Multicanonical Monte Carlo study of solid-solid phase coexistence in a model colloid

G. R. Smith* and A. D. Bruce
Department of Physics and Astronomy, The University of Edinburgh, Edinburgh EH9 3JZ, Scotland, United Kingdom
(Received 18 December 1995)

We describe a Monte Carlo approach to the determination of the relative stability of two phases, which is conceptually direct, potentially rather general, and particularly well suited to parallel computers. The approach exploits the information contained in the frequencies of the transitions between the macrostates of the order parameter distinguishing the two phases. The transition frequencies are observed in simulations initiated from macrostates with order-parameter values intermediate between those of the two phases; they are used to provide estimators of the macrostate transition probability matrix and thence estimators of the sampling distribution itself. The procedure allows one to construct a series of sampling distributions, weighted with respect to the canonical distribution, which approach the multicanonical limit, flat across order-parameter space. It entails only simulations that are short compared to the (multicanonical) relaxation time of the order parameter. Reweighting the transition-probability estimator of the multicanonical sampling distribution provides a good estimate of the canonical distribution of the order parameter for any value of the conjugate field, permitting the identification of the coexistence field in particular. The method is developed in the context of a system of hard spheres with short-range attractive interactions, described by a square potential well, which provides a simple model of the intercolloid depletion potential in colloid-polymer mixtures. In particular we explore the phase diagram in the region in which studies by others, based on free energy evaluation by thermodynamic integration, have shown the coexistence of two fcc solid phases of different densities. [S1063-651X(96)08705-3]

PACS number(s): 02.70.Lq, 05.70.Ce

I. INTRODUCTION

The problem of determining the phase behavior of a system is arguably *the* generic task of equilibrium statistical physics. The problem is traditionally couched in the language of thermodynamics [1]: one must calculate the free energy of candidate phases, the phase of minimum free energy (for given thermodynamic coordinates) being thermodynamically favored. Such calculations are not readily accomplished: the free energy (essentially the partition function) can virtually never be determined exactly for systems with interesting phase behavior. A variety of alternative strategies exist. Variational approximations [2], series expansions [3], and quasiharmonic approximations [4] provide varying degrees of insight and reliability. However, 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, the standard computational tool for dealing with many-body systems [5].

In its most commonly implemented form [Boltzmann importance sampling (BS)] the MC method allows one to generate a sequence (Markov chain) of microstates of the chosen system, with the assurance that, far enough along the sequence (at MC times large compared to the ergodic time over which the memory of the initial microstate is preserved), microstates will appear in the sequence with the equilibrium probabilities prescribed by the canonical Boltzmann form. In principle this standard MC framework allows one to address the problem of phase structure immediately and transparently, without the need to appeal explicitly to the concept of a free energy at all: the favored phase should be identified simply, in the simulation as in nature, by the macrostates that dominate once the initial transient process (''equilibration'') is complete. In practice, however, this simplistic approach founders because the MC procedure is plagued by long ergodic times near phase boundaries, so that the simulations remain trapped in the phase (group of macrostates) favored by the choice of initial state. This problem is actually a reflection of the very faithfulness with which the standard MC procedure realizes the canonical distribution: paths linking the groups of macrostates associated with the two phases have Boltzmann weights that are exponentially small in the system size, implying exponentially large interphase-crossing times. If, then, one wishes to operate within the BS framework one is forced to deal *separately* with the two phases (and to resort to the language of free energies). Even then, the problem remains awkward [6]. BS methods yield good estimators of canonical expectation values such as the energy as simple averages over the sampled states; however, while the free energy can be expressed in terms of a canonical expectation value, this expectation value is estimated unacceptably poorly by BS (see, for example, the discussion in [7]). If one is to use BS techniques one must perform a series of independent simulations to measure the energy (or some component thereof), along some path—in either thermodynamic coordinate space or in the space of model parameters—which links the system of interest to some reference system, whose free energy is known analytically or has been estimated in some prior calculation. The free energy of the system can then be determined by

*Present address: Laboratory of Molecular Biophysics, Rex Richards Building, The University of Oxford, South Parks Rd., Oxford OX1 3QU, United Kingdom. Electronic address: graham@biop.ox.ac.uk

integrating the energy along the simulation path. These integration methods (IM's) are commendably simple to implement and have become the standard basis for MC measurement of free energies. Nevertheless, perhaps because of certain practical limitations of the method (which we shall touch on in Sec. VI) or more probably because of the general conviction that such an important problem merits a somewhat more direct solution, there have been many attempts to develop an alternative MC approach.

The recurring idea in such studies is that the solution should be sought in some form of extended sampling, that is, a MC procedure designed to sample from a distribution other than the canonical form. The seminal contributions are probably those of Torrie and Valleau [8]; the most recent variants are the *multicanonical ensemble method* of Berg and Neuhaus [9,10] and the related *expanded ensemble method* (or ''simulated tempering'') of Lyubartsev *et al.* [11] and Marinari and Parisi [12]. These techniques can be deployed in a variety of ways. As described elsewhere [7], they can be developed to yield the absolute value of the free energy in a single phase region. Alternatively (our focus of interest here), they may be used to assess directly the relative stability of two competing phases. In either case the key issue (in fact, if resolved fully, the *only* issue) is how to *construct* a sampling distribution tuned to the problem in hand. In particular, to address the competing-phase problem it is clear that one requires a sampling distribution that is weighted (with respect to the canonical form) so as to enhance the likelihood of macrostates lying on the interphase path, that is, values of the order parameter ($\mathscr{M}$, say) intermediate between those associated with the two phases. In the multicanonical ideal [9,10] the sampling distribution is flat along the interphase path.

Notwithstanding much recent activity, it seems fair to say that existing techniques for evolving appropriate sampling distributions are less than satisfactory. The typical strategy is to make an initial guess for the weights required to secure a multicanonical distribution and then to refine the guess. To refine it one needs to estimate the implied sampling distribution, to determine the extent to which it fails to be multicanonical; it is customary to form that estimator, in the obvious way, simply on the basis of a histogram of the macrostates visited when one MC samples from it. This *visited-states* (VS) strategy provides only a somewhat crude and ponderous basis for refining the weights to be attached to regions that the sampling procedure has failed to visit. We have recently introduced [7] an alternative strategy for the construction of multicanonical sampling distributions, which estimates the sampling distribution on the basis of a histogram of the transitions made between macrostates, in simulations launched from macrostates chosen to ensure that the *complete* range of the macrostate space of interest is probed from the outset. In this paper (a short account of which has appeared elsewhere [13]) we show that this *transition probability* (TP) method can be applied effectively to the determination of a sampling distribution that is multicanonical along an interphase path.

It has been customary to distinguish between two aspects of the multicanonical program: the *construction* stage (just discussed) is followed by a *utilization* stage entailing further sampling from the putative multicanonical distribution; canonical averages are then determined by reweighting [14] the measured multicanonical averages to compensate for the bias introduced by the multicanonical weights. The utilization stage encounters a further problem: although the ergodic time for the multicanonical distribution is not exponentially large in the system size, it remains typically power law large, reflecting the fact that the MC dynamics through $\mathscr{M}$ space is essentially diffusive, in the multicanonical limit. In some instances this problem is acute: in the case of the particular phase boundary studied here it appears in such an extreme form that sampling over periods long compared to the ergodic time is altogether impractical. However, we shall show that the TP estimator of the multicanonical distribution can be determined, to the accuracy required to deal with the interphase problem, on the basis of simulations *short* compared to the ergodic time (they require equilibration only in a less-exacting sense) and this problem is skirted.

To date, multicanonical methods have been utilized in studies of phase coexistence in ferromagnets [9], fluids [15], and lattice gauge theories [16]. In the present study we focus on a problem of *structural* phase behavior. We consider a model system of hard spheres with attractive interactions described by a square potential well. This model is believed to provide a qualitatively authentic description of the behavior of polymer-colloid mixtures [17]. The hard spheres represent the colloid particles; the square well models the depletion potential of interaction between the colloid particles, mediated by the polymer molecules; the range and strength of this ''depletion potential'' reflect, respectively, the polymer length and concentration. Extensive MC studies [18], supported by mean field calculations [19], have shown that, as the interaction range is reduced, the conventional liquid-gas critical point moves in towards the triple point and the line of liquid-vapor coexistence shrinks to zero, consistent with the behavior observed in experimental studies of colloid-polymer mixtures [20,21]. For still smaller ranges of potential (in a regime that has as yet proved inaccessible to experimental work) the existing theoretical work [18,19] predicts the appearance of a new line of phase coexistence, stretching out from a triple point and again terminating in a critical point, but now separating two *solid* phases, with the same symmetry (both are fcc), but different lattice spacings; this line of phase coexistence is the object of the present study.

The motivation for this choice has three strands. First, the system itself is clearly of sufficient intrinsic interest to warrant studies extending and corroborating existing work: the square-well model provides the simplest context in which to explore important general issues such as the conditions that the interaction range has to satisfy if a liquid phase is to exist at all [22]. Second, tackling the solid-solid phase boundary in this system offers a sensible way to progress multicanonical studies: while going somewhat beyond the problems presented by liquid-vapor coexistence, it skirts the distinctive difficulties associated with melting and with solid-solid phase boundaries involving a change in symmetry. Third, it provides an appropriate context for comparing the multicanonical approach with established methods: the existing MC studies exploit integration methods to relate the free energy (the Helmholtz function) to that of a system of non-

interacting hard spheres, whose free energy has been independently calculated [23].

The layout of the paper is as follows. Section II defines the model, establishes the notation, and identifies the objectives of the calculations that follow. Section III describes the principles of our procedure for establishing a multicanonical sampling distribution. MC implementation details are described in Sec. IV. Section V presents the results of our MC studies, designed to explore both the general principles of the method and the particular physics of our model system, with emphasis on the solid-solid phase boundary. Section VI contains discussion and conclusions.

