Very fast averaging of thermal properties of crystals by molecular simulation

Sabry G. Moustafa, Andrew J. Schultz, and David A. Kofke*
Department of Chemical and Biological Engineering, University at Buffalo, The State University of New York, Buffalo, New York 14260-4200, USA

(Received 24 January 2015; revised manuscript received 12 August 2015; published 8 October 2015)

Knowledge of approximate harmonic behavior of crystals is introduced into a new "mapped averaging" framework to yield alternative expressions for the thermodynamic properties of crystalline systems. The expressions separate the known harmonic behavior from residual averages, which thus encapsulate anharmonic contributions to the properties. With harmonic contributions removed, direct measurement of these anharmonic contributions by molecular simulation can be accomplished without contamination by noise produced by the already-known harmonic behavior. We show with application to the Lennard-Jones model that first-derivative properties (pressure, energy) can be obtained to a given precision via this harmonically mapped averaging at least 10 times faster than by using conventional averaging, and second-derivative properties (e.g., heat capacity) are obtained at least 100 times faster; in more favorable cases, the speedup exceeds a millionfold. Free-energy calculations are accelerated by 50 to 1000 times. Data obtained using these formulations are rigorous and not subject to any added approximation, and in fact are less sensitive to inaccuracies relating to finite-size effects, potential truncation, equilibration, and similar considerations. Moreover, the approach does not require any alteration in how sampling is performed during the simulation, so it may be used with standard Monte Carlo or molecular dynamics methods. However, the mapped averages do require evaluation of first and second derivatives of the intermolecular potential, for evaluation of first and second thermodynamic-derivative properties, respectively. Apart from its usefulness to simulation, the formalism developed here may constitute a basis for new theoretical treatments of crystals.

DOI: 10.1103/PhysRevE.92.043303 PACS number(s): 05.10.−a, 65.40.−b

I. INTRODUCTION

The lattice dynamics (LD) method forms the foundation for our understanding of the properties of crystalline systems [1,2]. It is based on the assumption that the intermolecular energy of the crystal is, for the given volume, at a local minimum with respect to atom positions, and that lattice vibrations are small in magnitude compared to the intermolecular spacing. This permits the use of a harmonic approximation, in which the energy is expanded to second order in the atom displacements. The result is an approximate Hamiltonian that can be solved in closed form for the dynamic and thermodynamic behaviors, using either quantum or classical mechanics. Volume-dependent thermal effects can be evaluated with the quasiharmonic extension of this treatment.

At low temperature or high pressure, LD works very well and provides a description of crystal behavior that is adequate for many purposes. However, atomic vibrations grow with increasing temperature and decreasing density, and the harmonic approximation begins to fail. Consequently, there are many conditions of interest for which LD is inadequate. In such situations, the most reliable alternative is molecular simulation.

Molecular simulation as normally practiced for the evaluation of anharmonic thermal properties does not exploit the harmonic character of the crystal to improve calculation of averages. Thus, even at conditions where LD provides an excellent description, simulation essentially "starts from scratch" in evaluating the properties, making no use of the LD characterization. Consequently, the precision of the properties computed by molecular simulation is severely compromised, inasmuch as the stochastic averaging must contend with fluctuations contributed by the harmonic component of the energy. There is a clear inefficiency in computing stochastically a large contribution that is already known analytically.

In this work, we develop and demonstrate an approach that remedies this problem. By repurposing a class of methods formulated to improve free-energy (FE) calculations, we show that the harmonic character of the crystal can be leveraged to provide in essence a direct measurement of the anharmonic contributions to the properties. The FE methods are due to Jarzynski [3], who suggested that the convergence of FE calculations can be accelerated via "targeted perturbation." Here, a coordinate mapping $\mathbf{r} \to \mathbf{r}'$ couples well-defined collective atom displacements to a FE perturbation trial, $\lambda \to \lambda'$, such that the positions $\mathbf{r}'$ are more appropriate than $\mathbf{r}$ to the perturbed state λ'. This has the effect of increasing the overlap of the sets of configurations relevant to the λ and λ' systems, respectively. Good configuration-space overlap is critical to ensuring that FE perturbation methods provide accurate and precise results: precision is enhanced when many configurations generated in λ are statistically important in the λ' state; accuracy requires that all states important to λ' are represented while sampling λ [4,5]. It is difficult in general to formulate a mapping that is effective in improving overlap, and consequently the targeted-perturbation idea has not seen wide use. However, for crystalline systems, we show that the harmonic approximation can provide guidance.

