Adaptive integration method for Monte Carlo simulations

Marc Fasnacht* and Robert H. Swendsen†
Physics Department, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, USA

John M. Rosenberg‡
Department of Biological Sciences, University of Pittsburgh, Pittsburgh, Pennsylvania 15213, USA
(Received 10 November 2003; published 24 May 2004)

We present an adaptive sampling method for computing free energies, radial distribution functions, and potentials of mean force. The method is characterized by simplicity and accuracy, with the added advantage that the data are obtained in terms of quasicontinuous functions. The method is illustrated and tested with simulations on a high density fluid, including a stringent consistency test involving an unusual thermodynamic cycle that highlights its advantages.

DOI: 10.1103/PhysRevE.69.056704 PACS number(s): 05.10.Ln, 02.70.Rr, 02.70.Uu

I. INTRODUCTION

Many of the most important tasks for Monte Carlo simulations involve the computation of a distribution function, or, equivalently, an effective free energy. When investigating the temperature dependence of the specific heat or magnetic susceptibility in a magnetic model, one of the most efficient approaches is to compute a histogram as a function of energy and to then calculate the desired property as a continuous function of temperature from those data [1–4]. We might also want to compute a distance-dependent distribution function to compare with x-ray scattering studies, or we might be interested in the free energy of a particle as a function of its interaction with its surroundings as part of a calculation of its chemical potential.

In all cases mentioned, a standard difficulty is that the probability for a straightforward Monte Carlo simulation is often much too low to be practical to sample all parameter values of interest. It is well known that one of the most efficient ways of getting around this problem is to perform a multicanonical simulation [5], in which the low probability of an inaccessible region is compensated by adding an artificial effective potential. The effect of the added potential can then be subtracted out to calculate the desired information.

If the added potential is exactly equal to the desired free energy, the free energy barriers to inaccessible regions will be completely eliminated. Unfortunately, if we knew this free energy, we would not have to do the computation in the first place.

A number of adaptive methods have been developed to solve the problem of calculating an appropriate compensating potential [6–10]. The main idea is to obtain data about the desired free energy during the course of the Monte Carlo simulation itself. As the simulation progresses, the compensating potential in the simulation is changed as it becomes better known. Technically, this approach violates detailed balance. However, detailed balance is restored asymptotically as the adaptive potential converges to the correct function, and many studies have shown that the procedure is effective [9].

The question remains of how to perform adaptive sampling most efficiently. This paper presents a method of adaptive sampling that synthesizes the advantages of various earlier methods. It has several key features.

(1) Instead of calculating a histogram that is proportional to the probability of a parameter of interest, we calculate derivatives of this probability. This is based on the theoretical point that the information on which a Monte Carlo step is made involves only the relative probabilities of the two states; it does not involve the absolute probabilities. Consequently, the most efficient strategy should involve the direct estimates of ratios between neighboring bins, with the full probability distribution being constructed at the end of the calculation. These considerations have been validated by the success of the broad histogram method [11] and transition matrix Monte Carlo [10].

(2) The use of additional information from details of the configurations has proved very effective in increasing the efficiency of the broad histogram method and transition matrix Monte Carlo. Our method incorporates such information whenever possible.

(3) Yan and de Pablo have recently presented an interesting paper in which they showed some of the advantages of working with the derivative of the free energy [12]. They phrased their method in terms of an effective temperature, but we would like to expand the concept to include derivatives with respect to any parameter of interest. In addition to a derivative with respect to energy, we are particularly interested in derivatives with respect to nonlinear parameters in the interactions between particles, separations between particles, and generalized reaction coordinates.

A particularly important advantage of using derivatives is that the bins can be made very small without losing accuracy, as we will discuss in Sec. III. Since the generalized free

*Present address: Howard Hughes Medical Institute, Department of Biochemistry and Molecular Biophysics, Columbia University, New York, NY 10032. Electronic address: mf5k@andrew.cmu.edu; mf2206@columbia.edu
†Electronic address: swendsen@andrew.cmu.edu
‡Electronic address: jmr@jmr3.xtal.pitt.edu

energy of interest is reconstructed at the end of the calculation by numerical integration, making the bins smaller can produce a smoother function and greater detail, but will not introduce any distortion. This is of particular importance in biological applications where details of the changes in free energy can help with understanding the underlying processes. This is in contrast to the usual histogram of values that are proportional to the desired probability. If the bins are made smaller, the relative statistical fluctuations become larger, and the signal can disappear into the noise.

The proposed adaptive integration method (AIM) includes the computation of derivatives of the desired generalized free energy over a parameter region of interest with a large number of bins and small bin size. The updated estimates for the derivatives are then immediately incorporated into the Monte Carlo simulation to reduce the effective barriers and improve efficiency. This direct use of derivatives not only makes a direct use of important information available from the simulation, but it also highlights the close connection between free energies or potentials of mean force, and the generalized forces experienced by the particles during the simulation. This deep connection both enables us to achieve high accuracy and to increase the sensitivity of the results to details of the physical behavior. We will return to this point when we discuss the information contained in the average generalized forces.

The potential of mean force is well named because its gradient is the ensemble-averaged force. This makes it particularly useful in biophysical applications, both for the insight it is capable of providing and because these forces are becoming amenable to experimental manipulation.

We also introduce a different way of carrying out the sampling for radial distribution functions by doing the actual simulation without any direct interaction between the particles of interest. We show that the direct interaction has a purely additive effect on the total potential of mean force (pmf), so it can be added in at the end of the calculation. This enables us to do a single Monte Carlo simulation that allows us to generate potentials of mean force for a wide variety of direct interactions.

The goal of the calculations presented in this paper is to provide clear illustrations and tests of the method. To do this, we need a system that is computationally simple in order to study the convergence properties and accuracy of the method. This allows us to do multiple long runs to gather good statistics for an accurate description of the behavior. The test system should also be well understood, so that we can verify the correctness of our results. It is essential to establish a method on such a system before applying it to physically or biologically more interesting systems, which tend to be larger and more complex. A good test system is essential to verify that the method really works correctly, and to ensure that there are no undetected systematic errors.

The two-dimensional Lennard-Jones (LJ) fluid is an ideal model for our purpose. The pairwise potential has the form

$$\phi_{LJ}(r) = 4\epsilon\left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6}\right], \tag{1}$$

which makes it computationally simple. It is well understood, but still exhibits behavior complex enough to provide challenging problems [13–21]. As we will show, it allows us to test the efficiency of our methods by applying them in different ways.

We have performed a variety of tests. The first tests concern the calculation of a radial distribution function for a dense liquid. This is the simplest example of a more general calculation of potentials of mean force, which involve finding the free energy along a reaction coordinate. This is a very common type of calculation in biological systems [22,23].

The next set of tests involves the calculation of the chemical potential. The basic approach here is well known, but the advantage of AIM is the production of quasicontinuous integration curves that reduce numerical integration errors.

In the final set of tests we compute a thermodynamic cycle in which we can apply AIM in different ways and examine the accuracy and consistency of the method. The calculations on the cycle illustrate the advantages of AIM for the direct calculation of the chemical potential. They also show how to obtain the chemical potential through two independent paths while illustrating the usefulness of the solvent-induced potential of mean force, which is computed without any direct interaction between particles.

In the following section, we will discuss the general formalism for adaptive integration with histograms, and then explain the adaptive integration method in Sec. III. Explicit calculations of the radial distribution function will be given in Sec. IV. Sec. V contains the description of our analysis of a thermodynamic cycle, the first part of which also contains the calculation of the chemical potential.

II. ADAPTIVE METHODS

Adaptive methods are designed to find the free energy of a system as a function of some parameter λ, or to determine the probability distribution of an observable of interest.

For free energy problems, the parameter λ could be part of the original problem of interest, or it could be contained in a term in the Hamiltonian that has been added or modified for convenience. One example would be an overall multiplicative factor for the energy, in which case we would have $U(\lambda,\mathbf{X}) = \lambda U(\mathbf{X})$. Another example would be a parameter in a two-particle interaction energy, either between all pairs or between two specific particles. In simulations of biological molecules, the parameter could describe motion along a general reaction coordinate.

If the system is described by an internal energy $U(\lambda,\mathbf{X})$, then the probability of a microstate $\mathbf{X}$ is given by

$$P(\mathbf{X}|\lambda) = \frac{e^{-\beta U(\lambda,\mathbf{X})}}{\int_V e^{-\beta U(\lambda,\mathbf{X})} d\mathbf{X}}, \tag{2}$$

the partition function is

$$Z(\lambda) = \int_V e^{-\beta U(\lambda,\mathbf{X})} d\mathbf{X} \tag{3}$$

and the associated generalized free energy is

$$F(\lambda) = -\beta^{-1}\ln[Z(\lambda)]. \tag{4}$$

In the equations above $\beta = 1/k_B T$, where T is the temperature of the system and k_B is the Boltzmann constant.

As an example of the computation of the probability distribution for an observable of interest, we will discuss the radial distribution function for the separation $r_{ij}(\mathbf{X})$ between particles i and j. The parameter λ we are interested in is the separation between particles i and j. From Eq. (2), the distribution is then given by

$$P_{ij}(r) = \frac{\int_V e^{-\beta U(\mathbf{X})}\delta(r - r_{ij}(\mathbf{X}))d\mathbf{X}}{Z}. \tag{5}$$

Naturally, the same formalism applies with obvious changes to any reaction coordinate or observable of interest.

In each case, adaptive methods seek an approximation to $F(\lambda)$ or $P_{ij}(r)$, which we will denote as $\hat{F}(\lambda)$ and $\hat{P}_{ij}(r)$. As the simulation progresses, the estimates are modified by using the information obtained to smooth the effective interactions. In the case of the free energy, the simulated probability changes from Eq. (2) to being proportional to

$$\exp[-\beta U(\lambda, \mathbf{X}) + \beta\hat{F}(\lambda)] \tag{6}$$

so that the marginal distribution for λ becomes proportional to

$$\exp\{-\beta[F(\lambda) - \hat{F}(\lambda)]\}. \tag{7}$$

As $\hat{F}(\lambda)$ approaches $F(\lambda)$, this distribution becomes flat, making the simulation more efficient.