II. MODEL, NOTATION, AND OBJECTIVES

We consider a system of N particles with spatial coordinates $\{\vec{r}\}$ and configurational energy

$$\Phi(\{\vec{r}\}) = \sum_{i<j=1}^{N} \phi(|\vec{r}_i - \vec{r}_j|). \tag{1}$$

The pairwise interaction potential has a hard-core square-well form

$$\phi(r) = \begin{cases} \infty, & 0 \leqslant r < \sigma \\ -\epsilon, & \sigma \leqslant r < (1+\delta)\sigma \\ 0, & r \geqslant (1+\delta)\sigma. \end{cases} \tag{2}$$

The parameter δ controls the range of the potential. The existing MC studies [18] indicate that the liquid-vapor coexistence curve disappears for $\delta \lesssim 0.25$ and that a line of solid-solid phase coexistence appears for $\delta \lesssim 0.06$. Here we shall focus on the behavior for the single value $\delta = 0.01$, for which the existing data suggests [18] that the phase boundary terminates at a critical temperature $T_c \simeq 1.6$ [24].

Our simulation envisages particles confined in a *variable* cubic volume V, with periodic boundary conditions. Each particle is associated with the site of a fcc lattice: by this we mean that simulations are always initiated from fcc arrangements and that studies of the pair correlation function (to be reported below) confirm that crystalline order is preserved under all the conditions studied. The crystallinity constraint restricts the choice of particle number to the set $N = 4m^3$, with m integral; we have studied systems with $N = 32$, 108, and 256 particles, to provide the basis for a finite-size scaling analysis. The IM studies [18] were performed in a constant-V ensemble, for systems of $N = 108$ particles, over a wide range of values of T, V, and δ.

The canonical weight associated with the N particle coordinates $\{\vec{r}\}$ in volume $V = Nv$, at temperature T, and pressure p is prescribed by the probability density

$$P^c(\{\vec{r}\}, v|N,T,p) = \frac{1}{\mathcal{Z}(N,T,p)} \exp[-\beta p V - \beta \Phi(\{\vec{r}\})], \tag{3}$$

where $\beta = 1/kT$ and the partition function $\mathcal{Z}(N,T,p)$ is prescribed by

$$\mathcal{Z}(N,T,p) = \int_0^\infty dV \prod_j \left[\int_V d\vec{r}_j \right] \exp[-\beta p V - \beta \Phi(\{\vec{r}\})]. \tag{4}$$

The two solid structures coexisting at the phase boundary of interest are believed to differ simply by a homogeneous dilatation. One may therefore choose as the order parameter $\mathcal{M}$ the density or, equivalently (and our preference), the specific volume v and—the true content of this ''choice''—focus on the canonical probability density function (PDF) for volume macrostates

$$P^c(v|N,T,p) \equiv \prod_j \left[\int_V d\vec{r}_j \right] P^c(\{\vec{r}\}, v|N,T,p). \tag{5}$$

This probability density constitutes the natural microscopic finite-sized counterpart of the Helmholtz free energy. Specifically, combining Eqs. (3)–(5) we find

$$-\ln P^c(v|N,T,p) = \beta p V + \beta F(N,v,T) + \ln \mathcal{Z}(N,T,p), \tag{6}$$

where

$$F(N,v,T) = -\beta^{-1} \ln \left\{ \prod_j \left[\int_V d\vec{r}_j \right] \exp[-\beta \Phi(\{\vec{r}\})] \right\}$$

is the Helmholtz free energy for a system of N particles, whose thermodynamic limit identifies the free energy density

$$f(v,T) \equiv \lim_{N\to\infty} \frac{1}{N} F(N,v,T). \tag{7}$$

From a thermodynamic perspective the conditions for the coexistence of two phases, with specific volumes v_A and v_B, are, first, the requirement of a common pressure

$$p = -\left(\frac{\partial f(v,T)}{\partial v} \right)_{v_A} = -\left(\frac{\partial f(v,T)}{\partial v} \right)_{v_B} \tag{8}$$

and, second, the equality of the Gibbs free energy densities $g = pv + f$,

$$[pv + f(v,T)]_{v_A} = [pv + f(v,T)]_{v_B}. \tag{9}$$

Microscopic, finite-sized counterparts of these conditions can be realized in either of two principal ways. First one may seek conditions under which the PDF(5) displays two maxima, at v_A and v_B, of *equal heights* [25]. The condition that v_A and v_B locate maxima gives, on appeal to (6),

$$0 = \left[p + \frac{1}{N} \frac{\partial F(N,v,T)}{\partial v} \right]_{v_A} = \left[p + \frac{1}{N} \frac{\partial F(N,v,T)}{\partial v} \right]_{v_B}, \tag{10}$$

which gives the finite-sized counterpart of the common pressure condition (8). Similarly the ''equal heights'' requirement itself gives

$$\left[pv + \frac{1}{N} F(N,v,T) \right]_{v_A} = \left[pv + \frac{1}{N} F(N,v,T) \right]_{v_B},$$

which realizes the condition (9), in the thermodynamic limit. Alternatively one may establish the conditions under which the PDF(5) displays two peaks of *equal weights* [26]. The saddle-point equations for the locations of the peaks (asymptotically sharp) give Eq. (10) again. Equating the weights of the peaks gives

$$\left[pv + \frac{1}{N}F(N,v,T) + \frac{1}{2\beta N}\ln\left\{ \frac{\partial^2 F(N,v,T)}{\partial v^2} \right\} \right]_{v_A} = \left[pv + \frac{1}{N}F(N,v,T) + \frac{1}{2\beta N}\ln\left\{ \frac{\partial^2 F(N,v,T)}{\partial v^2} \right\} \right]_{v_B},$$

which again yields the equal Gibbs densities condition (9), in the thermodynamic limit. Both equal-heights and equal-weights conditions thus provide a valid basis for the identification of a phase boundary in a finite system, in the sense that the results obtained will approach the correct limit for large N.

Deferring (until Sec. V C) the question of which of these criteria is optimal, we see that the information required to determine the relative stability of the two phases (along with all properties that can be written as canonical averages of functions of v only) is contained within the PDF $P^c(v|N,T,p)$. While the *structure* of this PDF, in the vicinity of its peaks, is readily determined by standard constant N-p-T MC techniques [27], the relative *scale* of the peaks (essential to both equal-heights and equal-weights criteria) is not, because of the low canonical weight (we shall subsequently explore how ''low'') of v macrostates lying between them. In principle, the solution to this problem is to construct a non-Boltzmann sampling distribution for microstates, of the form [cf. Eq. (3)]

$$P^s(\{\vec{r}\},v|N,T;\{\eta\}) = \frac{1}{\mathcal{Z}(N,T;\{\eta\})}\exp[\eta(v) - \beta\Phi(\{\vec{r}\})]. \tag{11}$$

where $\{\eta\}$ denotes a set of dimensionless ''weights,'' one associated with each of the volume macrostates [28]. The associated PDF for the volume $P^s(v|N,T;\{\eta\})$ prescribes the canonical PDF of interest, for *any* desired pressure p, by appeal to the reweighting [14]

$$P^c(v|N,T,p) \propto P^s(v|N,T;\{\eta\})\exp[-\beta pV - \eta(v)] \tag{12}$$

with the overall proportionality constant determined by normalization. The task, then, is to construct an appropriate sampling distribution (in effect, a set of weights). The hallmark of such a distribution is simply that it can be estimated accurately by MC simulation, so that the canonical distribution (for the pressures of interest) can itself be determined to the required accuracy, by the reweighting (12).

III. ESTIMATING, REFINING, AND EXPLOITING THE SAMPLING DISTRIBUTION

The task we have just identified has three elements. First, one must define more explicitly the attributes of the sampling distribution. Second, one must devise a scheme by which it may be constructed. Third, one must exploit it.

In the approach to the phase-coexistence problem developed by Berg and Neuhaus [9] the target sampling distribution is chosen to be *multicanonical* (i.e., flat [29]) across macrostate space. This choice is clearly *reasonable* since it addresses the core of the problem presented by the canonical distribution, namely, the low probability of macrostates on the interphase path. While it is not clear that it is *optimal* for the task in hand, we shall also adopt it here; we shall find *a posteriori* rationalization for this choice [30].

The construction process itself has two aspects. One must choose a method for *estimating* the ''current'' sampling distribution, that is, the sampling distribution associated with one's current ''guess'' at a multicanonical set of weights. One must also choose an algorithm for *refining* the distribution, that is, refining the set of weights to bring the sampling distribution closer to the multicanonical ideal. We consider these two aspects in turn.

In [9] the sampling distribution is estimated by counting the visits to each macrostate in a MC exploration of that distribution [the ''visited states'' (VS) strategy]. Here we adopt a different strategy [the ''transition probability'' (TP) method], designed to use the information contained in measurements of the frequencies of transitions between the different macrostates. The idea is explored in general terms elsewhere [7]; we present it here in outline form and in the specific context of the current problem.