Targeted perturbation was formulated specifically to calculate FE differences for finite perturbations. We develop it here instead for differential perturbations, and rather than applying the technique toward FE calculations *per se*, we employ it to form FE derivatives. In this manner we devise new

*kofke@buffalo.edu

forms of ensemble averages for thermodynamic properties. The resulting framework provides the mechanism needed to introduce efficiencies based on the harmonic character of the crystal, which is used to formulate an appropriate mapping. Once specified, the approach is easily implemented and entails only a simple reformulation of the averages taken in an otherwise standard molecular simulation. Importantly, although the harmonic approximation enters into the development of the method, only the precision—not the accuracy—of the results is affected by the validity of this approximation. The alternative ensemble averages derived here are rigorous formulas for the associated properties, and the advantage they provide is in achieving the goal of isolating the anharmonic contributions from harmonic behavior.

In the next section, we present the general formulation of the "mapped averaging" method. We then introduce a mapping specific for crystalline systems, and with it we develop formulas for mapped averages of thermal and volumetric properties. In Sec. III, we demonstrate the effectiveness of these formulas in providing properties by molecular simulation, using the Lennard-Jones model as an example. The mapped averages yield anharmonic properties directly by simulation, and we show in this section that such quantities furthermore have appealing features relating to the robustness of the property estimates they provide; they are, for example, much less susceptible to finite-size effects than conventional averages. We conclude in Sec. IV, where we summarize the present work and offer some ideas for further development.

II. FORMALISM AND METHOD

A. Free-energy derivatives

We define a unitless FE as $\mathcal{A} \equiv \beta A$, where $\beta \equiv (k_B T)^{-1}$ with T the temperature and k_B the Boltzmann constant, and A is the FE of a d-dimensional system of N molecules in volume V; for a monatomic system, $\mathbf{r}$ is then a dN-dimensional vector. Likewise, we define $\mathcal{U}(\mathbf{r};\lambda) \equiv \beta U$, where U is the configurational energy. The FE is a function of λ, which we treat as multivariate. Jarzynski [3] showed that the FE difference $\Delta\mathcal{A} \equiv \mathcal{A}(\lambda') - \mathcal{A}(\lambda)$ between states λ' and λ can be expressed as a targeted-perturbation ensemble average in the λ system:

$$\Delta\mathcal{A} = -\ln\langle J e^{-\Delta\mathcal{U}}\rangle_\lambda, \tag{1}$$

where $\Delta\mathcal{U} \equiv \mathcal{U}(\mathbf{r}';\lambda') - \mathcal{U}(\mathbf{r};\lambda)$, and $J \equiv |\partial\mathbf{r}'/\partial\mathbf{r}|$ is the Jacobian of the mapping.

Our interest is in expressions for the thermodynamic properties of the λ system, and these are given as derivatives of the FE. Accordingly, we consider a differential perturbation in λ: $\lambda' = \lambda + d\lambda$, and expand to second order in $d\lambda$ to develop expressions for the first [3] and second derivatives of $\mathcal{A}$ with respect to elements of λ' (as indicated by subscripts):

$$\mathcal{A}_\nu = -\langle J_\nu\rangle + \langle\,\mathcal{U}_\nu\rangle, \tag{2a}$$