For the case of the radial distribution function, the probability of the configurations becomes proportional to

$$\frac{e^{-\beta U(\lambda, \mathbf{X})}}{\hat{P}(r_{ij}(\mathbf{X}))} \tag{8}$$

which also has the effect of smoothing the distribution of the distance $r_{ij}(\mathbf{X})$.

The adaptive integration method we present in the following section is based on the principles given in the Introduction and has several advantages over previous methods.

III. ADAPTIVE INTEGRATION METHOD

To give a concrete description of the AIM, we will first discuss the method for a calculation of the free energy. Assuming that the internal energy is a continuous function of λ, $U = U(\lambda, \mathbf{X})$, we can formally calculate the derivative of the free energy with respect to λ:

$$\frac{dF(\lambda)}{d\lambda} = -\frac{-\beta\int_V e^{-\beta[U(\lambda,\mathbf{X})]}\dfrac{dU(\lambda,\mathbf{X})}{d\lambda}d\mathbf{X}}{\beta Z} \tag{9}$$

or

$$\frac{dF(\lambda)}{d\lambda} = \left\langle \frac{dU(\lambda,\mathbf{X})}{d\lambda} \right\rangle_\lambda, \tag{10}$$

where the notation $\langle\cdot\rangle_\lambda$ indicates that the average is taken for a fixed value of λ.

In all adaptive sampling methods, intervals of values of λ are grouped together to divide up the continuous range of λ into bins, so that there are a finite number of quantities to calculate. An advantage of our method is that we are able to greatly increase the number of bins and decrease the bin size, so that we can obtain much finer resolution, as discussed below.

The average in Eq. (10) can easily be obtained from simulations. The free energy as a function of λ can then be calculated by integrating with respect to λ:

$$F(\lambda) = \int \left\langle \frac{dU(\lambda,\mathbf{X})}{d\lambda} \right\rangle_\lambda d\lambda. \tag{11}$$

This last equation and Eq. (10) are well known and form the basis of the thermodynamic integration method (see e.g., Ref. [24]).

In order to calculate the integral in Eq. (11), we carry out a Monte Carlo (MC) simulation that includes moves in λ. By recording the average $\langle dU/d\lambda \rangle_\lambda$ for all values (bins) of λ in the range of interest, we can reconstruct an estimate for the free energy differences between two values of λ, $\Delta\hat{F}(\lambda)$ by simple numerical integration.

The acceptance probability for a change of λ between an old value λ_o and a new value λ_n is then modified to

$$p(\lambda_o \to \lambda_n) = \min[e^{\{-\beta[U(\lambda_n,\mathbf{X})-U(\lambda_o,\mathbf{X})]+\beta[\hat{F}(\lambda_n)-\hat{F}(\lambda_o)]\}}, 1]. \tag{12}$$

MC moves between configurations are done in the usual way, using the current value of $\hat{F}(\lambda)$.

Similarly, if we are interested in calculating a radial distribution function, or any generalized free energy as a function of some reaction coordinate, we would modify the acceptance probability to be

$$p(\mathbf{X}_o \to \mathbf{X}_n) = \min\left[e^{-\beta[U(\mathbf{X}_n)-U(\mathbf{X}_o)]}\frac{\hat{P}(r_{ij}(\mathbf{X}_o))}{\hat{P}(r_{ij}(\mathbf{X}_n))}, 1 \right], \tag{13}$$

where $\mathbf{X}_o$ and $\mathbf{X}_n$ are the old and new configurations of the system. The AIM approach can be used for any quantity for which we can calculate a derivative. This makes it unsuitable for problems in which the relevant derivatives would diverge such as the calculation of the radial distribution function for a hard sphere gas. However, for most systems of interest with continuous degrees of freedom, derivatives are well behaved and the method is applicable.

One advantage of this approach is that we can obtain almost continuous estimates of the free energy as a function of the parameter of interest. Since we are integrating in the method, finer binning does not result in decreased accuracy. For the adaptive integration method, all the values between

two points are taken into account when calculating the difference in biasing potential, independent of how many bins are in between. So the bins can be made almost arbitrarily small, the main limitation being memory for storing all the values. This is not true for methods that use a biasing potential based on probabilities: For these methods, the potential difference will depend on the number of hits per bin. The fluctuations in that number increase with decreasing bin size. The errors for the potential will therefore increase as we make the bin size smaller and the biasing potential will converge more slowly. This is not only important for the smoothness of the final potential, but also during the simulation for the biasing potential, since these methods satisfy detailed balance only asymptotically when the potential has converged.

Another obvious advantage of AIM is that the calculation can easily be divided up into different windows, each of which is covered by an independent calculation. This is often an advantageous reformulation of the implementation of adaptive algorithms, since the characteristic time for a random walk in energy space is at least proportional to the square of the range of energies. Since the full range of energies is generally proportional to the number of particles or spins in the system, this can be quite important. Because AIM calculates derivatives, there is no matching problem in splicing together the results of independent simulations. Although there are good methods for dealing with the problem of joining histograms [1–3,25] it is still better not to have the problem at all.

A recent paper by Yan and de Pablo [12] introduces a measurement technique that is closely related to the method described in this section. The random-walk method they presented for taking data from a simulation is a special case of the general formalism introduced here, although they used a different representation of the derivative of the free energy than would arise directly from our formulation. They used an integral representation of the inverse temperature to calculate the density of states in a fluid. However, they did not use those data as part of the adaptive simulation, relying on the Wang-Landau method instead. They also restricted themselves to the special case mentioned, rather than generalizing the approach as done here. The results presented in Sec. IV A lead us to believe that the already excellent results of Yan and de Pablo's random-walk algorithm could be further improved by using derivative methods for directing the adaptive sampling as well as for the data collection.

There are also related approaches based on molecular dynamics and a generalization of thermodynamic integration [26–31]. The basic idea is that free energy can be seen as the potential of mean constraint force. The theoretical formulation in Ref. [31] is similar to the method we describe here. The methods give excellent results. However, the fact that in molecular dynamics we have to keep track of the dynamics of the system and have to introduce constraint forces makes these methods much more complicated than the Monte Carlo framework we propose.

A. Potentials of mean force

Consider a system of N particles interacting with a pairwise interaction potential ϕ. The total potential energy of the system is given by

$$U(\mathbf{X}) = \sum_{i<j} \phi(\mathbf{x}_i - \mathbf{x}_j). \tag{14}$$

The probability for a certain configuration $\mathbf{X}$ in the NVT ensemble is given by

$$P(\mathbf{X})_{\mathrm{NVT}} = \frac{e^{-\beta U(\mathbf{X})}}{\int e^{-\beta U(\mathbf{X})} d\mathbf{X}} \tag{15}$$

$$= \frac{e^{-\beta U(\mathbf{X})}}{Z}, \tag{16}$$

where Z is the partition function.

We want to calculate the pmf between two particles i and j. We define $\mathbf{r}_{ij} \equiv \mathbf{x}_i - \mathbf{x}_j$, where $\mathbf{x}_i$ and $\mathbf{x}_j$ are the respective positions of the particles, and $r \equiv \|\mathbf{r}_{ij}\|$. The potential of mean force $\hat{\phi}$ is the potential that would generate a probability distribution along a certain degree of freedom equal to the marginal distribution of the full probability distribution with respect to that degree of freedom. It is related to the pair distribution function $g(r)$ by

$$\hat{\phi}_{ij}(r) = -\frac{1}{\beta} \ln[g(r)]. \tag{17}$$

The pair distribution function is defined as the probability $P(r)$ of finding a pair of atoms a distance r apart, relative to the probability $P_{\text{ideal gas}}(r)$ expected for a completely random distribution of an ideal gas at the same density [32]:

$$g_{ij}(r) \equiv \frac{P(r)}{P_{\text{ideal gas}}(r)}. \tag{18}$$

$P_{\text{ideal gas}}(r)$ is given by

$$P_{\text{ideal gas}}(r) = \frac{r^{d-1} C_d}{V}, \tag{19}$$

where V is the total volume of the system, d is the number of dimensions, and C_d is the surface of a unit sphere in d dimensions. The probability of two particles in the interacting system being at a distance r is given by

$$P(r) = \frac{\int e^{-\beta U(\mathbf{x}_i, \mathbf{x}_j, \mathbf{X}_{N-2})} \delta(|\mathbf{r}_{ij}| - r) d\mathbf{x}_i d\mathbf{x}_j d\mathbf{X}_{N-2}}{Z}, \tag{20}$$

where $\mathbf{X}_{N-2}$ are the position vectors of all the particles other than i and j and $d\mathbf{X}_{N-2} = \Pi_{k \notin \{i,j\}} d\mathbf{x}_k$. Without loss of generality, we can set the origin of our coordinate system at $(\mathbf{x}_i - \mathbf{x}_j)/2$. Then $\mathbf{x}_i = \mathbf{r}_{ij}/2$ and $\mathbf{x}_j = -\mathbf{r}_{ij}/2$. Since we have translational invariance, we can integrate over all possible positions of the origin. Similarly, integrating over all orientations of $\mathbf{r}_{ij}$ is equivalent to rotating all the other particles, so the angular degrees of freedom can be integrated out. It follows

$$P(r) = \frac{V \int e^{-\beta U(r\hat{\mathbf{r}}_{ij}/2, -r\hat{\mathbf{r}}_{ij}/2, \mathbf{X}_{N-2})} r^{d-1} C_d d\mathbf{X}_{N-2}}{Z}. \quad (21)$$