Let $\alpha \equiv [\{\vec{r}\},v]$ label the *micro*states of the system; the microstates form a continuous set but, for simplicity, we shall adopt a discrete and abbreviated notation in which P^s_α represents the probability $P^s(\{\vec{r}\},v|N,T;\{\eta\})d^N\vec{r}\,dv$. Let $i=1,\ldots,N_m$ label the *macrostates* associated with a set of volumes $\{v\}$ with members v_i, which span the region of potential interest, with $P^s_i = P^s(v_i|N,T;\{\eta\})dv$ the associated macrostate probability. Denote by $\rho^s_{ii'}(t)$ the transition probability from macrostate i to macrostate i', at time t after the initiation of a MC process designed to generate the chosen sampling distribution. Then

$$\rho^s_{ii'}(t) = \sum_{\alpha\in i}\sum_{\alpha'\in i'} P^s(\alpha|i)(t)\rho^s_{\alpha\alpha'}, \tag{13}$$

where $\rho^s_{\alpha\alpha'}$ is the transition rate between *microstates*. By the defining property of the MC procedure this microstate transition rate must satisfy the detailed balance condition

$$\rho^s_{\alpha\alpha'} = \rho^s_{\alpha'\alpha}\frac{P^s_{\alpha'}}{P^s_\alpha}. \tag{14}$$

We shall *impose* the further requirement that the conditional probability in Eq. (13) has its stationary limiting form

$$P^s(\alpha|i)(t) \simeq P^s(\alpha|i). \tag{15}$$

This condition presupposes that the microstates from which macrostate transitions are attempted are selected with the true sampling probability, *for the given macrostate*. The most obvious way of fulfilling this condition is to ensure that the position coordinates $\{\vec{r}\}$ are always given time to equili-

brate for the given v [indeed, equilibrate to their *canonical* distribution, since $P^s(\alpha|i) \equiv P^c(\alpha|i)$, because the weights depend only on the macrostate coordinate v]. However, it is helpful to note a less obvious alternative. *If*

(i) the initial macrostate is chosen to be i with probability P_i^s and (ii) the initial microstate is chosen to be α with the conditional probability $P^s(\alpha|i)$ [so that Eq. (15) is satisfied *at* $t=0$], then it is easily seen that Eq. (15) is satisfied at $t=1$, and subsequently. In effect conditions (i) and (ii) together ensure that the initial microstate is chosen with the equilibrium sampling distribution and the detailed-balance condition ensures that this remains so for subsequent microstates too. As we shall discuss in Sec. V A, it is growingly easy to fulfill these conditions the closer the sampling distribution comes to the multicanonical limit.

Now, returning to the main argument, we substitute Eqs. (14) and (15) into (12) and appeal to the identity

$$\frac{P^s_{\alpha'}}{P^s_{\alpha}} = \frac{P^s_{i'} P^s(\alpha'|i')}{P^s_i P^s(\alpha|i)}$$

to find that

$$\rho^s_{ii'}(t) = \rho^s_{ii'} = \rho^s_{i'i} \frac{P^s_{i'}}{P^s_i}. \tag{16}$$

This result shows that a MC estimate of the macrostate TP matrix provides an alternative route to the sampling probability distribution. The potential advantages that this strategy has with respect to the VS method reside in the fact that it does not presuppose equilibrium over the space of macrostates. Thus one may make use of simulations that are targeted (through the choice of the initial macrostate) on specific regions of macrostate space, thereby allowing rapid gathering of information even in the region of low canonical weight. Moreover, one can bypass the ergodic problems that persist through to the multicanonical limit: we do not need simulations that extend for periods long compared to the ergodic time. This point is crucial to the present studies; we shall return to it.

An estimate of the sampling distribution provides the basis for refining its parameters (the second stage of the construction process) to bring it closer to the multi-canonical limit. The simplest scheme [9] takes an estimate $\widetilde{P}_i^{(n)}$ of the elements of the sampling distribution $P_i^{(n)}$ associated with a particular set of weights $\{\eta^{(n)}\} \equiv \eta_i^{(n)}, i=1 \ldots N_m$, and generates a new set of weights $\{\eta^{(n+1)}\}$ by

$$\eta_i^{(n+1)} = \eta_i^{(n)} - \ln \widetilde{P}_i^{(n)} + \text{const}. \tag{17}$$

The constant is arbitrary [the sampling probabilities (11) are invariant against a uniform increment of the weights]; we fix its value by the convention $\eta_i = 0$ for $i = N_m$. The iterative scheme (17) is by no means ideal: in particular it provides no way of taking account of confidence levels in the existing set of weights and in the estimator of the current sampling distribution. We have examined these issues elsewhere [7]; however, here we shall use this basic scheme.

Iterating the two-component process we have described (*estimation* of $\{P^{(n)}\}$ and *refinement* of $\{\eta^{(n)}\}$) completes the construction stage; from it there emerges an approximately multicanonical set of weights $\{\widetilde{\eta}^{\star}\}$ associated with an approximately multicanonical distribution with elements $\widetilde{P}_i^{\star} \simeq$ const. We now proceed to consider how (and why) we *use* this distribution. In principle Eq. (12) allows us to reweight an estimate of *any* sampling distribution, to provide an estimate of the canonical distribution. In the abbreviated notation of this section we write the elements of the estimator, derived from the estimator of the nth sampling distribution in the form

$$\widetilde{P}_i^c \propto \widetilde{P}_i^{(n)} \exp[-\beta N p v_i - \eta_i^{(n)}]. \tag{18}$$

In practice we *choose* to make use of our estimate of the multicanonical limiting distribution, determining our reweighted estimator by

$$\widetilde{P}_i^c \propto \widetilde{P}_i^{\star} \exp[-\beta N p v_i - \eta_i^{\star}]. \tag{19}$$

To rationalize this choice we note the identity

$$\eta_i^{\star} = \eta_i^{(n)} - \ln P_i^{(n)} + \text{const} \tag{20}$$

relating the true multicanonical weights to the true elements of the nth sampling distribution. Combining this result with Eq. (17) gives

$$\ln \widetilde{P}_i^{(n)} - \ln P_i^{(n)} = \eta_i^{\star} - \eta_i^{(n+1)}. \tag{21}$$

The faithfulness of the estimator $\widetilde{P}^{(n)}$ (its closeness to the true nth level distribution) is thus reflected in the closeness of the $(n+1)$th level weights to the multicanonical limit. Thus evidence that a given set of weights needs further refinement is *a posteriori* evidence of deviations of the estimate of the sampling distribution (used to establish those weights) from the true sampling distribution. The iterative process may thus be seen as producing a sequence of sampling distributions, which are more reliably estimated the closer they come to the multicanonical limit.

IV. IMPLEMENTATION OF THE MC PROCEDURE

Our simulations have been performed on the Connection Machine CM-200, using a mixture of geometrical decomposition and primitive parallelism to allow $N_r = O(10^3)$ independent replicas to be run in parallel. The simulations comprise two kinds of update procedures: updates of the particle coordinates, for a given system volume, and updates of the system volume itself. We discuss these in turn.

The particle coordinates $\{\vec{r}\}$ are updated using a standard Metropolis algorithm [5]. Trial ''new'' coordinates are chosen from within a sphere centered on the current particle position; an update from coordinates $\{\vec{r}\}$ to coordinates $\{\vec{r'}\}$ is accepted with probability

$$P^a(\{\vec{r}\} \to \{\vec{r'}\}) = \min[1, \exp(-\Delta E)], \tag{22}$$

where

$$\Delta E = \beta \Delta \Phi \equiv \beta [\Phi(\{\vec{r'}\}) - \Phi(\{\vec{r}\})].$$

To maintain the acceptance probability at appropriate levels (~ 0.5) requires a value for the radius of the particle displacement sphere of order $\delta/2$, where δ is the width of the potential well Eq. (2).

We now turn to the volume updates. We found it convenient to restrict the values of v explored to a discrete set $\{v\} \equiv v_i, i=1,\ldots,N_m$, extending across the range of values of potential interest. (We shall discuss the spacing between the members of this set below.) Changes in volume are most conveniently realized as homogeneous dilatations that preserve the scaled coordinates

$$\vec{s} \equiv L^{-1}\vec{r} \text{ with } L \equiv (Nv)^{1/3}. \tag{23}$$

In a volume transition a candidate ''new'' volume v_j is chosen with equal likelihood from the two members of the set $\{v\}$ adjacent to the current volume v_i; the acceptance probability is computed according to the Metropolis algorithm [31]

$$P^a(v_i \rightarrow v_j) = \min[1, \exp(-\Delta E)], \tag{24}$$

where

$$\Delta E \equiv \beta \Delta \hat{\Phi} - \Delta \eta - N \Delta \ln v, \tag{25}$$

with

$$\hat{\Phi} \equiv \hat{\Phi}(\{\vec{s}\}, v) \equiv \Phi(\{L\vec{s}\}),$$

$$\Delta \eta \equiv \eta(v_j) - \eta(v_i) \equiv \eta_j - \eta_i .$$

The final term in Eq. (25) reflects the Jacobian of the transformation from $\{\vec{r}\}$ to $\{\vec{s}\}$ coordinates in the partition function (4).

The acceptance probability for volume transitions is strongly dependent on the volume step size: since a trial volume change resulting in even a single hard-core overlap will be rejected, only ''small'' volume changes have a reasonable chance of being accepted. The step size $\Delta v \equiv v_j - v_i$ must therefore be correspondingly small. Since the extent of the problem varies with the density (it is most acute at high densities), so must the step size in order to keep the acceptance probability roughly constant over the range covered by the set $\{v\}$ [32]. We constructed suitable sets (maintaining acceptance probabilities close to 0.5) guided partly by experiment and partly by simple statistical arguments. For $N=256$, for example, the chosen set $\{v\}$ comprised $N_m \sim 3\times 10^3$ members with spacings ranging from $\Delta v \simeq 10^{-5}$ (near $v_1 = 0.7171$ [24]) to $\Delta v \simeq 10^{-4}$ (near $v_{N_m} = 0.818$). As we shall see, the need to handle large values of N_m is a key feature of this problem.