$$\mathcal{A}_{\mu\nu} = \langle J_\mu\rangle\langle J_\nu\rangle - \langle J_{\mu\nu}\rangle + \langle\,\mathcal{U}_{\mu\nu}\rangle + \mathrm{Cov}[\,\mathcal{U}_\mu, J_\nu] + \mathrm{Cov}[J_\mu,\mathcal{U}_\nu] - \mathrm{Cov}[\,\mathcal{U}_\mu,\mathcal{U}_\nu], \tag{2b}$$

where $\mathrm{Cov}[\mathrm{X},\mathrm{Y}] \equiv \langle\mathrm{XY}\rangle - \langle\mathrm{X}\rangle\langle\mathrm{Y}\rangle$. In the case where J is independent of $\mathbf{r}$, it may be removed from the ensemble averages, and its covariances are zero; then

$$\mathcal{A}_\nu = -J_\nu + \langle\,\mathcal{U}_\nu\rangle, \tag{3a}$$

$$\mathcal{A}_{\mu\nu} = J_\mu J_\nu - J_{\mu\nu} + \langle\,\mathcal{U}_{\mu\nu}\rangle - \mathrm{Cov}[\,\mathcal{U}_\mu,\mathcal{U}_\nu]. \tag{3b}$$

The derivatives of $\mathcal{U}$ can be separated into parts due to the mapping and parts due to any direct dependence on λ' (for fixed $\mathbf{r}$). Using ∂_ν to represent the latter,

$$\mathcal{U}_\nu = \partial_\nu\mathcal{U} - \beta\mathbf{F}\cdot\mathbf{r}_\nu, \tag{4a}$$

$$\mathcal{U}_{\mu\nu} = \partial_{\mu\nu}\mathcal{U} - \beta\mathbf{F}\cdot\mathbf{r}_{\mu\nu} + \beta\mathbf{r}_\mu\cdot\Phi\cdot\mathbf{r}_\nu - [\partial_\mu(\beta\mathbf{F})\cdot\mathbf{r}_\nu + \partial_\nu(\beta\mathbf{F})\cdot\mathbf{r}_\mu], \tag{4b}$$

where $\beta\mathbf{F} \equiv -\nabla_\mathbf{r}\mathcal{U}$ is the force vector and $\beta\Phi \equiv \nabla_\mathbf{r}\nabla_\mathbf{r}\mathcal{U}$ is the force-constant matrix. All derivatives in Eqs. (2)–(4) are evaluated at $\mathbf{r}' = \mathbf{r}$ and $\lambda' = \lambda$.

B. Temperature and volume mapping

We consider now the specific case in which the FE derivatives are with respect to (reciprocal) temperature and volume, i.e., $\lambda \equiv (\beta, V)$. We then arrive at expressions for the energy and pressure (via first derivatives), heat capacity, bulk modulus, and thermal pressure coefficient (via second derivatives), all based on a mapping informed by the harmonic approximation.

Before proceeding, we give special consideration to the treatment of the volume-change mapping. The conventional way to evaluate derivatives of the FE with volume in fact already uses a mapped average: when performing a volume change, molecule positions are scaled homogeneously with the size of the box, thereby preventing any molecules from landing outside the box when the volume is decreased, or leaving empty gaps when the volume is increased. It has long been recognized that this formulation provides a better average than one where the positions are not mapped. Care must be taken in the handling of periodic boundary conditions (PBCs) in connection to this mapping—particularly when using nonpairwise potentials—but the necessary procedures are well understood [6,7] and correctly implemented in most simulation codes. To avoid such complications in connection to the proposed harmonic mapping, we take measures to separate the uniform scaling from it. Specifically, we express the new mapping in terms of position coordinates $\mathbf{r}$ that are, by definition, divided by the box length L (such that $V = L^d$). In this formulation, the potential energy is written as $\mathcal{U}(L\mathbf{r};\lambda)$. Hence, the conventional uniform scaling (and any PBC complications that accompany it) enters the framework for $\mathcal{A}_V$ via the term $\partial_V\mathcal{U}$ in Eq. (4a), while the proposed harmonic mapping is isolated in the $\beta\mathbf{F}\cdot\mathbf{r}_V$ contribution; a corresponding separation is made for the second derivatives.