Using Eqs. (18) and (19), and the fact that r is constant for the integration, we find

$$g(r) = \frac{V^2 \int e^{-\beta U(r\hat{\mathbf{r}}_{ij}/2, -r\hat{\mathbf{r}}_{ij}/2, \mathbf{X}_{N-2})} d\mathbf{X}_{N-2}}{Z} \quad (22)$$

$$= \frac{V^2 \int e^{-\beta U(\mathbf{X})} d\mathbf{X}_{N-2}}{Z} \quad (23)$$

$$= K \int e^{-\beta U(\mathbf{X})} d\mathbf{X}_{N-2}. \quad (24)$$

From Eq. (17) it follows that the derivative of $\hat{\phi}_{ij}$ with respect to r is

$$\frac{d\hat{\phi}_{ij}(r)}{dr} = -\frac{dg(r)/dr}{\beta g(r)} \quad (25)$$

and from Eq. (24)

$$\frac{dg(r)}{dr} = -\beta K \int e^{-\beta U(\mathbf{X})} \left[\frac{(\mathbf{F}_j - \mathbf{F}_i) \cdot \hat{\mathbf{r}}_{ij}}{2} \right] d\mathbf{X}_{N-2}. \quad (26)$$

$\mathbf{F}_i$ and $\mathbf{F}_j$ are the total forces on particles i and j, respectively, given by

$$\mathbf{F}_i = -\nabla_i U = -\sum_{k=1}^{N} \nabla_i \phi(\mathbf{x}_i - \mathbf{x}_k). \quad (27)$$

This means that the derivative of the potential of mean force is given by

$$\frac{d\hat{\phi}_{ij}(r)}{dr} = \frac{\int e^{-\beta U(\mathbf{X})} \frac{(\mathbf{F}_j - \mathbf{F}_i) \cdot \hat{\mathbf{r}}_{ij}}{2} d\mathbf{X}_{N-2}}{\int e^{-\beta U(\mathbf{X})} d\mathbf{X}_{N-2}} \quad (28)$$

$$= -\left\langle \frac{(\mathbf{F}_i - \mathbf{F}_j)}{2} \cdot \hat{\mathbf{r}}_{ij} \right\rangle_r \quad (29)$$

$$= \left\langle \frac{dU}{dr} \right\rangle_r. \quad (30)$$

This last expression says that we only need to calculate the average forces on each particle projected onto the line between the two particles.

Given this, we recover the potential of mean force by integrating with respect to r:

$$\hat{\phi}_{ij}(r) = \int \left\langle \frac{dU}{dr} \right\rangle_r dr. \quad (31)$$

Equations (28) and (31) underscore the intuitive physical interpretation of the pmf as the negative integral of the average forces.

Volume normalization and uniform sampling. In more than one dimension, sampling will not be uniform as a function of distance if we just subtract out the potential of mean force. The particles are more likely to be at larger separations due to the larger available volume. If we want to sample distances uniformly, we have to use an additional bias term that corrects for this. The probability for particles to be at a certain distance r, if they do not interact and are uniformly distributed, is given by Eq. (19). To achieve uniform sampling in r, we add a term of the form $(d-1)\ln(r)/\beta$ to the biasing potential.

B. The chemical potential

The chemical potential of a fluid can be calculated in a number of ways. In the particle insertion method proposed by Widom [33], a particle is randomly inserted into the fluid to sample the interaction energy ΔU of the test particle with the rest of the fluid. The average $\langle \exp(-\beta \Delta U) \rangle$ can be used to calculate the excess chemical potential $\mu' \equiv \mu - \mu_{\text{ideal gas}}$ using $\mu' = -k_B T \ln\langle \exp(-\beta \Delta U) \rangle$. The average $\langle \cdot \rangle$ here is over all configurations of the fluid without the test particle and all random positions of the test particle. This method works well for dilute fluids. For dense fluids, however, the averages will be dominated by sampling of physically irrelevant configurations where the test particle overlaps with another particle, so that sampling is not very accurate. One way around this is to insert the test particle gradually by incrementally increasing the interaction of the test particle with the rest of the fluid, as described by Mon and Griffiths [34]. The chemical potential can then be calculated by using thermodynamic integration (see, e.g., Ref. [24]) over the different stages.

This formalism can be adapted to calculate the excess chemical potential of a particle using the adaptive integration method. In order to do this, we could modify the potential energy such that all the interactions of a certain particle k are scaled by a factor of $\lambda \geqslant 0$. The modified total potential energy is given by

$$U_\lambda(\mathbf{X}) = \sum_{i<j; i,j \neq k} \phi(r_{ij}) + \lambda \sum_{i; i \neq k} \phi(r_{ki}) \quad (32)$$

$$= U_0(\mathbf{X}) + \lambda U_1(\mathbf{X}). \quad (33)$$

To simplify notation, we will denote the first term in the potential energy as U_0 and the second term as λU_1. The free energy, as a function of λ, of our system is

$$F(\lambda) = -\frac{1}{\beta} \ln(Z_\lambda) \quad (34)$$

$$= -\frac{1}{\beta}\ln\left(\frac{\int e^{-\beta[U_0(\mathbf{X})+\lambda U_1(\mathbf{X})]}d\mathbf{X}}{\Lambda^{dN}N!}\right), \tag{35}$$

where $\Lambda=(\beta h^2/2\pi m)^{1/2}$ is the de Broglie wavelength of the particles, which results from integrating over the momentum degrees of freedom. The factor of $N!$ is a result of having N particles of the same species. It follows that

$$\frac{dF(\lambda)}{d\lambda} = -\frac{1}{\beta}\frac{d}{d\lambda}\ln(Z_\lambda) \tag{36}$$

$$= -\frac{-\beta\int e^{-\beta[U_0(\mathbf{X})+\lambda U_1(\mathbf{X})]}U_1(\mathbf{X})d\mathbf{X}}{\beta\Lambda^{dN}N!\,Z_\lambda} \tag{37}$$

$$= \langle U_1\rangle_\lambda. \tag{38}$$

Integrating this from $\lambda=0$ to $\lambda=1$ then gives the free energy change due to interactions when adding an extra particle, which is the definition of the excess chemical potential μ'. The number of particles is constant. To get the total chemical potential μ, we need to adjust the number of particles, corresponding to the ideal gas component $\mu_{\text{ideal gas}}$ which is known exactly: $\mu_{\text{ideal gas}}=k_BT\ln(N/V)$.

A well-known problem with this approach is poor convergence [34,35]. The difficulty is that U_1 and the derivative that we wish to calculate diverge as $r\to 0$. This makes convergence near $r=0$ very difficult, although the method works well for other ranges of values.

A way around this divergence problem is to change the λ dependence of the potential energy $U(\lambda)$. As long as the two limits $\lambda=0$ and $\lambda=1$ evaluate to the right values, we can choose any well-behaved parametrization [34]. We have chosen to rescale the interaction radius of the test particle so that it interacts with other particles with

$$\phi_\lambda(r) \equiv \phi_{LJ}\left(\frac{r}{\lambda}\right) = 4\epsilon\left[\left(\frac{\lambda\sigma}{r}\right)^{12} - \left(\frac{\lambda\sigma}{r}\right)^{6}\right]. \tag{39}$$

Then

$$U_\lambda(\mathbf{X}) = \sum_{i<j;i,j\neq k}\phi(r_{ij}) + \sum_{i;i\neq k}\phi_\lambda(r_{ki}). \tag{40}$$

With this parametrization, reducing λ corresponds to decreasing the interaction radius of the test particle with other particles. The resulting free energy $F(\lambda)$ is well behaved, with finite derivatives even at $\lambda=0$. Another advantage of this formulation is that we have a direct relation between the chemical potential and the particle size, so we need to do only one calculation for different test particle sizes.

IV. ADAPTIVE METHODS APPLIED TO THE RADIAL DISTRIBUTION FUNCTION

In order to test the simulation methods introduced in the preceding section, we have applied them to the Lennard-Jones model of a simple fluid to calculate the radial distribution function, and consequently the effective potential of mean force, between two selected particles.

The choice was based on several considerations. The radial distribution function is a simple example of a more general class of effective potentials of mean force, which means that this calculation demonstrates the implementation of a wide variety of problems of interest. The model is relatively simple and well understood. The computational effort involved in making convincing tests is not extravagant.

Finally, there is also an easily computed standard by which to judge the results. For a Lennard-Jones fluid, the radial distribution function can be calculated very accurately by taking data from all pairs of particles instead of just two, because all particles are identical. This will increase the available data by a factor of $N(N-1)/2$. In addition, the calculation of an effective potential from just two particles is fundamentally challenging because the diffusion of particles is slow. To obtain an accurate radial distribution function, it is necessary for the two particles to diffuse over large distances relatively rapidly. This problem does not arise if all pairs of particles are used in the calculation and the pmfs obtained converge very rapidly. For these reasons we use the potential of mean force obtained from an average over all pairs of particles as reference pmf to which to compare our results below.

We have to emphasize, however, that this approach of averaging over all particles to obtain the pmf is only valid as long as all particles are identical. The adaptive integration method introduced is aimed at cases where this is not true. We chose a system of identical particles merely as a test case to evaluate the method.

The performance of AIM is compared to a state of the art method, the Wang-Landau method. The Wang-Landau method has received a lot of attention since its introduction because of its high efficiency and flexibility. Apart from the original work on spin lattice systems [9,36], it has also been used for sampling along chemical reaction coordinates [37], the study of phase transitions in fluids [38], and the simulation of proteins [39,40].

We performed the tests for a dense two-dimensional Lennard-Jones fluid with $\rho^*=0.8$ and $T^*=0.625$, since calculations of potentials of mean force at high density are difficult due to long diffusion times and therefore make for a stringent test case. We used a relatively small system of $N=36$ particles to allow for long runs with high statistics. In all cases, we use a square box with periodic boundary conditions.