The two update procedures we have described are combined to give the following MC program, allowing one to estimate the macrostate probabilities associated with a given sampling distribution. Each replica simulation (RS) is initialized with the particles assigned to the sites of a fcc lattice, with volume chosen randomly from the set $\{v\}$. Each RS is equilibrated using the constant-volume updating scheme for the spatial coordinates $\{\vec{r}\}$ specified in Eq. (22). Each RS is then allowed to undergo ''volume'' transitions (dilatations) following the algorithm (24). Each volume transition is followed by a further N_e updates of the spatial coordinates alone to allow equilibration within the new volume macrostate. The procedure (volume update followed by coordinate equilibration) is repeated N_c times for each RS, pooling the records of all volume transitions from all N_r replicas in a histogram. We denote the elements of the histogram gathered at sampling stage n (i.e., that associated with a MC procedure with sampling distribution characterized by weights $\{\eta^{(n)}\}$) by $C_{ij}^{(n)}$. The elements of this histogram provide estimators $\tilde{\rho}_{ij}^{(n)}$ for the elements of the associated macrostate TP matrix $\rho_{ij}^{(n)}$ [33]:

$$\tilde{\rho}_{ij}^{(n)} = \frac{C_{ij}^{(n)} + 1}{\Sigma_j (C_{ij}^{(n)} + 1)}. \tag{26}$$

Estimators $\tilde{P}_i^{(n)}$ of the macrostate probabilities $P_i^{(n)}$ can then be determined by appeal to Eq. (16). Since the TP matrix is tridiagonal, this equation is readily solved by a simple iterative scheme. In the first one or two iterations the sampling probabilities span many orders of magnitude and it is necessary to work with the logarithm of $\tilde{P}_i^{(n)}$ to prevent arithmetic overflow. To limit the buildup of rounding errors it is necessary to use Eq. (16). to generate the probabilities $\tilde{P}_i^{(n)}$ in *increasing* order. If one attempts to generate the macrostate probabilities in *decreasing* order, the uncertainties in the resulting estimates become unacceptably large once one has traversed six or seven orders of magnitude (postponed to 14 with double precision arithmetic).

The precision of the estimator that emerges is controlled, predominantly, by the number N_c of the transitions recorded. It is illuminating to have some *a priori* estimate of the value to be assigned to this parameter. To that end, let

$$Q \equiv \frac{P_a^s}{P_b^s} \tag{27}$$

denote the ratio of the probabilities of two macrostates $i=a$ and $i=b$, at the extremes of $\{v\}$ space (lying close to the modes associated with the two phases), for some sampling distribution P^s. Our objective, in this problem of phase coexistence, requires that we determine the ratio Q to within a fractional uncertainty of order unity:

$$\frac{\delta Q}{Q} \sim O(1). \tag{28}$$

Given Eqs. (6) and (7), this condition is equivalent to the requirement that the difference between the Helmholtz free energy *densities* of the two phases is determined to within corrections $O(1/N)$, which are comparable with the finite-size corrections we must generally expect to contend with. Now, appealing to Eq. (16) we may write

$$Q = \prod_i^{N_m - 1} q_i \text{ where } q_i = \frac{\tilde{\rho}_{i,i+1}^{s}}{\tilde{\rho}_{i+1,i}^{s}}. \tag{29}$$

If we neglect correlations among the variables r_i, we may then implement the requirement (28) in the form

$$1 \sim \frac{\delta Q}{Q} \sim \sqrt{N_m}\frac{\delta q}{q} \sim \sqrt{\frac{N_m}{N_v}},$$

where

$$N_v \sim \frac{N_r N_c}{N_m}$$

is the mean number of visits per macrostate [34]. Combining these two results, we conclude that the precision demanded in Eq. (28) requires that

$$N_c \sim \frac{N_m^2}{N_r}. \tag{30}$$

This result provides not only a practical guide to the parameterization of the MC procedure but also an immediate insight into the advantages of the TP estimator (relative to the VS estimator) of the sampling distribution. Use of the VS estimator presupposes a simulation that extends over times that are at least comparable with the interphase-crossing time. Even in the multi-canonical limit, where this evolution is diffusion limited rather than barrier limited, this implies a number of steps N_c through v space of order N_m^2. This constraint applies to *each* contributing replica. Thus, while the large values of N_r falling within the compass of parallel computers such as the CM-200 enhance the statistics gathered over such a period, they do not allow us to reduce this period itself, within the VS framework [35]. Given the need to handle large values of N_m—as we have seen, an unavoidable corollary of the hard-core potential—the VS strategy is then quite impractical here [36]. In contrast, Eq. (30) shows that in the TP method the simulation time required is reduced by of order $1/N_r$. In this approach, each replica simulation typically explores only a small region of macrostate space; the method provides a simple framework within which the information that each simulation provides (about the relative probabilities of macrostates in the region it samples) can be pooled to provide an estimator of the sampling distribution for the whole macrostate space.

V. RESULTS

A. The multicanonical sampling distribution

We proceed to explore the results emerging from the procedure described in the preceding section, focusing first on the generation of the multicanonical sampling distribution. Figure 1 shows the results of the weight-generation procedure for a system of $N=32$ particles, at a temperature $T=1.1$ expected to lie below the critical temperature for solid-solid phase coexistence [18]. The first level weights $\{\eta^{(1)}\}$ are set to zero (the initial sampling distribution is canonical). The general structure of the multicanonical weights over the entire region is apparent already after one iteration; convergence to a form very close to a fixed point of Eq. (17) is achieved within seven or eight iterations requiring only about 25 min of CM-200 processing time. Further convergence is secured in the course of the final ''production-and-refinement'' iterations described below. The shape of the weight function $\eta(v)$ is recognizable as that of the Helmholtz function for a finite-sized system, a correspondence that is implied by Eqs. (6) and (12), in the multicanonical limit [37]. The same procedure is practicable on larger systems: for $N=108$ an *ab initio* determination of multicanonical weights (spanning 125 *decades* of probability) required some 2 h of CM-200 processing time. The procedure can be streamlined by using a finite-size-scaling (FSS) estimate of an initial set of weights [38]; typically we found that this FSS estimate is quite close to the multicanonical limit (the discrepancy gets larger with increasing T), so only one or two iterations are required before it is possible to move to the production-and-refinement stage.

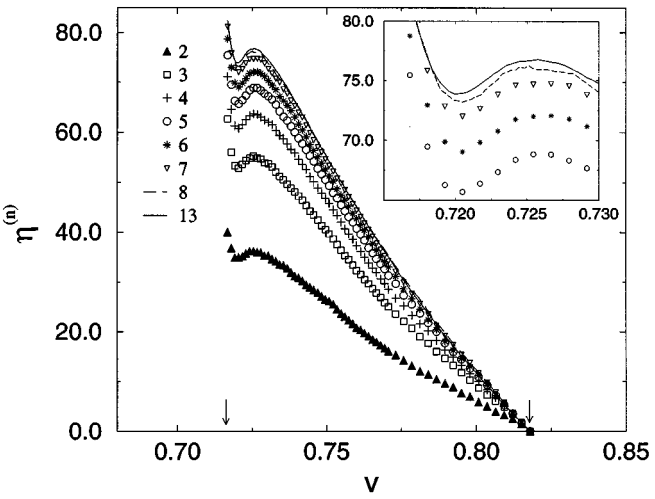

FIG. 1. Multicanonical weights $\eta^{(n)}(v)$ for various iterations n, with $N=32$, $T=1.1$ [24]; the arrows identify the bounding volumes in the set $\{v\}$. The MC control parameters are $N_c=250$, $N_m=237$, and $N_r=4096$. The inset shows a detail of the high-density region. The units are specified in [24].

The results shown in Fig. 1 exploit the TP estimator of the sampling distribution. To illuminate the workings of the TP method (and to contrast it with the VS approach) we show, in Fig. 2, histograms of the counts of macrostate visits (accumulated over all N_r replicas) during (a) an early and (b) a late stage in the iterative weight-generation process. These results are *not* used in the weight update procedure. The purpose of showing Fig. 2(a) is to emphasize that the VS histogram initially provides virtually *no* useful information about most of the macrostate space (it does pick up a local attractor in the low-v regime). The structureless form of the

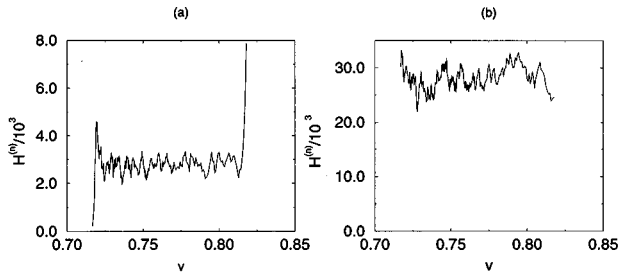

FIG. 2. Histograms $H^{(n)}$ of the volume macrostates visited at sampling stages (a) $n=2$ and (b) $n=13$ in the course of the weight generation shown in Fig. 1. The units are specified in [24].

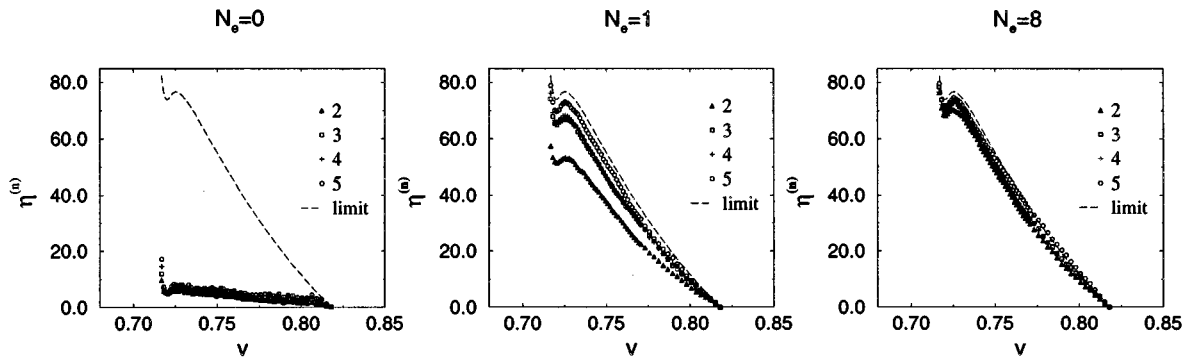

FIG. 3. Convergence of the weights over iterations $n=2,\ldots,5$, for different values of N_e. Other parameters are as in Fig. 1.

histogram simply reflects the fact that the various replicas have evolved relatively little from their initial states in the course of the simulation. The ''useful information'' resides in the *flow* through macrostate space (the macrostate transitions), not its *population*. Figure 2(b) shows that a late-stage histogram is *consistent* with multicanonical behavior; but this ''consistency'' is less than meets the eye. Again each replica has typically migrated across only a small part ($\sim 1/30$) of v space; the flatness of the histogram is thus a reflection of the initially uniform distribution of the replicas over this space. The relatively large fluctuations are also testimony to the strong correlations within each RS.