The structural effects resulting from changes of state in a harmonic system can be expressed in terms of the deviation of each molecule from its lattice site. Accordingly, we define the mapping in terms of $\Delta\mathbf{r} \equiv \mathbf{r} - \mathbf{R}$, where $\mathbf{R}$ is the vector of lattice-site coordinates (which, in our L-scaled coordinate system, is independent of volume, so $\mathbf{R}' = \mathbf{R}$). Specifically,

TABLE I. Ensemble averages required for first- and second-derivative properties, showing both conventional and harmonically mapped formulations.

Property	Definition	Conventional average	Harmonically mapped average
Configurational energy	$U = \mathcal{A}_\beta$	$\langle U \rangle$	$\frac{d(N-1)}{2\beta} + \langle U + \frac{1}{2}\mathbf{F}\cdot\Delta\mathbf{r}\rangle$
Pressure	$P = -\dfrac{\mathcal{A}_V}{\beta}$	$\frac{\rho}{\beta} + \langle P_{\rm vir}\rangle$	$\Delta\hat{P} + \langle P_{\rm vir} + f_V \mathbf{F}\cdot\Delta\mathbf{r}\rangle$
Isochoric Heat capacity	$TC_V = -\beta\mathcal{A}_{\beta\beta}$	$\beta{\rm Var}[U]$	$\frac{d(N-1)}{2\beta} + \beta\ {\rm Var}[U_{\rm anh}] - \frac{1}{4}\langle \mathbf{F}\cdot\Delta\mathbf{r} + \Delta\mathbf{r}\cdot\Phi\cdot\Delta\mathbf{r}\rangle$
Isothermal bulk modulus	$\dfrac{B}{V} = \dfrac{\mathcal{A}_{VV}}{\beta}$	$\frac{\rho}{\beta V} - \langle \partial_V P_{\rm vir}\rangle - \beta{\rm Var}[P_{\rm vir}]$	$\frac{1}{V}\Delta\hat{B} - \beta\ {\rm Var}[P_{\rm anh}] - \langle \partial_V P_{\rm vir} + f_{VV}\mathbf{F}\cdot\Delta\mathbf{r} + 2 f_V\ \partial_V \mathbf{F}\cdot\Delta\mathbf{r} - f_V^2\ \Delta\mathbf{r}\cdot\Phi\cdot\Delta\mathbf{r}\rangle$
Isochoric thermal pressure coefficient	$T\gamma_{\rm v} = \mathcal{A}_{\beta V} - \dfrac{\mathcal{A}_V}{\beta}$	$\frac{N}{V\beta} + \beta{\rm Cov}[U, P_{\rm vir}]$	$\Delta\hat{P} + \beta\ {\rm Cov}[U_{\rm anh}, P_{\rm anh}] + \frac{1}{2}\langle f_V\ \mathbf{F}\cdot\Delta\mathbf{r} + \partial_V\mathbf{F}\cdot\Delta\mathbf{r} - f_V\ \Delta\mathbf{r}\cdot\Phi\cdot\Delta\mathbf{r}\rangle$

$U_{\rm anh} \equiv U - U_{\rm lat} + \frac{1}{2}\mathbf{F}\cdot\Delta\mathbf{r}$; $P_{\rm anh} \equiv P_{\rm vir} + f_V\ \mathbf{F}\cdot\Delta\mathbf{r} - P_{\rm lat}$; $P_{\rm vir} = -\partial_V U$; $P_{\rm lat} = -\partial_V U_{\rm lat}$; $B_{\rm lat} = -V\ \partial_V P_{\rm lat}$; $\rho \equiv N/V$;
$f_V = \frac{\beta\Delta\hat{P}-\rho}{d(N-1)}$; $f_{VV} = f_V^2 - \frac{1}{V}\frac{\beta\Delta\hat{B}-\rho}{d(N-1)}$; $\Delta\hat{P} \equiv \hat{P} - P_{\rm lat}$; $\Delta\hat{B} \equiv \hat{B} - B_{\rm lat}$; $\partial_V\mathbf{F} = \frac{1}{Vd}(\mathbf{F} - \mathbf{r}\cdot\Phi)$ (consult [7] in treating $\mathbf{r}$).