Conformations of the fluid were sampled using a Metropolis Monte Carlo scheme, with acceptance probabilities modified according to the respective algorithm (Wang-Landau or AIM). The pseudorandom numbers used in the calculations were generated using an implementation of the the ran2() pseudorandom number generator described in Ref. [41]. This pseudorandom number generator, which is based on a paper by L'Ecuyer [42], has a very long period ($>2\times10^{18}$) and it is generally agreed to have very good properties [41].

A. AIM compared to Wang-Landau method

We performed 15 simulation runs of 10^7 Monte Carlo sweeps (after equilibrating the system) for both AIM and the

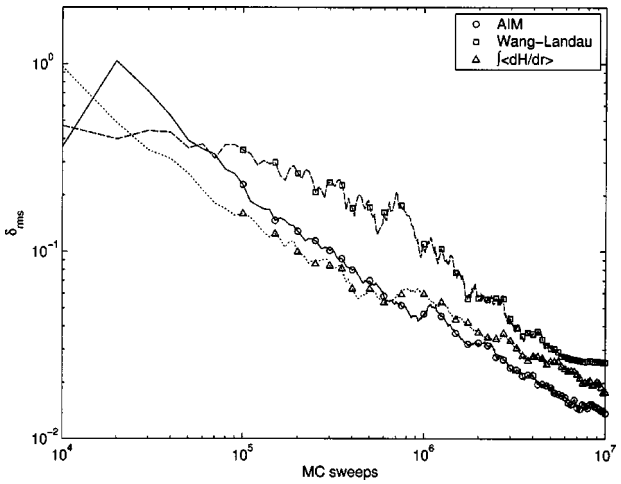

FIG. 1. Average rms difference δ_{rms} of the pmf calculated by different methods from a reference pmf (see Sec. IV) as a function of number of MC sweeps. δ_{rms} is averaged over 15 runs for the Wang-Landau (squares) and adaptive integration methods (circles). The third line (triangles) shows the error when calculating the pmf using derivatives during the simulation with the Wang-Landau method.

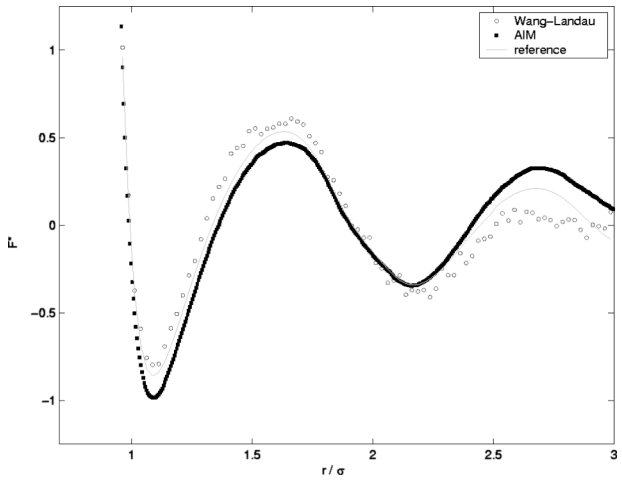

FIG. 2. Potential of mean force after 5×10^5 MC sweeps. The filled squares are from adaptive integration, the circles from Wang-Landau. The continuous gray line is the reference potential. The curves on the figure represent raw data, not smoothed curves.

Wang-Landau Method and used each algorithm to calculate the potential of mean force between two test particles.

To compare the errors for the potential obtained with AIM and the Wang-Landau method, we calculated the rms difference δ_{rms} between the corresponding pmf and the reference pmf from an average over all particles (see Sec. IV). For the Wang-Landau method, we performed preliminary runs with different settings of the parameters. The accuracy of the method, as measured by the rms difference from our reference pmf, did not vary significantly for settings close to the values suggested in the original papers by Wang and Landau [9,36]. For data shown below we chose an initial multiplication factor of $f=1.5$, with 0.8 as the threshold for flatness of the histogram. With these settings the Wang-Landau method converged after roughly 10^7 steps. Beyond this point, the multiplication factor f is too small to improve the pmf effectively [36].

Figure 1 shows the rms difference δ_{rms}, our measure of error for the potentials, as a function of number of MC sweeps for the different methods. The squares represent the results for the Wang-Landau algorithm. The circles show the results for AIM. The triangles were obtained by calculating the potential by integrating the average force while running the Wang-Landau algorithm, which is essentially the approach used by Yan and de Pablo [12].

The results indicate that the adaptive integration method converges somewhat faster than the Wang-Landau method. It also shows that even when using the Wang-Landau algorithm, the accuracy of the potential can be increased by recording the average force: The average error at the end of 10^7 MC sweeps for the adaptive integration method was $\delta_{rms}=0.0137\pm0.0019$. For the Wang-Landau calculations we obtained $\delta_{rms}=0.0256\pm0.0047$ and for the calculation of the pmf from derivatives while running with Wang-Landau we found $\delta_{rms}=0.0177\pm0.0022$. All numbers are in units of reduced energy. To show the nature of the errors more, Fig. 2 depicts the pmf at an early stage in the simulation. After only 5×10^5 MC sweeps, the potentials have not converged completely. This makes it easier to see that the curves obtained using AIM are almost continuous, since we were able to use very small bin sizes. The smoother potentials are valuable when exploring the local structure of the potential of mean force.

Naturally, AIM also automatically records the derivative of the pmf, which can be useful for determining the location of peaks. As can be seen from Fig. 2, derivatives would be more difficult to calculate from the results of the Wang-Landau method.

From these data we can see that AIM converges faster than the Wang-Landau method in terms of Monte Carlo sweeps. The actual execution time for the simulations is almost the same for both methods. The additional steps needed for either algorithm take up only an insignificant fraction of the total computation time, as the most time-intensive part of the calculation is the calculation of the potential energy differences.

B. Modified direct interaction

We have used AIM to calculate the potential of mean force between two special particles in the unusual case in which they have a zero direct interaction potential and only feel the effect of the solvent particles. An interesting feature of this calculation is that it captures all the information about the effect of the solvent on the behavior of the two special particles, even when they do have a direct interaction. We will refer to this part of the potential of mean force as the *solvent-induced pmf*.

Once the solvent-induced pmf is known, we can determine the potential of mean force for any other direct interaction between these two particles without any further simulations. The direct interaction ϕ_{ab} between the special

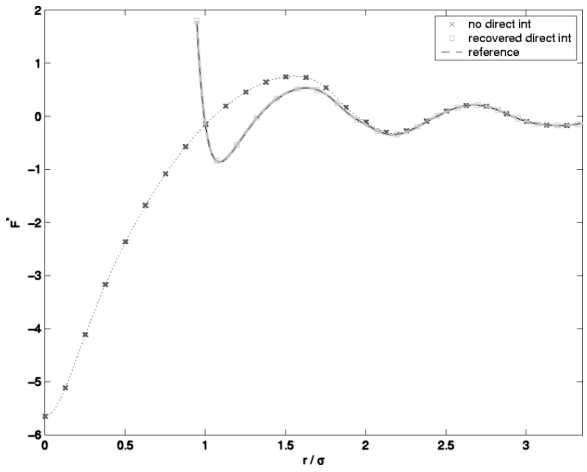

FIG. 3. Potential of mean force for regular Lennard-Jones particles. The dotted curve with × markers shows the solvent-induced pmf. The black dashed curve is a reference potential calculated from the radial distribution functions of all the particles (see above). The gray curve with square markers was calculated by adding the direct Lennard-Jones interaction to solvent-induced pmf. The black and gray curves are essentially on top of each other.

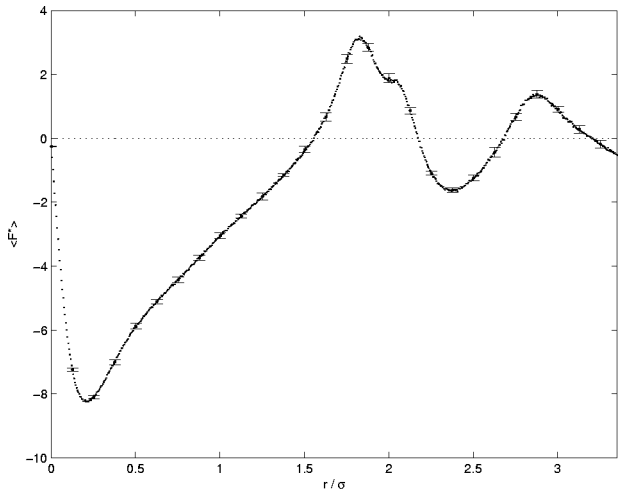

FIG. 4. Average, solvent-induced force on the special particles as a function of separation. The plot shows the average values of the force over ten independent MC runs of 10^7 MC sweeps each, with different initial random numbers. We plot the individual points without connecting them, so that fluctuations between neighboring bins are more easily visible. The error bars, shown for a subset of the points, indicate the standard deviation of the values from different simulation runs. The standard deviations for the remainder of the points are of comparable size. The solvent-induced pmf shown in Fig. 3 was obtained by integration of the forces used to calculate the average force shown here. (The integration in Fig. 3 was done for each MC run to obtain individual pmfs before averaging over the ten simulations to obtain the final plot.)

particles a and b is an additive term in the total potential energy:

$$U(\mathbf{X}) = \sum_{i<j} \phi(r_{ij}) \tag{41}$$

$$= \phi_{ab}(r_{ab}) + \sum_{i<j;i,j\notin\{a,b\}} \phi(r_{ij}) \tag{42}$$

$$= \phi_{ab}(r_{ab}) + U^*(\mathbf{X}), \tag{43}$$

where U^* consists of the remaining terms. It is then easy to see that

$$\hat{\phi}_{ab}(r_{ab}) = \phi_{ab}(r_{ab}) + \int \left\langle \frac{dU^*}{dr} \right\rangle_r dr. \tag{44}$$

Since the average in the second term does not depend on the direct interaction and is for a fixed separation of the special particles, we need only calculate it for the case in which the special particles have no direct interactions. This allows us to sample all possible separations, including those in which the special particles overlap strongly. Any direct interaction of interest can then be added analytically to the solvent-induced pmf.