Next, let us consider the MC control parameters associated with the TP procedure. An indication of the interplay between the parameters N_c, N_r, and N_m is provided by Eq. (30). In fact, the results for various test runs imply that the algorithm is robust down to a value of N_c rather smaller even than this. Moreover, there are effects [which Eq. (30) does not capture; cf. [34].] that actually *reduce* the efficiency of the procedure as N_c is increased. In the early stages of the procedure (when the sampling distribution is far from multicanonical) the low-v macrostates tend to empty relatively rapidly. It is advantageous to set N_c low enough to preempt this process in order to avoid the need to reinitialize and reequilibrate the replicas, prior to gathering information for the next iteration.

Consider now the role of N_e, the number of constant-volume coordinate-updating sweeps (through all replicas) allowed between volume-update attempts. Figure 3 shows the evolution of the weights for three different values of N_e, over the early stages of the iterative procedure. It is clear that increasing the time devoted to equilibration between v updates increases the speed of convergence to the multicanonical limit. The principle here is easy to identify: the larger N_e is the closer one comes to satisfying the condition on the macrostate transition probability matrix [Eq. (15)] implicit in Eq. (16); the better the TP estimator of the current sampling distribution is, and the closer [cf. Eq. (21)] the implied weights are to their multicanonical limit. To understand the behaviour in more detail, in particular, why it is that the TP estimator continually *underestimates* the change in weights required to reach the multicanonical limit, it is necessary to consider what is occurring physically in the simulations. Suppose first that N_e is sufficiently large that equilibrium at each v is established and compare the transition probabilities of dilations and contractions. Transitions that increase v can be made relatively freely since a dilation introduces no hard-core overlaps, but transitions that decrease v *are* likely to be suppressed by the occurrence of such overlaps. Now suppose that N_e is reduced so that equilibration at each v is incomplete. While dilations are hardly affected, moves to lower volume become *more* probable: to the extent that equilibration is imperfect, the scaled particle coordinates (i.e., the set $\{\vec{s}\}$, defined in Eq. (23)] are preserved through to the next v update and a lower v macrostate may be recovered by a transition that reverses a preceding dilation and thereby effectively restores an earlier microstate. It follows that, in these circumstances, the ratio $\tilde{\rho}^{\,s}_{i,i+1}/\tilde{\rho}^{\,s}_{i+1,i}$ is underestimated, as is then the sampling probability gradient [cf. Eq. (29)], and thence the change to be made to the weights [cf. Eq. (17)].

Of course, while increasing N_e decreases the number of iterations required to reach near-multicanonical behavior it increases the time required for each iteration. We have not sought to establish the optimal conditions, but we expect efficiency gains from tuning N_e to be relatively small.

Now let us turn to the later stages of the iterative procedure. For large enough n ($n \simeq 8$ in Fig. 1) we reach a situation where the noise in the TP estimator drowns any remaining signal that the multicanonical limit has not been fully realized and the weights undergo only random fluctuations between iterations. At this point we move into the second (production-and-refinement) stage of the procedure. In this stage the advantages of low values of N_c noted above do not apply; accordingly it is reset to substantially higher values, typically of the order of a few thousand rather than a few hundred, per iteration. We continue to update the weights after each iteration, using the TP estimator of the current sampling distribution, but we also use this TP estimator, in conjunction with the current set of weights, to provide an estimate of the canonical distribution (for a pressure, or pressures, or interest), using the reweighting prescribed by Eq. (12). Thus we accumulate estimates of the canonical distribution from which, at the end of the simulation, we produce a ''best estimate'' of P^c, together with its uncertainty [32].

We have seen that the parameter N_e controls the quality of the TP estimator of the sampling distribution (and thence the rate of convergence of the weights) in the early stages of the procedure. It is important to establish what effects it has on the quality of the later estimators. To do so we ran a series of simulations employing a sampling distribution associated

with a refined set of multicanonical weights (in the particular case $N=32$, $T=1.1$). These simulations each performed the same number of volume updates, but varied in the value assigned to N_e. TP estimators of the sampling distribution were formed in each case and new values of weights computed in the usual way [by appeal to Eq. (17)]; the procedure was repeated (without weight updating) to provide uncertainties on the predicted new weights. We found that the ''new'' weights did not differ systematically from the ''refined'' estimate: the estimate of the sampling distribution is relatively insensitive to the value assigned to N_e, in the multicanonical limit. To see why this should be so we need to recall the alternative way of securing the validity of Eq. (15), expressed in the two conditions noted following that equation. While the second of these conditions (equilibration within the starting macrostate) is always fulfilled, the first is not; it is, however, fulfilled more closely the closer the sampling distribution comes to the multicanonical limit, because, in this limit, our ''random'' set of initial v macrostates *does* match the (flat) sampling distribution. Full (multicanonical) equilibrium is secured initially and preserved; Eq. (15) is satisfied without further intervention.

B. The canonical distribution: Initial applications

We now turn to the results for the canonical distribution and the physical properties that it defines. In this section we present the results for one particular system size ($N=108$); the size dependence (finite-size scaling) of our results is examined separately in the following subsection.

Figure 4 shows the results for the canonical PDF $P^c(v)$ at $T=1.1$, for a range of pressures p. Two ranges of pressures are clearly identified. Pressures in excess of $p_{cx}\simeq 29$ support a dense structure, characterized by a relatively sharp peak around $v\simeq 0.719$; pressures below p_{cx} support an expanded structure, with a substantially broader peak (note the different scales in the two figures) at higher v. The incipient dense structure is just discernible at $p=29$, in Fig. 4(b)

Figure 5 provides an alternative perspective of the canonical PDF, at three pressures, including a refined estimate of the coexistence pressure p_{cx}, identified by the equal-weights criterion (cf. Sec. II), for reasons to be discussed in Sec. V C. below. The inset shows the low canonical weight of macrostates along the interphase path; in some applications (for larger N and lower T) we have used the TP method to evolve distributions surmounting probability differentials of 260 decades.

On the face of it, the error bars on the PDFs in Figs. 4 and 5 are substantial [the fractional uncertainties in the PDFs are $O(1)$]. However, the bulk of this uncertainty originates in uncertainties in the relative weights of the two phases; the shapes of the peaks of P^c are individually well outlined and single-phase averages (i.e., averages calculated over the peaks separately) are substantially more precisely defined than these error bars might suggest. Moreover, as we have discussed [cf. remarks following Eq. (28)] the $O(1)$ uncertainties are, in fact, consistent with the determination of the free energy density (and thus the parameters of the coexistence curve) to within $O(1/N)$ corrections.

The differences between the two phases, already visible in Figs. 4 and 5, are made more explicit in Fig. 6, which shows the specific volume [Fig. 6(a)] and the compressibility [Fig. 6(b)] as functions of the pressure. Both quantities are evaluated from the canonical averages (of the volume and its fluctuations) over the range of v associated with the phase that is favored at the pressure under consideration.

C. Finite-size scaling

To establish (and optimize) the extent to which the MC procedure captures the limiting (large-N) behavior of ultimate interest requires a systematic analysis of the N dependence of the results. Figure 7 shows the results for the volume PDF, for various system sizes, at $T=1.0$. In each case the PDF has been evaluated at the pressure (an estimator for p_{cx}) that accords equal weights to the two phases. The sharpening of the peaks with increasing N is apparent and is also reflected in the behavior of the canonical average of the specific volume v, as the pressure is varied through the coexistence region, shown in Fig. 8(a). In contrast to Fig. 6(a), these averages were taken over the entire sampling distribution, at a given p, in order to display more clearly the growing precision with which the coexistence pressure is defined.

Now let us consider the merits of the different estimators of the coexistence boundary, discussed in Sec. II. Figure 8(b) shows the N dependence of the estimators for p_{cx} (at $T=1.0$) based on both the equal-weights (EW) and the equal-heights criteria. Although the limited range of N values studied allows no definitive conclusions to be drawn, the data suggest that *both* estimators are subject to $O(1/N)$ corrections with respect to the thermodynamic (infinite-volume) limit. The least-squares fits to both sets of data (dashed lines)

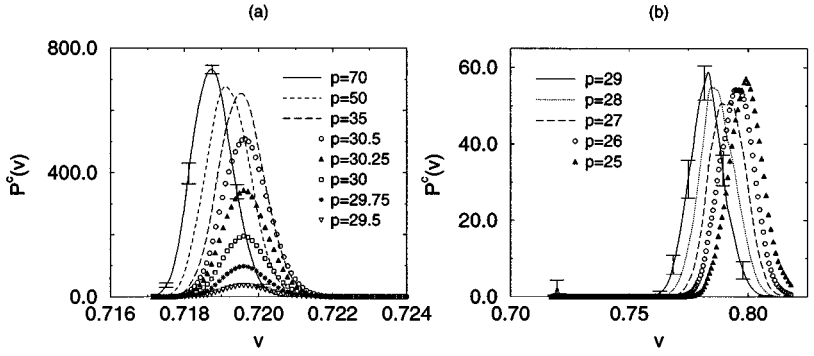

FIG. 4. Canonical PDF $P^c(v)$ for $N=108$, $T=1.1$, at various pressures p. The PDF was smoothed using a moving average over a window of 50 v states; some typical error bars are shown. The two figures show, respectively, the peaks corresponding to (a) the dense phase and (b) the expanded phase. The units are specified in [24].

are equally good and both have the same intercepts, within error. The ordinate intercepts, which are the best estimates of the infinite-volume transition point, are both consistent with the assignment $p_{cx}=22.78(6)$.