we have, for a classical monatomic crystalline system,

$$\Delta\mathbf{r}' = f(\beta', V')\Delta\mathbf{r}, \tag{5a}$$

where the function f is

$$f(\beta', V') \equiv \frac{L}{L'}\left(\frac{\beta}{\beta'}\right)^{1/2}\left(\frac{V}{V'}e^{-\beta[\delta A - \delta U_{\rm lat}]}\right)^{1/d(N-1)}. \tag{5b}$$

Here, $\delta X \equiv X(V') - X(V)$ is the change in property X resulting from this scaling, and $U_{\rm lat}(V)$ is the perfect-lattice energy; $N-1$ rather than N appears in the exponent because we work with a system having fixed center of mass (c.m.) (this is enforced in the mapping by ensuring that the c.m. of the atoms coincides with the c.m. of the lattice). The temperature dependence of this scaling is based on the $T^{1/2}$ dependence of the vibration amplitude of a classical harmonic system; this scaling was used previously to formulate an efficient temperature-perturbation scheme to measure the FE of crystals [8,9]. For the volume dependence, we estimate $\delta A \approx -\hat{P}\delta V + \hat{B}\delta V^2/2V$, where $\hat{P}$ and $\hat{B}$ are parameters selected to approximate the pressure and bulk modulus, respectively. The measured properties do not depend on those parameters, but their uncertainties do; there is always an optimum value that minimizes the uncertainty. For reasons related to the treatment of finite-size effects (see Sec. III C), we used quasiharmonic analysis to estimate $\hat{P}$ and $\hat{B}$, and we observed that this choice is not far from the optimum. A similar scaling was used recently to formulate an efficient volume-change move for isobaric simulations of crystals [10]. Both the volume and temperature scaling can be reasoned as distributing the nonlattice free-energy change of the crystal uniformly and isotropically among all the molecules.

Given this mapping, the Jacobian of transformation and the necessary energy derivatives can be easily derived. First, the Jacobian (expressible as $|\partial\Delta\mathbf{r}'/\partial\Delta\mathbf{r}|$ because $\mathbf{R}$ is constant) is

$$J(\beta', V') = f(\beta', V')^{d(N-1)} \tag{6}$$

(the full Jacobian used in Eq. (3) multiplies this by a factor of $(V'/V)^{N-1}$ to account for the uniform coordinate scaling). This is independent of $\mathbf{r}$, so we can use the form of the derivatives given in Eq. (3). Second, the energy derivatives can be obtained by plugging the first and second derivatives of $\mathbf{r}$ (viz., $\mathbf{r}_\nu = f_\nu\ \Delta\mathbf{r}$ and $\mathbf{r}_{\mu\nu} = f_{\mu\nu}\ \Delta\mathbf{r}$) into Eq. (4), yielding

$$\mathcal{U}_\nu = \partial_\nu\mathcal{U} - \beta f_\nu\ \mathbf{F}\cdot\Delta\mathbf{r}, \tag{7a}$$

$$\mathcal{U}_{\mu\nu} = \partial_{\mu\nu}\mathcal{U} - \beta f_{\mu\nu}\ \mathbf{F}\cdot\Delta\mathbf{r} + \beta f_\nu f_\mu \Delta\mathbf{r}\cdot\Phi\cdot\Delta\mathbf{r} - [f_\nu\partial_\mu(\beta\mathbf{F}) + f_\mu\partial_\nu(\beta\mathbf{F})]\cdot\Delta\mathbf{r}. \tag{7b}$$