C. Recovery of the original radial distribution function

Naturally, we can recover the usual Lennard-Jones radial distribution function from the solvent-induced pmf, which provides an important consistency check for the method. In this case, all the particles, including the special particles, interact with the same potential. For this Monte Carlo simulation, we performed 10 runs with 10^7 MC sweeps each for a 36-particle Lennard-Jones fluid with $\rho^*=0.8$ and $T^*=0.625$.

The dotted curve with × markers in Fig. 3 shows the solvent-induced pmf. The potential was calculated by integrating the average force, which is shown in Fig. 4. The gray curve with squares as markers on the same figure was obtained by adding the Lennard-Jones potential to the solvent-induced pmf.

The black dashed curve is a reference potential of mean force obtained from using all pairs of particles. The error bars shown are for an equally spaced subset of the data points, and indicate the variation of the calculated pmfs from different runs. The errors for the remaining data points are similar, but were omitted to make the plot more readable.

Since the difference between the reference potential and the recovered potential is too small to be visible in Fig. 3, we plot the difference between the two curves for a subset of the points in Fig. 5. The deviations are quite small compared to the actual function value and well within the error bars, with the single exception of the first data point, with $r=0.95\sigma$. Since for small r values the repulsive part of the Lennard-Jones potential becomes large, the standard simulation using all particles has relatively low statistics. In this region, the reconstructed potential is actually more accurate than the reference pmf. We note that the deviations from zero between neighboring values of r are strongly correlated in Fig. 5. This is not surprising, given that the values are calculated by integration to give smooth curves.

The plot of the average forces in Fig. 4 allows us to accurately determine the positions of the maxima and minima of the potential of mean force. It also shows interesting structure in the variation of the potential, which is hard to see in

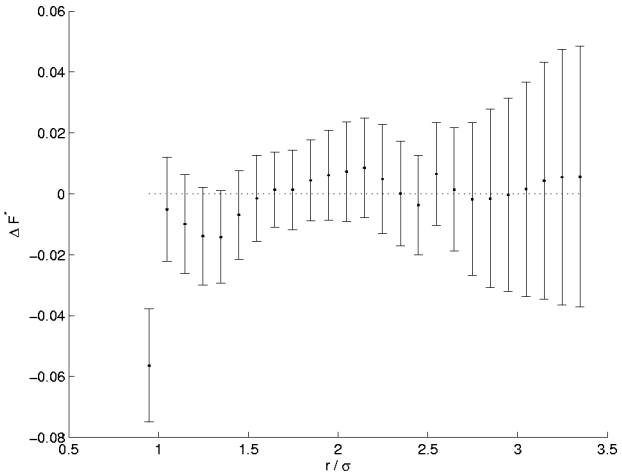

FIG. 5. Difference between the reconstructed potential of mean force and the reference pmf (gray curve with squares and black dashed curves in Fig. 3). The reconstructed pmf was calculated by adding the direct Lennard-Jones interaction to the solvent-induced pmf. The values are averaged over ten MC runs. The error bars represent the standard deviation of the potentials of mean force from different runs. The differences represent a small fraction of the actual function values. The deviations between neighboring values are correlated because the pmf is calculated by integration. The largest difference is for $r<\sigma$, where the reference potential is difficult to calculate, because this is a very low probability configuration.

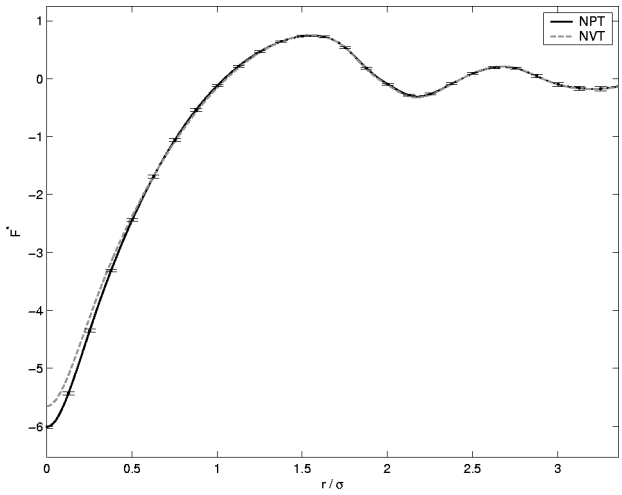

FIG. 6. Comparison of potential solvent-induced pmf for regular Lennard-Jones particles for the NVT (gray dashed curve) and NPT ensembles (black continuous curve with error bars). The curves are very similar. The main difference is the lower value of the pmf for overlap of the two special particles ($r=0$) for the NPT curve, because this enables the reduction of the total volume of the system.

the potential itself. In particular, the shoulder in the plot of the average forces around $r=2\sigma$ is much easier to see than the subtle variations of the slop of the pmf plot in Fig. 3. A more detailed study reveals that the structure in the average forces around $r=2\sigma$ is due to the second and third rings of particles around the center particle: In a solid, there would be two peaks in $g(r)$ at $r=\sqrt{3}\sigma r^{1/6}$ and $2\sigma r^{1/6}$. In a fluid these peaks will be broadened and overlap since they are very close together. To verify that these peaks would lead to the structure seen in Fig. 4, we generated a simple model of the radial distribution function that uses weighted Gaussian functions at $r=\sqrt{3}\sigma r^{1/6}$ and $2\sigma r^{1/6}$ to model the peaks. We then calculated the corresponding pmf and force curves. The corresponding curve for the average force showed a clearly visible shoulder around $r=2\sigma$, while the potential curve displays only small variations in the slope, very similar to what we see in Figs. 3 and 4, which strongly supports our interpretation.

This demonstrates another advantage of AIM: in a lot of applications, such as in calculations on biological systems, weak signals can very often point to subtle but important effects. These are very hard to detect in poorly understood systems. The resolution and additional information provided by the average force plot in AIM could be important tools in detecting such effects. It is important to note that this information would be very difficult to obtain with other methods. Calculating accurate derivatives is very difficult with histogram-based methods.

Figure 6 compares the solvent-induced potential of mean force for the NVT and the NPT ensembles. The parameters of the NPT simulation were chosen to make the average volume close to that of the previous NVT calculation. The pressure imposed was $p^*=1.585$, giving an average volume of $\overline{V}^*_{\mathrm{NPT}}=45.2$ versus $V^*_{\mathrm{NVT}}=45.0$. We can see from the graph that the results are almost identical for values of $r>\sigma$. However, the curve for the NPT calculation is systematically lower for small values of r. This is to be expected, since the average volume of the system will be reduced when the two test particles overlap for the constant pressure ensemble. Since a smaller volume is favorable in the NPT ensemble, overlapping particles will be more favorable in this ensemble in comparison with the constant volume ensemble.

V. A Thermodynamic Cycle

In this section we apply AIM in a more general context by using it to calculate a complete thermodynamic cycle that would be very difficult to simulate with ordinary methods. This cycle will explicitly link the chemical potential of a particle with the radial distribution function of the fluid. A schematic depiction of the sequence is shown in Fig. 7. We begin with a system of N particles, and insert one additional particle with the same properties. The free energy needed for this step is the chemical potential. Next, we remove the direct interaction between the new particle and one other particle. The potential of mean force between those two particles is the solvent-induced pmf that we have already discussed. We then constrain these two particles to overlap and reduce the particle number by one. This state is equivalent to having a test particle in the system that interacts twice as strongly with the remaining particles. By reducing this interaction to the normal particle-particle interaction, we complete the cycle.

As before, the calculations will be done on relatively small systems to allow for long runs, since the main goal of the calculation is to test the methods.

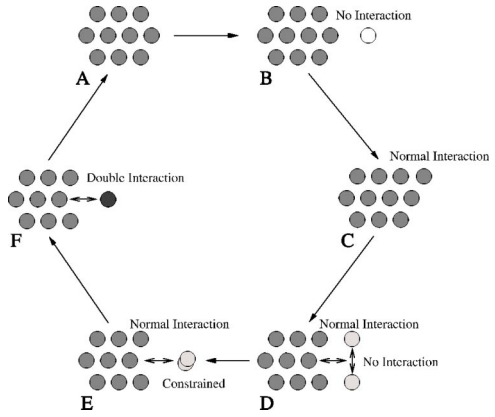

FIG. 7. Schematic view of cycle. We define the following sequence of states: A—N particle system; B—$(N+1)$ particle system that is composed of a normally interacting N particle system plus one particle that does not interact with any other particle, i.e., an ideal gas particle; C—$(N+1)$ particle system with uniform interactions between all pairs of particles; D—$(N+1)$ particle system with one pair of special particles without direct interaction with each other; E—$(N+1)$ particle system with one pair of special particles that have no direct interaction with each other and are constrained to be at the same position; F—N particle system with a single special particle whose interactions with the other particles are double the normal interaction. From state F, reducing the strength of the interactions of the special particle with the rest of the particles returns the system to state A.

There are several reasons for carrying out this calculation: First it provides a stringent test of the consistency and accuracy of the method. The errors in each step can be calculated, and the errors from different stages of the calculation can be compared, and, of course, this provides the consistency check that the total free energy difference must vanish.

The second purpose is to illustrate the advantages of AIM for the direct calculation of the chemical potential. The small bin sizes that can be used with AIM eliminate a significant source of error in such calculations. A similar calculation in the last part of the cycle shows how to use the method to calculate changes in the free energy when the strength of the interactions is varied. This is relevant to biochemical and materials science calculations that explore the effects of modifying the interactions.

We use AIM to calculate both, the potential of mean force and the chemical potential. Because this cycle returns to the original state, the net sum of the free energy differences between all individual stages must be zero. This provides a stringent test for the accuracy and efficiency of our methods. This type of test of the accuracy of the method is essential for biological applications. Very often the quantity of interest in calculations in this field is a difference between large quantities that almost cancel each other out. To get meaningful results, it is very important to have methods that give precise results.