The question of which of the two estimators is the more appropriate *a priori* has attracted considerable attention. Arguments based on cluster expansions [39], supported by empirical analysis of MC data [39], have provided strong evidence that, at least for lattice models with periodic boundary conditions, the EW estimator is to be preferred in that it is correct to within terms that are *exponentially* small in the system size. It is not clear to us what these arguments imply (if indeed they are applicable at all) to the off-lattice N-p-T ensemble explored here. We choose to adopt the EW estimator, in part because it does appear to give the smaller corrections and in part because it seems to us to be the more natural choice.

It is of interest to explore the system-size dependence of the canonical probability of macrostates lying on the interphase path. With this in mind, let $Q^{\star}\equiv P^{c}_{\max}/P^{c}_{\min}$ denote the ratio of the canonical probabilities of two macrostates, one coinciding with a maximum of $P^{c}(v)$ (in practice we chose the maximum associated with the expanded phase) and the other associated with the minimum of $P^{c}(v)$, lying between the two maxima. This ratio measures the probability differential (''free energy barrier'') that the multicanonical weights are designed to offset. Figure 9 shows its behavior as a function of system size $L\equiv V^{1/3}$, for temperature $T=1.0$. The result is somewhat unexpected. One might anticipate that the macrostates along the interphase path would be dominated by mixed-phase arrangements, comprising macroscopic regions of each of the two phases, separated by an interface, with an area of order $L^{d-1}\equiv L^{2}$. This viewpoint leads immediately to the prediction $\ln Q^{\star}\sim L^{2}$. This prediction is certainly consistent with studies of lattice-based models, for which it has been exploited to provide a measurement of the interface tension [40,41]. However, the behavior measured here is much closer to $\ln Q^{\star}\sim L^{3}$. The reason for this difference is that, at least in the systems we have studied, the interphase states do not appear to be inhomogeneous in the way the above argument envisages: examination of the distribution of free volume along the interphase path shows that these states are no less homogeneous than those that dominate the peaks in the PDF. In part this may simply be a reflection of the relatively small system sizes we have been able to study: in this regime the ''interfaces'' separating the two phases have not only to patch one structure to the other,

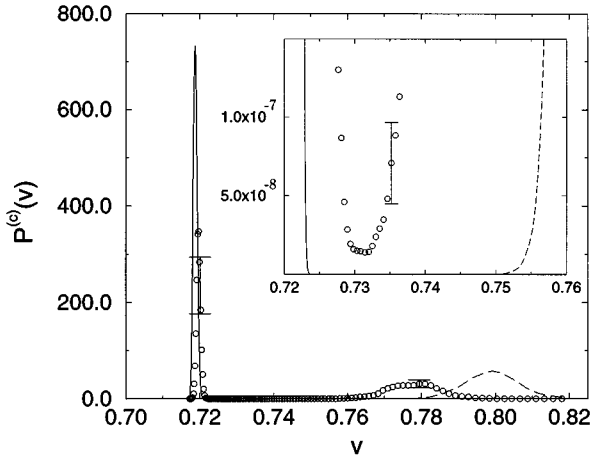

FIG. 5. Canonical distribution function $P^{c}(v)$ for $N=108$, $T=1.1$, at pressures $p=70$ (solid line), $p=25$ (dashed line), and $p=30.19\pm0.12=p_{cx}$ (circles). The inset shows a detail in the regime of intermediate density. The units are specified in [24].

but also to absorb the additional strains (relatively significant, for small systems) imposed by the requirement that the overall structure can be accommodated in a cubic sampling volume.

D. Phase diagram

We now turn to consider the solid-solid phase boundary in more detail. We have used the techniques described in previous sections to evolve multicanonical sampling distributions, and thence TP estimators for the canonical volume PDFS, over a range of temperatures (from $T=1.0$ up to the vicinity of the critical point believed to lie around $T\simeq1.6$ [18]) and system sizes ($N=32$, $N=108$, and $N=256$). Figure 10 shows some of the results, for systems of size $N=108$, at the coexistence pressure (for the corresponding N and T) determined by the equal weights criterion. With increasing T, the canonical probability of the region between the two modes increases and the modes merge together. For $N=108$ the PDF becomes unimodal around $T=1.7$; for $N=32$ the merger occurs at lower temperature. Clearly, implementation of the equal-weights criterion becomes progressively harder as one moves into this regime. To identify the equal-weights pressure we used an arbitrary division of the range of v at or near the point where P^{c} is minimal, although a fitting of two overlapping Gaussians is arguably

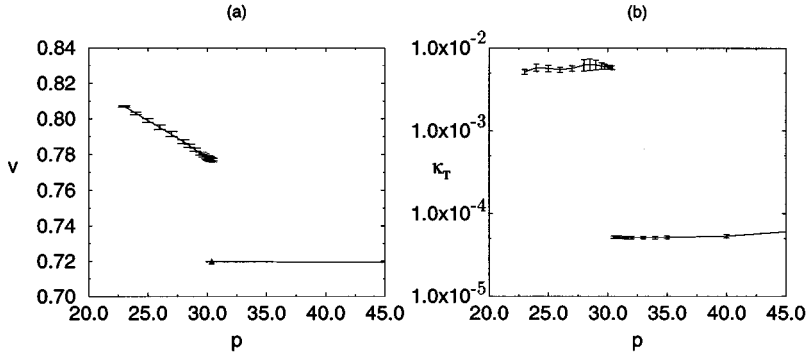

FIG. 6. Canonical expectation values for a system with parameters $N=108$, $T=1.1$, as a function of the pressure p. (a) The specific volume $v\equiv V/N$; the two data points represent estimates of the specific volume of each phase, at coexistence. (b) The isothermal compressibility κ_T; note the logarithmic scale on the ordinate.

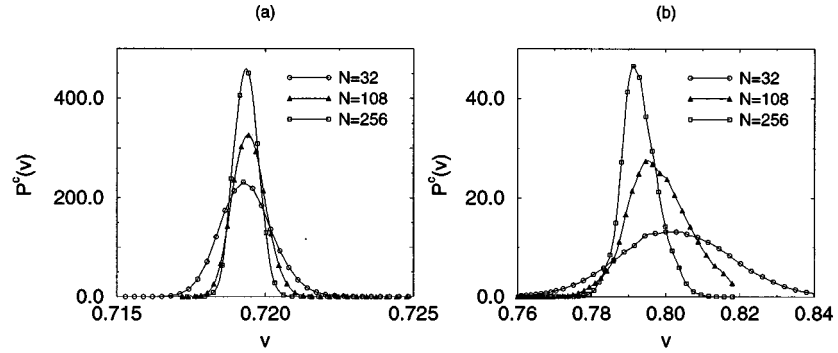

FIG. 7. Canonical PDF $P^c(v)$ for $T=1.0$ and three different values of N, showing separately (and on different scales) the regions corresponding to (a) the dense phase and (b) the expanded phase. The units are specified in [24].

more appropriate. To assign values to the specific volumes of the two phases, in this region, we simply used the location of the modes to provide finite-size estimators. The finite-size estimators of p_{cx} and the specific volumes were then extrapolated to give estimates of their thermodynamic limits.

The results are represented in Figs. 11(a) and 11(b), which show, respectively, the phase boundaries in the p-T and T-v planes. These figures also show the results established independently [18,42] using integration methods, on a system of size $N=108$. There is good agreement between the two estimates of the p-T coexistence curve; discrepancies are at most of the order of 1% and are generally within the error bars on the multicanonical points. The agreement in the location of the phase boundary in the v-T plane is also fairly good, though there are more obvious differences over the specific volume of the expanded solid—differences that grow with increasing temperature. There are several possible origins for these discrepancies. In part they may reflect the difference between the equal-heights and equal-weights criteria: the double tangent construction used in the integration method is, implicitly, an equal-heights estimator. Another contributory factor is the N dependence of the canonical PDF at coexistence, in particular the shift in the expanded-phase mode to lower v as N increases (cf. Fig. 7). Third, one might expect that the appeal to a hard-sphere solid reference system becomes more problematic the closer the physical system is to criticality: the typical configurations of the physical system and the reference system (which does not have a critical point) will become less well matched and the errors in the integration method harder to control.

Taken at face value, our results suggest that the critical temperature is rather lower than that implied by the IM results. However, to handle the critical region adequately it is necessary to appeal to a more sophisticated finite-size scaling framework, along the lines of that applied to the study of liquid-vapor coexistence [43], which have been developed to yield results of remarkable precision [15]. To do so would, we believe, be relatively straightforward.