Expressions for some thermodynamic properties derived from these formulas (using harmonic estimates of $\hat{P}$ and $\hat{B}$) along with their conventional counterparts are summarized in Table I. These include the energy U and pressure P as first-derivative properties, and the three $\{T, V\}$ second-derivative properties: the isochoric heat capacity, $C_V \equiv (\partial U/\partial T)_V$; the isothermal bulk modulus, $B \equiv -V(\partial P/\partial V)_T$; and the isochoric thermal pressure coefficient, $\gamma_{\rm v} \equiv (\partial P/\partial T)_V$.

To illustrate that these "harmonically mapped averages" give directly the anharmonic contribution, let us consider the expression for the configurational energy, U. For a harmonic system, $-1/2\,\mathbf{F}\cdot\Delta\mathbf{r}$ is, in every configuration, exactly equal to the total energy (beyond $U_{\rm lat}$), so the mapped average $(U + 1/2\,\mathbf{F}\cdot\Delta\mathbf{r})$ is identically $U_{\rm lat}$ if applied to a harmonic system. Added to this average is the analytic expression for the energy of a harmonic system ($1/2\,k_{\rm B}T$ per degree of freedom), which emerges naturally from the Jacobian derivative in Eq. (3a). The volume-perturbation mapping is less rigorous—it assumes all phonon modes scale equally with volume—so for a harmonic system we cannot expect a contribution of zero to the volume-derivative averages for every configuration. For the pressure, the primary contribution is $\hat{P}$, and added to this is an average comprising terms that tend to offset to the extent that the system is harmonic.

A noteworthy feature of Eqs. (7) is that, in the case of volume derivatives, the quantities directly relating to the harmonic mapping (i.e., those other than $\partial_V\mathcal{U}$ and $\partial_{\mu V}\mathcal{U}$) are given in terms of $\Delta\mathbf{r}$ (rather than $\mathbf{r}$) and the total forces and their derivatives. The complications from PBC [6,7] afflicting quantities such as $\mathbf{F}\cdot\mathbf{r}$—which prevent it from being evaluated as written, in terms of the total force $\mathbf{F}$ on each molecule—are

not present in $\mathbf{F} \cdot \Delta\mathbf{r}$. Moreover, the PBC-affected uniform-scaling contributions are compartmentalized (appearing as P_{vir} and $\partial_V P_{\text{vir}}$ in Table I) and unchanged by the harmonic mapping. These features are especially helpful when harmonically mapped averaging is implemented using established modeling packages (e.g., LAMMPS and VASP), which typically provide means to compute the uniform-scaling quantities with PBC, while also providing access to the total force on each atom. Incidentally, for molecular dynamics (MD) simulations, the forces are computed in the process of advancing the atom positions. Consequently, the force vector $\mathbf{F}$ is available "for free," allowing the first-derivative quantities [Eq. (7a)] to be evaluated with little added computational overhead.

C. Simulation details

We demonstrate the mapped averaging approach with application to the Lennard-Jones (LJ) model, defined by the pair potential $u(r) = 4\epsilon_{\text{LJ}}[(\sigma_{\text{LJ}}/r)^{12} - (\sigma_{\text{LJ}}/r)^6]$, where σ_{LJ} and ϵ_{LJ} are the LJ size and energy parameters, respectively, and r is the pair separation. In what follows, all quantities and results are presented in units such that σ_{LJ} and $\epsilon_{\text{LJ}}/k_{\text{B}}$ are unity ("LJ units").

