideal value.

[43] If the switch is required to accommodate a significant change in some quantity other than the lattice type (e.g., the volume) the order parameter needs to be defined so as to carry an explicit phase label (and it becomes less appropriate to refer to it as an ''order parameter''). The weight function $\eta[\mathcal{M}]$ then splits into two distinct branches $\eta_\alpha[\mathcal{M}]$ and the switch involves a discontinuous change in the weight; this discontinuity can be tuned to cancel the typical energy cost of the change in the ''other'' quantity.

[44] By $E(\{\vec{u}\}, \{\vec{R}\}_\alpha)$ we mean the total configurational energy $E(\{\vec{r}\})$ reparameterized (for a configuration of phase α) by $\{\vec{u}\}, \alpha$.

[45] G. R. Smith and A. D. Bruce, J. Phys. A **28**, 6623 (1995); Phys. Rev. E **53**, 6530 (1996).

[46] The multicanonical weights are defined only to within an arbitrary additive constant; we fix that constant by the convention that the macrostate of *lowest* probability has *zero* weight. The weights may then be regarded as logarithmic measures of the macrostate probabilities *relative* to that of the lowest-probability macrostate. The height of the ''peak'' in the weight function for a given phase then measures the probability barrier to be surmounted, starting from the equilibrium states associated with that phase.

[47] Indeed for the LJ system [of the size, and under the conditions defined in Fig. 5(a)] switches are occasionally possible even *without* the aid of multicanonical weighting; in the corresponding HS system such occurrences are some 10^{11} times less frequent.

[48] Switches could also be launched successfully from states for which $\Delta\tilde{E}$ is *negative* but the scheme implemented here does not need to locate such states and does not do so.

[49] See for example R. J. Hardy, J. Geophys. Res., [Space Phys.] **85**, 7011 (1980).

[50] S. Somasi, B. Khomani, and R. Lovett, J. Chem. Phys. **113**, 4320 (2000).

[51] The (fractional) density difference at the phase boundary is at most 2×10^{-4}.

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

[53] D. A. Kofke, J. Chem. Phys. **98**, 4149 (1993).

[54] An adaptation of an argument of Ref. [49] suggests that the phase boundary temperature at density ρ will differ from the harmonic prediction by an amount $\Delta T(\rho)$ that varies as $[\phi''']^2/[\phi'']^3$ where the derivatives (of the LJ potential ϕ) are evaluated at the nearest-neighbor separation for the chosen density. This is the source of the dashed-dotted line in Fig. 11.

[55] G. Jackson and F. van Swol, Mol. Phys. **65**, 161 (1988).

[56] K. Singer and W. Smith, Chem. Phys. Lett. **140**, 406 (1987).

[57] A. R. Philpott and J. M. Rickman, J. Chem. Phys. **94**, 1454 (1991).

[58] J. M. Rickman and A. R. Philpott, J. Chem. Phys. **95**, 7562 (1991).

[59] S. Y. Sheu, C. Y. Mou, and R. Lovett, Phys. Rev. E **51**, R3795 (1995).

[60] As an (admittedly extreme) example, the difference between the LJ fcc and hcp free energies at (say) $kT/\epsilon = 0.5$ and $p = 0$ is some four orders of magnitude smaller than the absolute (excess) free energies of the two phases.

[61] S. Pronk and D. Frenkel, J. Chem. Phys. **110**, 4589 (1999) used the integration methods of Ref. [5] to determine the ''absolute'' entropies of hcp and fcc phases of hard spheres and thence the difference between these entropies, which is smaller by some four orders of magnitude. They also performed studies using the LS method: they report comparable accuracy for comparable computational resource.

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

[63] M. C. Moody, J. R. Ray, and A. Rahman, J. Chem. Phys. **84**, 1795 (1985).

[64] B. Kuchta and R. D. Etters, Phys. Rev. B **47**, 14 691 (1993).

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

[66] We refer to the ''shear implementation'' of Ref. [65]. In this case the ''physical'' path comprises configurations generated by a gradual shear (mutual translation) of lattice planes.

[67] The similarities between LS and the work of Ref. [63]—not so apparent in the LS formulation appropriate for hard spheres—become more obvious when LS is reformulated to deal with ''soft potentials.''

[68] C. Jarzynski (unpublished), has provided a general formulation of this idea.

[69] D. M. Ceperley and E. L. Pollock, in *Monte Carlo Methods in Theoretical Physics*, edited by S. Caracciolo and A. Fabrocini (ETS, Pisa, 1991).

[70] By switching between ensembles with truncated and untruncated potentials we investigated the effects of the truncation of the potential [Eq. (7)] on *anharmonic* effects. As one might expect these effects are similar in scale to that apparent in the harmonic results summarized in Ref. [2]. We found that the consequences for the phase diagram were small compared to the statistical uncertainties shown in Fig. 11.

[71] For example, Argon containing small concentrations of oxygen crystallizes as hcp: C. S. Barrett and L. Meyer, J. Chem. Phys. **42**, 107 (1965); see also T. Bricheno and J. A. Venables, J. Phys. C **9**, 4095 (1975). Note that (in contrast to the implications of Fig. 11) these results suggest that the stability of hcp relative to fcc is enhanced by *raising* the temperature.

[72] S. Alexander and J. P. McTague, Phys. Rev. Lett. **41**, 702 (1978).

[73] P. R. tenWolde, M. J. Ruiz-Montero, and D. Frenkel, Phys. Rev. Lett. **75**, 2714 (1995).

[74] W. C. Swope and H. C. Andersen, Phys. Rev. B **41**, 7042 (1990).

[75] Following quenches to $kT/\epsilon \simeq 0.45$, which is close to the fcc-hcp phase boundary suggested in Fig. 11.

[76] S. Auer and D. Frenkel, Nature (London) **409**, 1020 (2001).