The third purpose is to show how to obtain the chemical potential through two independent paths. This provides an additional way to compute chemical potentials. Although the first (well-known) path has the smaller statistical errors for this example, it is not completely clear that this would be the case for all problems.

The fourth purpose is to illustrate the usefulness of the solvent-induced potential of mean force discussed in Sec. IV B, which is computed without any direct interaction between particles. Because the solvent-induced pmf can be used to generate the full pmf for any direct interaction without further simulations, it allows us to perform a numerical integration to calculate the free energy involved in removing the interactions between particles. The value of the solvent-induced pmf at zero separation is also physically meaningful, since it is related to the chemical potential.

A. Sequence of states in the cycle

In this section, we will look at this cycle more rigorously and describe what quantities must be obtained from the Monte Carlo simulations to calculate the free energy differences.

Because the sequence of states is cyclic, going from A through F and back to A, the net change in free energy must be zero. Consequently, the Helmholtz free energy for the insertion of a particle (chemical potential) can be found in two ways. We can either follow path ABC, or in the opposite direction, use the path $AFEDC$.

State D is particularly interesting, because it provides a system that is well suited for adaptive sampling calculations of the effective interaction between particles. By using this thermodynamic cycle, we have a consistency condition for the free energies that can be exploited to improve calculations.

As we describe the calculations of the free energy differences between states, we will illustrate them with a numerical example for a Lennard-Jones system with $N=35+1$ particles at $T^*=0.625$ and $\rho^*=0.8$. All numerical quantities, including the free energy differences, are reported in reduced units (see, e.g., Ref. [32]).

1. From A to B

This step adds a noninteracting particle to an N particle system to give a $N+1$ particle system. For clarity in the following discussion, let us also assume that the extra particle is a different color, and make it a white particle to distinguish it from the other particles which we assume to be gray.

The partition function for a system of N particles at constant volume V and constant temperature T is given by

$$Z_N^A = \frac{1}{\Lambda^{dN}} \frac{1}{N!} \int_V d\mathbf{X}_\mathrm{N} \exp[-\beta U_N^A(\mathbf{X}_\mathrm{N})], \tag{45}$$

where Λ is the de Broglie wavelength, $\mathbf{X}_\mathrm{N}$ is the vector of all particle coordinates, $\mathbf{X}_\mathrm{N}=\{\mathbf{x}_1,\mathbf{x}_2,\ldots,\mathbf{x}_\mathrm{N}\}$, and $d\mathbf{X}_\mathrm{N}=d\mathbf{x}_1 d\mathbf{x}_2 \cdots d\mathbf{x}_\mathrm{N}$. The potential energy is

$$U_N^A(\mathbf{X}) = \sum_{j=1}^{N-1} \sum_{i>j}^{N} \phi(r_{ij}). \tag{46}$$

where $\phi(r_{ij})$ is the potential between particles separated by a distance $r_{ij}=|\mathbf{x}_i-\mathbf{x}_j|$.

The partition function for a system with N particles interacting with potentials $\phi(r_{ij})$ and one noninteracting particle is then

$$Z_{N+1}^{B} = \frac{V}{\Lambda^{d}} Z_{N}^{A}. \tag{47}$$

The change in the Helmholtz free energy between states A and B is given by

$$F_{N+1}^{B} - F_{N}^{A} = -\frac{1}{\beta} \ln\left(\frac{V}{\Lambda^{d}}\right), \tag{48}$$

$$F_{N+1}^{B} - F_{N}^{A} = \frac{1}{\beta} \ln\left(\frac{1}{V}\right) + C. \tag{49}$$

The first step of the cycle does not require simulation, since it is just the chemical potential of an ideal gas particle. For convenience we choose the constants and units such that the constant $C=0$ in Eq. (49) and find $F_{N+1}^{B}-F_{N}^{A}=-2.3792$ for our test system.

2. From B to C

To go from state B to state C, we turn on the interactions of the original particles with the new particle. This corresponds to calculating the excess chemical potential and can be handled as described in Sec. III B by changing λ from 0 to 1 in Eq. (40). We also have to account for the fact that at the beginning of this step we have two species of particles: N regular Lennard-Jones particles (gray) and one special particle (white). When $\lambda=1$, all pairs of particles have the same interactions, but there are still two species. We call this state C'. We must then account for the entropy change in going from this system to state C, in which all particles are of the same species (gray). Note that there is no difference in the energies or relative probabilities of configurations in states C' and C.

We use AIM to evaluate the free energy difference between B and C' by changing the size of the particle as explained in Sec. III B. The potential energy is

$$U_{N+1}^{BC',\lambda}(\mathbf{X}) = \sum_{j=2}^{N-1} \sum_{i>j}^{N} \phi_{LJ}(r_{ij}) + \sum_{j=2}^{N} \phi_{\lambda}(r_{1j}) \tag{50}$$

with

$$\phi_{\lambda}(r) = 4\epsilon\left[\left(\frac{\lambda\sigma}{r}\right)^{12} - \left(\frac{\lambda\sigma}{r}\right)^{6}\right]. \tag{51}$$

We then use AIM to calculate the change in free energy as function of λ. It follows

$$\frac{\partial}{\partial\lambda} F = \frac{\partial}{\partial\lambda} F^{BC',\lambda}(N+1,V,T) = \left\langle \frac{\partial}{\partial\lambda} U_{N+1}^{BC',\lambda} \right\rangle_{\lambda}, \tag{52}$$

which is the usual equation for thermodynamic integration for particle insertion to calculate the excess chemical potential.

For the free energy difference between C' and C we have

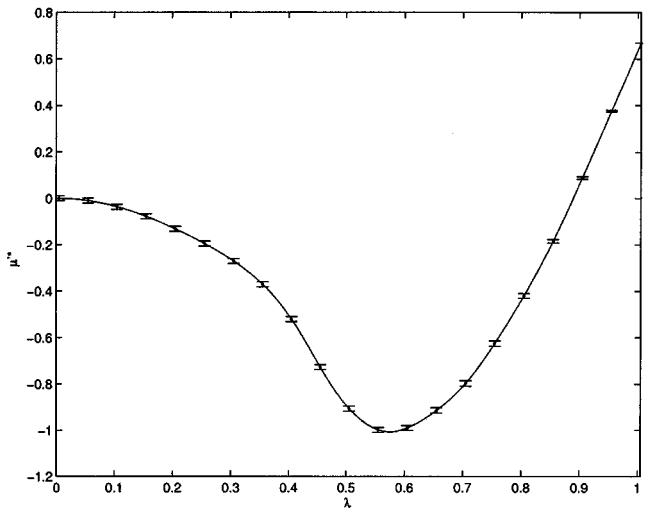

FIG. 8. Excess chemical potential μ' vs λ for adaptive integration for stage B to C. λ modifies the interaction radius. $N=36$ particle Lennard-Jones system at $T^{*}=0.625$, $V^{*}=45$, and $\rho^{*}=0.8$. The fall and rise of the curve can be explained by the configurations shown in Fig. 9 (see text).

$$Z_{N+1}^{C'} = \frac{\int_{V} d\mathbf{X}_{N+1} e^{-\beta U_{N}^{C}(\mathbf{X}_{N+1})}}{\Lambda^{d(N+1)} N!\,1!} \tag{53}$$

and

$$Z_{N+1}^{C} = \frac{\int_{V} d\mathbf{X}_{N+1} e^{-\beta U_{N}^{C}(\mathbf{X}_{N+1})}}{\Lambda^{d(N+1)} (N+1)!} \tag{54}$$

so that

$$F_{N+1}^{C} - F_{N+1}^{C'} = \frac{1}{\beta} \ln(N+1). \tag{55}$$

Figure 8 shows the free energy of the system as a function of λ between stages B and C', which scales the extra particle's interaction radius from zero to normal size. The free energy initially decreases, which is expected since the particle is small enough to fit in between other particles. As the radius increases, it will move into a low energy position near the minimum of the pairwise interactions with the surrounding particles, as shown in Fig. 9. When the interaction radius of the special particle is about half the size of that of the other particles, it can sit in between two other particles and be at the minimum of the interaction potentials with both. As the size further increases, it can simultaneously interact with three particles. As the particle radius further increases, the low energy configuration at the right of Fig. 9 is no longer possible and the energy starts to increase. Note that Fig. 9 shows a close packed configuration, which is certainly not the normal case for a fluid. However, because the density of the fluid is very high, local configurations will look roughly like this. In fact, the location of the minimum $\lambda = 0.575 \pm 0.005$ is very close to the geometric value for the

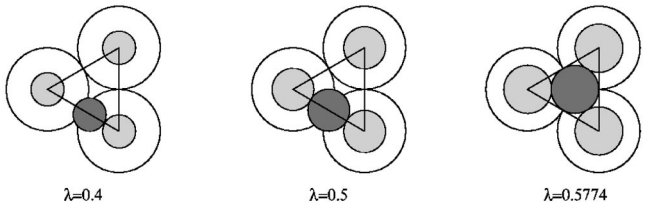

FIG. 9. Geometric interpretation of the free energy as function of interaction radius. The light gray particles are regular Lennard-Jones particles. The dark gray circle depicts the special particle. The filled circles represent the interaction radii between the solvent and special particle. The larger open circles represent the interaction radii of the Lennard-Jones interaction between the solvent particles.

close packed configuration. For the reduced free energy difference between stages B and C' we find

$$F_{N+1}^{C'} - F_{N+1}^{B} = 0.6687 \pm 0.0035 \tag{56}$$

for our test system. Between stages C' and C we have $F_{N+1}^{C} - F_{N+1}^{C'} = 2.2397$. Consequently, the total change in the Helmholtz free energy for this stage is then $F_{N+1}^{C} - F_{N+1}^{B} = 2.9084 \pm 0.0035$.