We performed standard canonical-ensemble (NVT) Monte Carlo (MC) simulations [11] over thermodynamic conditions within the region of stability of the LJ fcc solid [12]. Most calculations were performed for cubic box of $N = 500$ atoms, although some simulations of larger systems ($N = 4n^3$ atoms, with $n = 6,7,\ldots,10$) were conducted to study finite-size effects. The LJ potential was truncated at $r_{\text{c}} = 3.0$, except for simulations performed to analyze the long-range correction, which considered cutoffs as large as 7.5. We collected both conventional and mapped averages as defined in Table I, with samples taken after every N MC trials (which we define as one "cycle"). The quasiharmonic estimates of the pressure and bulk modulus ($\hat{P}$ and $\hat{B}$) were obtained through polynomial fitting of the harmonic free-energy versus volume, and then analytic differentiation of the fitting function. Runs of 1×10^7 MC trials were performed (after 2×10^6 steps of equilibration) at each state point, and each run was repeated 100 times with different random-number seeds, to collect error statistics. We also conducted canonical-ensemble MD simulations to study the effect of time step on the averages. For these simulations, the temperature was controlled using an Andersen thermostat [13] applied every 100 MD steps, and samples for the averages were taken every 10 MD steps (which thus defines a cycle for the MD simulations).

Most of the simulations were conducted to follow an isochore of density $\rho = 1.0$ or an isotherm of temperature $T = 1.0$. According to [12], the melting temperature of that isochore is $T_{\text{melt}} = 0.930$ while the solid melting density of the isotherm is $\rho_{\text{melt}} = 1.011$ (both in LJ units).

It is worth emphasizing that sampling of configurations is unaffected by the use of harmonically mapped averaging. The entire difference in the proposed method is the collection of averages in the fourth column of Table I, versus the averages given in the third column. The MC and MD methods used to conduct the simulations were unchanged from conventional techniques.

All calculations were performed on 2.83-GHz Intel® Core™ 2 Quad CPU Yorkfield Q9550 processors using the ETOMICA molecular simulation package [14].