E. Coexisting phases: A comparison

Finally, we turn to a comparison of the two solid phases and to consider the physical basis of their coexistence. First, we present the evidence corroborating that the structure of both phases is indeed fcc. Figure 12 shows the measured pair correlation function $\mathcal{G}(r)$, at coexistence (for $N=108$, $T=1.1$). In both phases $\mathcal{G}(r)$ exhibits well-defined peaks at positions (increasing as $\sqrt{m}$, with m the peak index) characteristic of a fcc lattice. The discontinuities in $\mathcal{G}(r)$ (apparent in the inset) can readily be traced to corresponding discontinuities in the potential 2, at $r/\sigma=1$ and $r/\sigma=1+\delta$. The measured discontinuity at $r=(1+\delta)\sigma\equiv r_o$ is consistent with the requirement that

$$\lim_{\Delta r\to 0^+}\frac{\mathcal{G}(r_0+\Delta r)}{\mathcal{G}(r_0-\Delta r)}=e^{1/T}.$$

To illuminate the circumstances in which the two phases coexist, we consider the elements of the Gibbs free energy density $g=f+pv=e-Ts+pv$, where e is the energy density and s the entropy density. The form of the Helmholtz function $f(v)$ can be deduced from the measured canonical PDF for v, by appeal to Eq. (6); the energy density is measurable directly, as a canonical average; the entropy density can then be inferred. Figure 13 shows the three contributing elements (e, s, and $p_{cx}v$) for $N=108$ at $T=1$, together with the Gibbs functional itself. The elements s, $p_{cx}v$, and g have been arbitrarily shifted vertically so that they equal zero at

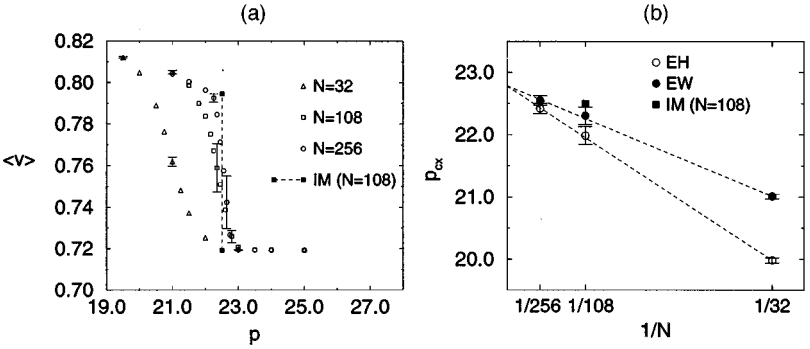

FIG. 8. Finite-size behavior for $T=1.0$ The IM data are taken from [18]. (a) The behavior of $\langle v\rangle$ as function of pressure p, near p_{cx}. (b) The coexistence pressure determined by equal-weights (EW) and equal-heights (EH) criteria, for different values of N. The dashed lines represent least-squares fits to the data and share the ordinate intercept $p=22.78(6)$.

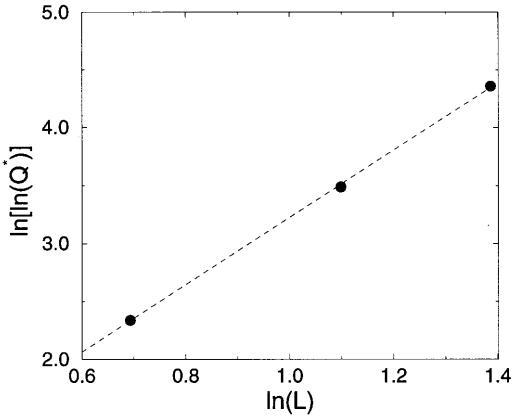

FIG. 9. Height of the interphase barrier measured by the ratio $Q^\star$ for systems of linear dimension $L=2,3,4$ (with $N=32,108,256$, respectively) at $T=1.0$. The dashed line represents a least-squares fit to the data and has a gradient of 2.91(1).

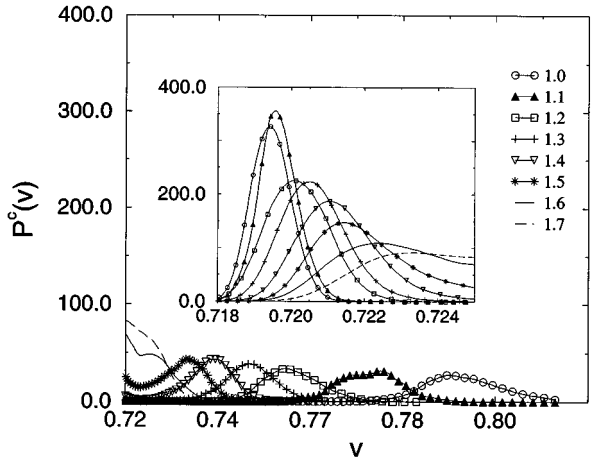

FIG. 10. Canonical PDF $P^c(v)$ for $N=108$ and a range of T values, at the corresponding coexistence pressure. The inset shows the dense phase on an expanded scale. The units are specified in [24].

the lowest v studied. The differences between the values of each element of g for the two phases (i.e., evaluated at the estimators of the two specific volumes at $T=1$) are represented by the arrows; to within $O(1/N)$ corrections they satisfy the condition $\Delta e - T\Delta s + p_{cx}\Delta v = 0$.

The *energy* has a close-packed ($v=0.707$) limit of $e=-6$; it begins to increase rapidly above this limit at $v\simeq 0.717$, where particles begin to lie outside the potential well of a significant fraction of their neighbors. It exhibits a point of inflection, close to the point $v=(1+\delta)^3/\sqrt{2}\approx 0.7285$ at which the energy of a fully ordered ($T=0$) crystal would step from $e=-6$ to $e=0$, under a homogeneous dilation. The *entropy* must display a logarithmic singularity in the close-packed limit; it increases steeply through the region characteristic of the dense solid, but its *curvature* is less than that of the energy. By contrast, in the rare-phase region the reverse is true: $s(v)$ exhibits the dominant curvature. The results express quantitatively what is qualitatively clear *a priori*: the dense phase owes its existence to, and its properties are dominated by, energetic considerations; for the rare phase, entropic factors are dominant. Note that the underlying Helmholtz function does not display the fully convex behavior required in the thermodynamic limit; the hump in the free energy functional g, straddling the interphase region, must be attributed to finite-size effects, which will persist until one reaches the regime of L values large enough that (cf. Sec. V C) $\ln Q^\star \sim L^2$.

As the temperature is raised towards the critical point, the functions $s(v)$ and $e(v)$ change relatively little; the evolution of the properties of the phases reflects, principally, the temperature-induced shift in the balance between entropy and energy: the low-v dominance of the energy is reduced and the high-v dominance of s enhanced, and the change-over between these regions occurs at lower v. This accounts for the key features observed: the dense phase becomes more compressible, the expanded phase less so, and the two phases coalesce as the specific volume of the rare phase falls.

For *lower* temperatures, on the other hand, the influence of the energy persists to higher v, pushing up the specific volume of the expanded phase and eventually taking the system into the region (below the triple point) where the expanded phase is unstable with respect to the liquid phase, which we have not explored in this work.

VI. DISCUSSION AND CONCLUSIONS

Although the system we have chosen to study in this paper is of considerable interest in its own right, this work has been principally concerned with *generic* issues arising in the MC study of phase behavior. These will remain the focus in our final discussion, which we shall organize into two stages. First we shall review the particular (transition-probability) implementation of the multicanonical method, developed and

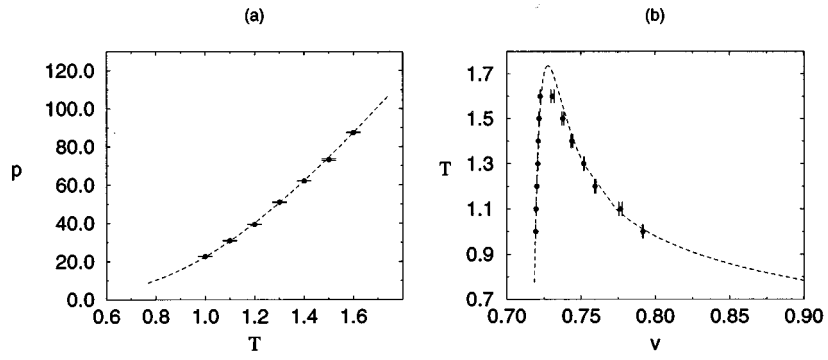

FIG. 11. Solid-solid phase boundary in (a) P-T and (b) T-v space. The data points are produced by extrapolating $N=32$ and $N=108$ (and for $T=1$, $N=256$) data against $1/N$. The dashed line shows thermodynamic integration results for an $N=108$ system [18,42]. The units are specified in [24].

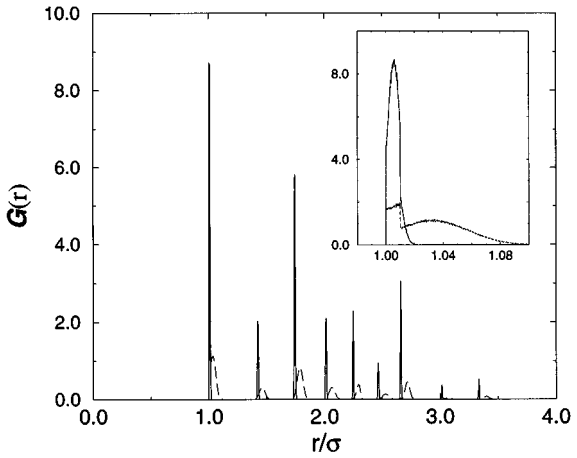

FIG. 12. Pair correlation function $\mathscr{G}(r)$ for $N=108$, $T=1.1$. The inset shows detail at small r. Full line, dense phase; dashed line, expanded phase.

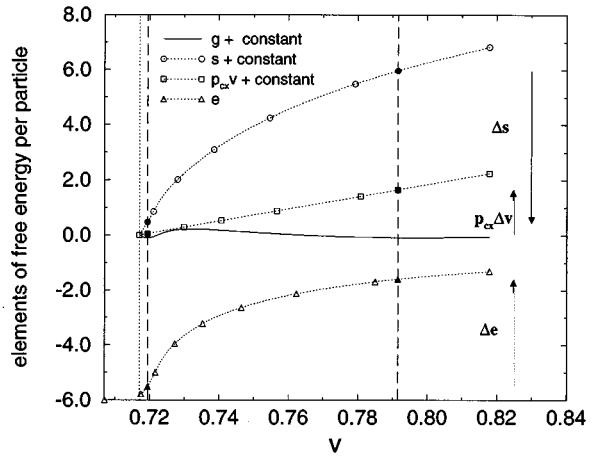

FIG. 13. Contributions to the free energy density at $T=1.0$. The entropy density s, the Gibbs density g, and the contribution $p_{cx}v$ have been translated by arbitrary constants so that they are zero at the low end of the range of v explored, shown by the dotted vertical line. The dashed vertical lines identify estimators of the specific volumes of the coexisting phases at $T=1.0$; the filled symbols identify the values of the free energy contributions for these values of v. The units are specified in [24].

applied here. Second, we shall compare the multicanonical approach to the study of phase boundaries with the traditional MC method based on integration along a path.