3. From C to D

To go to state D, we select two particles and remove the interactions between them. We again break this stage up into two parts: from C to C'' and from C'' to D. In state C'', all the particles interact with a regular Lennard-Jones potential, but two special particles are given a different identity, which is again indicated by a different color (light gray). The corresponding partition functions are

$$Z_{N+1}^{C''} = \frac{\int_V d\mathbf{X}_{N+1} e^{-\beta U_N(\mathbf{X}_{N+1})}}{\Lambda^{d(N+1)}(N-1)\,!\,2!}, \tag{57}$$

$$Z_{N+1}^{C} = \frac{2!}{N(N+1)} Z_{N+1}^{C''}, \tag{58}$$

so that

$$F_{N+1}^{C''} - F_{N+1}^{C} = \frac{1}{\beta} \ln\left(\frac{2!}{N(N+1)}\right). \tag{59}$$

To go from stage C'' to D, we label the positions of the two special particles $\mathbf{x}_1$ and $\mathbf{x}_2$. The potential energy for state D, where particles 1 and 2 do not interact, is

$$U_{N+1}^{D} = \sum_{i=1}^{N+1} \sum_{j=i+1, j\neq 2}^{N+1} \phi(r_{i,j}) \tag{60}$$

so that for state C'' the potential energy can be written as

$$U_{N+1}^{C''} = U_{N+1}^{D} + \phi(r_{1,2}). \tag{61}$$

The partition functions are now given by

$$Z_{N+1}^{D} = \frac{\int_V d\mathbf{X} e^{-\beta U_{N+1}^{D}(\mathbf{X})}}{\Lambda^{d(N+1)}(N+1)!} \tag{62}$$

and

$$Z_{N+1}^{C''} = \frac{\int_V d\mathbf{X} e^{-\beta U_{N+1}^{D}(\mathbf{X})} e^{-\beta\phi(r_{1,2})}}{\Lambda^{d(N+1)}(N+1)!}. \tag{63}$$

This gives us

$$\frac{Z_{N+1}^{C''}}{Z_{N+1}^{D}} = \langle e^{-\beta\phi(r_{1,2})} \rangle_D \tag{64}$$

$$= \int_0^\infty dr V(r) e^{-\beta\phi_{eff,D}(r_{1,2})} e^{-\beta\phi(r_{1,2})}. \tag{65}$$

The average in the last equation is for the internal energy where the two special particles do not interact. The potential of mean force ϕ_{eff} for state D is equivalent to the solvent-induced potential of mean force discussed in Sec. IV B and IV C. The free energy change for stages C to D is therefore

$$F_{N+1}^{D} - F_{N+1}^{C''} = \frac{1}{\beta} \ln[\langle e^{-\beta\phi(r_{1,2})} \rangle_D]. \tag{66}$$

For the Helmholtz free difference between C and C'', we find $F_{N+1}^{C''} - F_{N+1}^{C} = -4.0286$.

In order to evaluate the free energy difference between C and D, we need the radial distribution function of the two noninteracting particles, which was calculated using AIM. We performed ten independent runs of 10^7 MC sweeps each with $T^* = 0.625$, $V^* = 45$, and $\rho^* = 0.8$. The resulting solvent-induced potential of mean force is the same as the one calculated before in Sec. IV C. It is shown as dotted curve with $\times$ markers in Fig. 3.

The figure shows that it is very favorable for the two particles to overlap, as expected, since other particles around the overlapping particles see essentially one particle with twice the interaction strength while the density of the remaining fluid is reduced. We can also see that the potential of mean force is still relatively large at the edge of the simulation box. Evaluating the integrals yields

$$F_{N+1}^{D} - F_{N+1}^{C''} = \frac{1}{\beta} \ln[\langle e^{-\beta\phi(r_{1,2})} \rangle_D] \tag{67}$$

$$= -1.692 \pm 0.012 \tag{68}$$

for the change in free energy between C'' and D. The total free energy difference for this stage is $F_{N+1}^{D} - F_{N+1}^{C} = -5.720 \pm 0.012$.

4. From D to E

In state D, we have an $N+1$ particle system with uniform interactions between all pairs of particles except for particles 1 and 2, which do not interact with each other but can move freely. To go to state E, we constrain the two special particles

to occupy the same position, although that position can still move freely. For stages E and D, the partition functions are given by

$$Z_{N+1}^{E} = \frac{\int_V d\mathbf{X} e^{-\beta U_{N+1}^{D}(\mathbf{X})} \delta(\mathbf{x}_2 - \mathbf{x}_1)}{\Lambda^{d(N+1)}(N+1)!} \tag{69}$$

and

$$Z_{N+1}^{D} = \frac{\int_V d\mathbf{X} e^{-\beta U_{N+1}^{D}(\mathbf{X})}}{\Lambda^{d(N+1)}(N+1)!}. \tag{70}$$

Therefore

$$\frac{Z_{N+1}^{E}}{Z_{N+1}^{D}} = \frac{\int_V d\mathbf{X} e^{-\beta U_{N+1}^{D}(\mathbf{X})} \delta(\mathbf{x}_2 - \mathbf{x}_1)}{\int d\mathbf{X} e^{-\beta U_{N+1}^{D}}} \tag{71}$$

$$= \int d\mathbf{x}_{12} P_D(r_{1,2}) \delta(\mathbf{x}_2 - \mathbf{x}_1) \tag{72}$$

$$= \int d\theta\, r dr P_D(r_{1,2}) \frac{\delta(r_{1,2}) \delta(\theta_2 - \theta_1)}{r} \tag{73}$$

$$= P_D(0), \tag{74}$$

where $P_D(0)$ is the probability density for exact overlap of the test particles. So

$$F_{N+1}^{E} - F_{N+1}^{D} = -\frac{1}{\beta} \ln[P_D(0)] \tag{75}$$

$$= -\frac{1}{\beta} \ln\left(\frac{e^{-\beta\phi_{eff}(0)}}{\int_V V(r) dr\, e^{-\beta\phi_{eff}(r)}} \right) \tag{76}$$

$$= \phi_{eff}(0) + \frac{1}{\beta} \ln\left(\int_V V(r) dr\, e^{-\beta\phi_{eff}(r)} \right), \tag{77}$$

where ϕ_{eff} is the same solvent-induced pmf used for stage C to D.

No extra simulation is needed for this step, since we again use the solvent-induced potential of mean force, ϕ_{eff}, calculated earlier, to evaluate the integrals. We find a free energy change of

$$F_{N+1}^{E} - F_{N+1}^{D} = -1.5338 \pm 0.0011 \tag{78}$$

for constraining the two particles to occupy the same position.

5. From E to F

In stage E we have the two overlapping special particles and $N-1$ regular particles. In stage F we have $N-1$ regular particles plus a single special particle (shown in black) with a doubled interaction with the rest of the particles. It follows that the partition functions are related by

$$Z_{N+1}^{E} = \frac{1}{2!} Z_{N+1}^{F} \tag{79}$$

so that

$$F_{N}^{F} - F_{N+1}^{E} = \frac{1}{\beta} \ln\left(\frac{1}{\Lambda^{d} 2!} \right). \tag{80}$$

This step keeps track of the entropy change involved when we reduce the number of particles by one, with all the interactions remaining the same. We have $F_N^F - F_{N+1}^E = -0.4332$ for the difference in the free energy.

6. From F to A

The final step from F back to A involves reducing the interaction between the test particle and the other particles from double the interaction strength to the normal interaction strength. To again account correctly for different particle species, we divide this step into two parts: from F to A' and from A' to A. In both states F and A', one particle is of a different species. Since the path of integration does not include $\lambda = 0$, we do not encounter the convergence problems discussed in Sec. III B, and we are free to use a linear factor multiplying the energy of the test particle:

$$U_N^{FA',\lambda}(\mathbf{X}) = \sum_{i<j; i,j\neq k} \phi_{LJ}(r_{ij}) + \lambda \sum_{i; i\neq k} \phi_{LJ}(r_{ki}). \tag{81}$$

The difference between the two states is then that in state F, $\lambda = 2$, while in state A', $\lambda = 1$.

Using AIM, the change in the Helmholtz free energy as λ changes from 2 to 1 is

$$\frac{\partial}{\partial\lambda} F^{FA',\lambda}(N,V,T) = -\frac{\int_V d\mathbf{X} \frac{\partial}{\partial\lambda} e^{-\beta U_N^{FA',\lambda}(\mathbf{X})}}{\beta Z_N^{FA',\lambda} \Lambda^{dN} N!} \tag{82}$$

or

$$\frac{\partial}{\partial\lambda} F = F^{FA',\lambda}(N,V,T) = \left\langle \frac{\partial}{\partial\lambda} U_N^{FA',\lambda} \right\rangle_{\lambda}. \tag{83}$$

In stage A', we have two different species with $N-1$ particles and one particle respectively. So the free energy difference between the stages is given by

$$F_N^A - F_N^{A'} = \frac{1}{\beta} \ln(N). \tag{84}$$

The change in free energy as a function of the coupling constant λ is shown in Fig. 10. We only use the part of the data between $\lambda = 1$ and $\lambda = 2$, but we have shown the singularity at $\lambda = 0$ for completeness. The result is

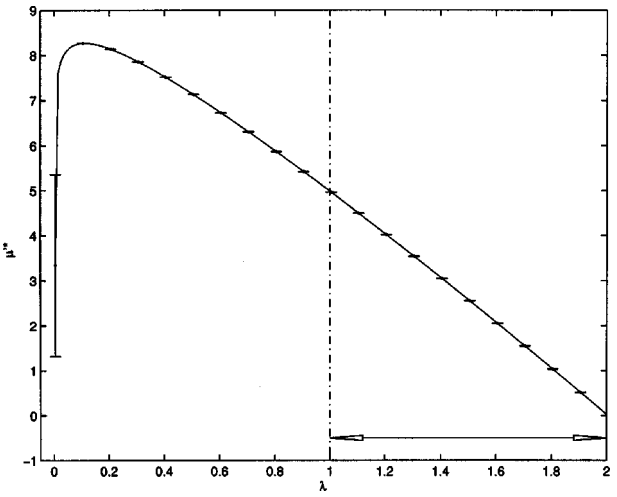

FIG. 10. μ' vs λ for adaptive integration (doubling of interaction). The data are for a $N=34+1$ particle Lennard-Jones system at $T^*=0.625$, $V^*=45$, and $\rho^*=0.778$. The part of the data, $\lambda \in [1,2]$, used to calculate $F_N^{A'}-F_N^F$ is indicated by the double headed arrow. The range of $\lambda \in [0,1]$ could be used to calculate the excess chemical potential. However, it is obvious from the plot that there is a divergence of the derivative at $\lambda=0$, as discussed in Sec. III B.

$$F_N^{A'} - F_N^F = 4.9626 \pm 0.0002. \tag{85}$$

An integration from $\lambda=0$ to $\lambda=1$ would, of course, correspond to the calculation of the excess chemical potential. However, we can see that such a calculation diverges for $\lambda \to 0$, as discussed in Sec. III B. The error bars for $\lambda \to 0$ get very large, indicating the numerical problems that would be associated with such a calculation.

For the free energy change between A' and A, we find $F_N^A - F_N^{A'} = 2.2221$. The change in the Helmholtz free energy for the full step is then given by $F_N^A - F_N^F = 7.1847 \pm 0.0002$.

7. Summary

If we add up all the contributions, we should get zero, since this is a closed thermodynamic cycle. Table I summarizes the results. Note that all free energy differences shown are in units of reduced energy. The accuracy of the method can be evaluated by looking at the sum of the individual free energy differences for the full cycle:

$$\Delta F_{cycle} = \sum_{IJ \in \{AB,\ldots,FA,AA'\}} \Delta F_{IJ}. \tag{86}$$

We can see from Table I that the total computed free energy difference, $\Delta F_{cycle}=0.027$, is very small in comparison with the values of the energy changes between the different steps in the calculation [44].

The numerical errors in Table I were calculated by performing each simulation step multiple times. The reported values are the mean values. The errors given are the errors on the average. Each of the steps is independent, except for steps *CD* and *DE*, which use the same data. The overall error calculation therefore has to take into account the correlation matrix between these two steps. The final free energy difference from zero over the full cycle is also consistent with the rather small calculated error.

TABLE I. Results for $T^*=0.625$, $V^*=45$ with $N=36$ particles.

Stage	ΔF	$\sigma(\Delta F)$
$A \to B$	−2.3792	0.0000
$B \to C'$	0.6687	0.0035
$C' \to C$	2.2397	0.0000
$C \to C''$	−4.0286	0.0000
$C'' \to D$	−1.692	0.012
$D \to E$	−1.5338	0.0011
$E \to F$	−0.4332	0.0000
$F \to A'$	4.9626	0.0002
$A' \to A$	2.2221	0.0000
Full cycle	0.027	0.013

TABLE II. Summary of results for thermodynamic cycles under different conditions.

T	N	V	ρ	ΔF	$\sigma(\Delta F)$
0.625	36	45.00	0.8	0.027	0.013
0.625	81	101.25	0.8	−0.005	0.023
0.70	81	101.25	0.8	0.013	0.024
0.45	81	112.19	0.722	−0.022	0.013
0.45	81	2700.0	0.03	0.016	0.017
0.5	81	124.61	0.65	−0.0189	0.0096

B. Other tests of the thermodynamic cycle

In addition to the calculation of a thermodynamic cycle discussed above, several additional such calculations were performed on other systems with various temperatures, densities, and numbers of particles.

Table II shows a summary of the results for these calculations. In all cases, the results gave very good accuracy and were consistent with the expected zero sum of the free energy differences around the cycle. The details of these calculations are too extensive to present in this paper, but those interested can find them in Ref. [43].

VI. CONCLUSIONS

In this paper, we have introduced the adaptive integration method and demonstrated its advantages for practical calculations. We have shown its flexibility in computing radial distribution functions and potentials of mean force, and we have shown its accuracy for thermodynamic integrations in the calculation of chemical potentials.

Our calculations of the thermodynamic cycle provide a very stringent test for AIM's accuracy and consistency. The results show that the method performs very well. In all cases examined, the total free energy differences over the cycle are very close to the theoretical value of zero. In all cases we

have looked at, including several calculations that we have not had space to present here in detail, the total free energy change over a full cycle is within one or two calculated standard deviations of zero.

These calculations also illustrate the flexibility and convenience of AIM. All computations were done with the same code. The only lines in the program that needed to be changed were those specifying the parameters used for the free energy calculations (e.g., λ or r) and the definition of the internal energy.

We believe that AIM will prove to be a useful tool for a wide variety of simulations in statistical mechanics, materials science, biochemistry, and biophysics.

ACKNOWLEDGMENTS

We thank Dr. John F. Nagle, Dr. Robert F. Sekerka, Öner Kozan and Andrew Peterson for helpful comments and discussions. This work was supported in part by a grant from NIGMS, Grant No. RO1-GM62221 (J.M.R., PI).

[1] A. Ferrenberg and R. Swendsen, Phys. Rev. Lett. **61**, 2635 (1988).
[2] A. Ferrenberg and R. Swendsen, Phys. Rev. Lett. **63**, 1195 (1989).
[3] S. Kumar, D. Bouzida, R. Swendsen, P. Kollman, and J. Rosenberg, J. Comput. Chem. **13**, 1011 (1992).
[4] S. Kumar, J. Rosenberg, D. Bouzida, R. Swendsen, and P. Kollman, J. Comput. Chem. **16**, 1339 (1995).
[5] B. Berg and T. Neuhaus, Phys. Rev. Lett. **68**, 9 (1992).
[6] M. Mezei, J. Comput. Phys. **68**, 237 (1987).
[7] G. Smith and A. Bruce, J. Phys. A **28**, 6623 (1995).
[8] O. Engkvist and G. Karlström, Chem. Phys. **213**, 63 (1996).
[9] F. Wang and D. Landau, Phys. Rev. Lett. **86**, 2050 (2001).
[10] J. Wang and R. Swendsen, J. Stat. Phys. **106**, 245 (2002).
[11] P. deOliveira, T. Penna, and H. Herrmann, Eur. Phys. J. B **1**, 205 (1998).
[12] Q. Yan and J. de Pablo, Phys. Rev. Lett. **90**, 035701 (2003).
[13] J. Barker, D. Henderson, and F. Abraham, Physica A **106**, 226 (1981).
[14] L. Rull, F. Cuadros, and J. Morales, Phys. Lett. A **123**, 217 (1987).
[15] F. Cuadros and A. Mulero, Chem. Phys. **159**, 89 (1992).
[16] A. Mulero, F. Cuadros, and M. Perezayala, Am. J. Phys. **61**, 641 (1993).
[17] F. Cuadros and A. Mulero, Chem. Phys. **177**, 53 (1993).
[18] F. Cuadros, A. Mulero, and W. Okrasinski, Physica A **214**, 162 (1995).
[19] F. Cuadros, A. Mulero, and W. Ahumada, Mol. Phys. **85**, 207 (1995).
[20] F. Cuadros, A. Mulero, and W. Okrasinski, Physica A **223**, 321 (1996).
[21] B. Smit and D. Frenkel, J. Chem. Phys. **94**, 5663 (1991).
[22] C. Bartels and M. Karplus, J. Comput. Chem. **18**, 1450 (1997).
[23] C. Reynolds, P. King, and W. Richards, Mol. Phys. **76**, 251 (1992).
[24] D. Frenkel and B. Smit, *Understanding Molecular Simulation* (Academic Press, San Diego, 2002).
[25] R. Swendsen, J. Wang, and A. Ferrenberg, Top. Appl. Phys. **71**, 75 (1992).
[26] T. Mulders, P. Kruger, W. Swegat, and J. Schlitter, J. Chem. Phys. **104**, 4869 (1996).
[27] W. denOtter and W. Briels, J. Chem. Phys. **106**, 5494 (1997).
[28] W. denOtter and W. Briels, J. Chem. Phys. **109**, 4139 (1998).
[29] W. denOtter, J. Chem. Phys. **112**, 7283 (2000).
[30] W. denOtter and W. Briels, Mol. Phys. **98**, 773 (2000).
[31] E. Darve and A. Pohorille, J. Chem. Phys. **115**, 9169 (2001).
[32] M. Allen and D. Tildesley, *Computer Simulation of Liquids* (Oxford Science, Oxford, 1987).
[33] B. Widom, J. Chem. Phys. **39**, 2808 (1963).
[34] K. Mon and R. Griffiths, Phys. Rev. A **31**, 956 (1985).
[35] A. Cross, Chem. Phys. Lett. **128**, 198 (1986).
[36] F. Wang and D. Landau, Phys. Rev. E **64**, 056101 (2001).
[37] F. Calvo, Mol. Phys. **100**, 3421 (2002).
[38] J. de Pablo, Q. Yan, and F. Escobedo, Annu. Rev. Phys. Chem. **50**, 377 (1999).
[39] N. Rathore and J. de Pablo, J. Chem. Phys. **116**, 7225 (2002).
[40] N. Rathore, T. Knotts, and J. de Pablo, J. Chem. Phys. **118**, 4285 (2003).
[41] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery, *Numerical Recipes in C* (Cambridge University Press, Cambridge, UK, 1993).
[42] P. L'Ecuyer, Commun. ACM **31**, 742 (1988).
[43] M. Fasnacht, Ph.D. thesis, Carnegie Mellon University, Pittsburgh, 2003.
[44] We note that steps involving the counting statistics due to different particle species ($C' \rightarrow C$, $C \rightarrow C''$, $E \rightarrow F$, and $A' \rightarrow A$) cancel each other out. So these steps do not contribute to the overall calculation of the energy difference of the cycle as would expect.