The TP estimator of the sampling distribution provides a tool that, we believe, should prove of rather general value in multicanonical studies. It addresses two key problems. First it allows one to begin to gather information about the whole of a targeted range of macrostate space, virtually from the outset of the procedure; in this respect it offers significant advantages with respect to VS, which is slow to deal with poorly sampled regions and vulnerable to inaccurate initial guesses at the values to be assigned to the weights. Second, it provides a way of addressing the residual ergodicity problems posed by the task of estimating a near-multicanonical distribution. The TP method allows one to combine the information obtained through a number of simulations, each exploring only a limited range of macrostate space and each therefore requiring a simulation time much smaller than the ergodic time. The method thus allows one to capitalize on computing architectures that allow many replicas to be run in parallel.

At this juncture we should note that the method we have described bears some resemblance to the method of multistage sampling [44]. If implemented in the present context, this approach would also utilize a set of simulations; each simulation would be constrained to walk (possibly multicanonically) within a narrow section of the full range of macrostates, overlapping with its neighbors. From a VS histogram, the PDF of the order parameter, within each section, would then be estimated, and by imposing continuity between the sections, the PDF could be reconstructed for the whole range of macrostates. Despite the similarities, the multicanonical approach retains several advantages. To use multistage sampling, we must decide *a priori* how to divide up the range of macrostates: how wide each section should be and how much it should overlap with its neighbors. We also must decide how to match the results from the various histograms, perhaps using just the overlapping states or perhaps a function fitted to the whole histogram. The use of the single pooled histogram of macrostate transitions in the TP approach handles all of this transparently. Moreover, to allow full equilibration (in the VS sense) over the range of macrostates in each section of the multistage sampling simulation would require that each section should contain substantially fewer macrostates even than are explored by one of the multicanonical replicas in the course of its run, thus requiring correspondingly more independent simulations.

As a final remark on the TP method, we note that the same approach may also be applied to determine the weights associated with the subensembles that feature in the expanded ensemble techniques developed by Lyubartsev *et al.* [11] and Marinari and Parisi [12].

Now let us turn to compare the multicanonical approach to the phase-boundary problem with the established integration methods. It is undeniable that, as Frenkel has argued [6], the IM approach is commendably simple, requiring little or no extension beyond the framework (for the evaluation of canonical averages) needed for the basic MC procedure. Nevertheless, the principal obstacle to any extended-sampling approach—the absence of a *systematic* way of *building* the desired sampling distribution—seems now to be surmountable and the multicanonical strategy is correspondingly more viable. It has some attractive features. First, in our view, it provides a framework in which MC error bounds are more readily assessed and controlled. The error bounds inherent in IMs are clouded by uncertainties about the influence of the integration-parameter spacing. Second, the method requires no appeal to a reference system (or systems) and is thus free of the additional uncertainties that may arise in the IMs when the integration path has to connect systems with grossly different configurational structure, in particular when the path has to traverse a near-critical region. Third, the multicanonical MC method surely provides the better framework for handling the subtle finite-size effects associated with critical points. Finally, it is couched in terms—

probability distributions rather than free energy functions—that relate most directly to what is actually measured in MC studies.

Nevertheless, it remains to be seen whether the multicanonical strategy will prove as versatile as the integration approach. In particular (a commentary on the limitations of the present study [45]) it remains to be seen whether it can be developed to deal directly with the phase boundaries (such as that between solid and liquid or between two solids of different symmetries) that involve more complex changes of microstructure.

ACKNOWLEDGMENTS

We have appreciated lively discussions with Dr. W. C.-K. Poon and are grateful to Dr. Peter Bolhuis for providing further details of the results of Ref. [18]. G.R.S. acknowledges financial support from the SERC.

[1] H.B. Callen, *Thermodynamics and an Introduction to Thermostatics* (Wiley, New York, 1985).
[2] R.P. Feynman, *Statistical Mechanics: A Set of Lectures* (Addison-Wesley, Reading, MA, 1972).
[3] J.E. Mayer and M.G. Mayer, *Statistical Mechanics* (Wiley, New York, 1941).
[4] R.A. Cowley, Adv. Phys. **12**, 421 (1963).
[5] K. Binder and D.W. Heermann, *Monte Carlo Simulations in Statistical Physics* (Springer-Verlag, Berlin, 1988).
[6] D. Frenkel, in *Molecular-Dynamics Simulation of Statistical-Mechanical Systems*, edited by G. Ciccotti and W.G. Hoover (North-Holland, Amsterdam, 1986), p 151.
[7] G.R. Smith and A.D. Bruce, J. Phys. A **28**, 6623 (1995).
[8] G.M. Torrie and J.P. Valleau, Chem. Phys. Lett. **28**, 578 (1974).
[9] B.A. Berg and T. Neuhaus, Phys. Lett. B **267**, 249 (1991); Phys. Rev. Lett. **68**, 9 (1992).
[10] B.A. Berg, Int. J. Mod. Phys. C **4**, 249 (1993).
[11] A.P. Lyubartsev, A.A. Martsinovski, S.V. Shevkunov, and P.N. Vorontsov-Velyaminov, J. Chem. Phys. **96**, 1776 (1992).
[12] E. Marinari and G. Parisi, Europhysics Lett. **19**, 451 (1992).
[13] G.R. Smith and A.D. Bruce (unpublished).
[14] A.M. Ferrenberg and R.H. Swendsen, Phys. Rev. Lett. **61**, 2635 (1988); **63**, 1658(E) (1989).
[15] N.B. Wilding, Phys. Rev. E **52**, 602 (1995).
[16] B. Grossmann, M.L. Laursen, T. Trappenberg, and U.J. Wiese, Phys. Lett. B **293**, 175 (1992).
[17] J.N.W. Lekkerkerker, W.C.-K. Poon, P.N. Pusey, A. Stroobants, and P.B. Warren, Europhys. Lett. **20**, 559 (1992).
[18] P. Bolhuis and D. Frenkel, Phys. Rev. Lett. **72**, 2211 (1994); P. Bolhuis, M. Hagen, and D. Frenkel, Phys. Rev. E **50**, 4880 (1994).
[19] C.F. Tejero, A. Daanoun, J.N.W. Lekkerkerker, and M. Baus, Phys. Rev. Lett. **73**, 752 (1994).
[20] W.C.K. Poon, J.S. Selfe, M.B. Robertson, S.M. Ilett, A.D. Pirie, and P.N. Pusey, J. Phys. (France) II **3**, 1075 (1993).
[21] S.M. Ilett, A Orrock, W.C.K. Poon, and P.N. Pusey, Phys. Rev. E **51**, 1344 (1995)
[22] M. Hagen, E.J. Meijer, G.C.A.M. Mooij, D. Frenkel, and H.N.W. Lekkerkerker, Nature **365**, 425 (1993).
[23] R. Hall, J. Chem. Phys. **57**, 2252 (1972).
[24] We quote temperatures in units of ϵ/k, volumes in units of σ^3, pressures in units of $\epsilon\sigma^{-3}$, and energies (per particle) in units of ϵ.
[25] K. Binder and D.P. Landau, Phys. Rev. B **30**, 1477 (1984).
[26] M.S.S. Challa, D.P. Landau, and K. Binder, Phys. Rev. B **34**, 1841 (1985).
[27] I. R. McDonald, Mol. Phys. **23**, 41 (1972).
[28] It is convenient to work with a *discrete* set of volume macrostates, whose construction is discussed below.
[29] We use the term *multicanonical* to describe an ensemble ''extended'' by the weighting of the macrostates of *any* macroscopic variable.
[30] See remarks following Eq. (19).
[31] D. Frenkel, in *Computer Modelling of Fluids, Polymers & Solids*, edited by C.R.A. Catlow, C.S. Parker, and M.P. Allen (Kluwer Academic, Dordrecht, 1990), p. 83.
[32] One must allow for the nonuniformity of the macrostate spacing when constructing the probability density for v from the discrete macrostate probabilities.
[33] J. J. Martin, *Bayesian Decision Problems and Markov Chains* (Wiley, New York, 1967).
[34] We assume, implicitly, that the visits are essentially uniformly distributed through the macrostates.
[35] See, however, the discussion of multistage sampling in Sec. VI.
[36] Consider, for example, the simulation of a system of $N=108$ particles. To secure an average acceptance ratio for v transitions of $\sim 1/2$ requires $N_m \sim 1.3\times 10^3$ macrostates. Simple random walk arguments then imply that on average some 2×10^6 volume updates would be required for a single simulation to traverse the whole range of macrostates. The highest speed we can achieve on the CM is 6000 sweeps per simulation per hour (for $N=108$, $N_r=64$), implying the need for over 300 h for just one traverse per replica.
[37] Indeed, although we shall proceed somewhat differently, the coexistence pressure, as well as the specific volumes of the coexisting phases, for this temperature can immediately be estimated by constructing the double tangent of the limiting multicanonical weight function.
[38] In view of the observed L dependence of the probability of interphase macrostates, discussed in Sec. V C, a simple L^d scaling of the weights (determined for a smaller system) provides a fairly good initial estimate.
[39] C. Borgs and R. Kotecký, Phys. Rev. Lett. **68**, 1734 (1992).
[40] K. Binder, Phys. Rev. A **25**, 1699 (1982).
[41] B.A. Berg, U.H.E. Hansmann, and T. Neuhaus, Z. Phys. B **90**, 229 (1993).
[42] P. Bolhuis (private communication).
[43] A.D. Bruce and N.B. Wilding, Phys. Rev. Lett. **68**, 193 (1992).
[44] J.P. Valleau and D.N. Card, J. Chem. Phys. **57**, 5457 (1972).
[45] The present study itself gives no assurance that the solid phases studied are anything more than metastable with respect to the liquid phase.