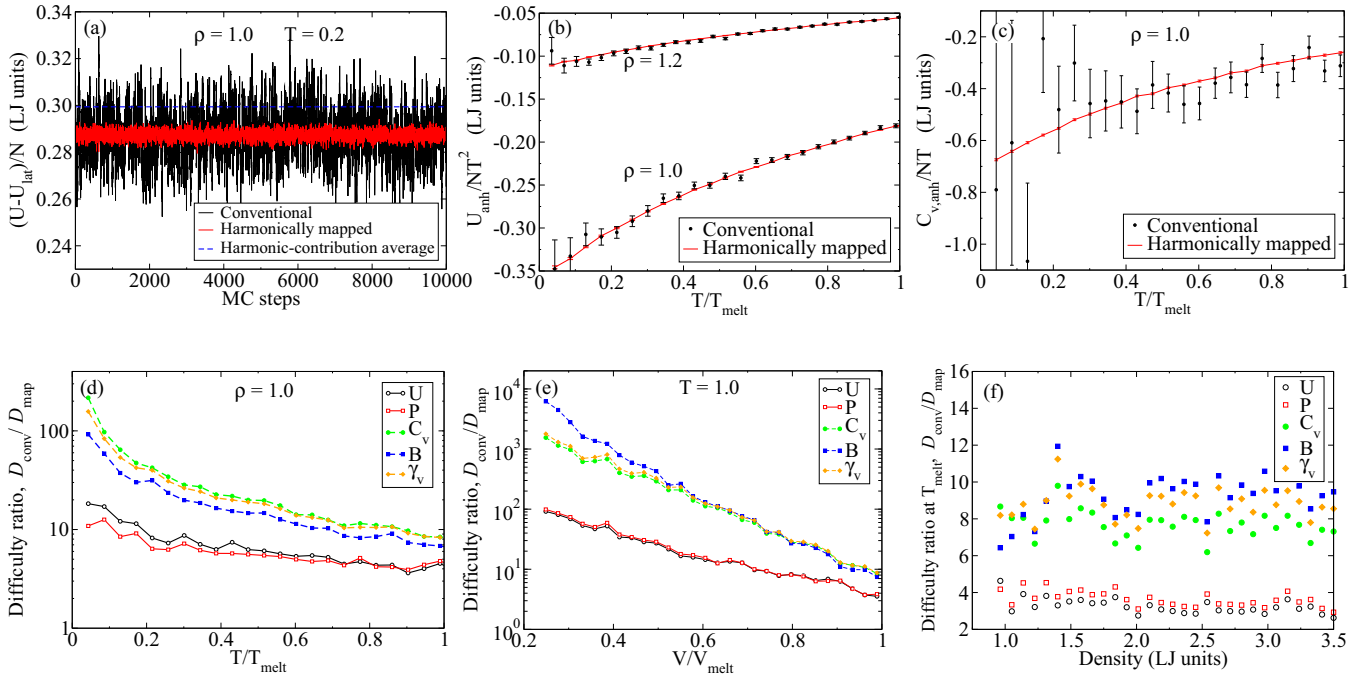

FIG. 1. (Color online) Views of the precision of harmonically mapped averaging in comparison to conventional averaging. Values plotted in (a) are for each configuration; (b) and (c) are averages over 1×10^7 MC steps. Lines in (b) and (c) simply join the points, and bars are 68% confidence intervals. Difficulty ratio in (d), (e), and (f) is ratio of $D \equiv t^{1/2}\sigma$ for conventional versus mapped averaging, where t is CPU time required to obtain a result with uncertainty σ [15]. To give a physical scale, using argonlike parameters, the pressure for states in (d) ranges from about -150 to $+150$ MPa, and for (e) about 1 TPa to 200 MPa. Data in (f) show performance along the melting line.

III. RESULTS AND DISCUSSION

A. Difficulty of measuring free-energy derivatives

In Fig. 1 we present several plots that illustrate the improved performance of mapped averaging versus conventional averaging. First, Fig. 1(a) provides a direct picture of the fluctuations in calculation of the energy per molecule. The comparison shows clearly that the conventional method suffers greatly from fluctuations in the harmonic contribution to the integrand. What is particularly notable in this case is how small the mapped-average fluctuations are relative to the difference from the harmonic-contribution average ($3/2\ k_B T$). The mapped average almost immediately provides a value of the anharmonic contribution that is statistically significant on this scale. In contrast, the conventional method would require a long period of averaging to be able to discern this difference, i.e., to provide an energy that is meaningfully different from the simple harmonic energy.

Figure 1(b) plots the average energy for two isochores, each from near-zero temperature up to the melting temperature for the given density ($T_{\rm melt} = 2.615$ at $\rho = 1.2$ [12]). The data are presented as $U_{\rm anh}/T^2$, which remains finite as $T \to 0$; this is the integrand used to compute FE differences via thermodynamic integration (TI) in temperature along an isochore (see Sec. III B). The size of the error bars and the scatter in the data clearly show that the result from the mapped average is much more precise, for the same amount of computation. As a consequence, the precision of the FE difference computed via integration of these data is markedly improved as well (see Sec. III B).

Figure 1(c) illustrates the same, but for the isochoric heat capacity C_V, a second-derivative property. The trivial $3(N-1)k_B/2$ harmonic contribution is subtracted and the data are presented as $C_{V,\rm anh}/T$, which is finite as $T \to 0$. The conventional and mapped averages are mutually consistent (as they must be, because neither involves any approximation), but the data for conventional averaging are widely scattered, far more than the mapped average. Such behavior is typical for stochastic averages of second-derivative properties. C_V is given as a variance of the energy (see Table I), which is noisy by itself as shown in Fig. 1(a). On the other hand, the mapped-average data show less scatter because $C_{V,\rm anh}$ is given as a variance of the less noisy anharmonic energy plus a canonical average that vanishes near zero temperature (see Table I).

Figures 1(d)–1(f) show the performance of the methods in terms of the "difficulty," $D \equiv t^{1/2}\sigma$, defined in terms of the central-processing unit (CPU) time t required to obtain a stochastic average with uncertainty σ. This measure accounts for any extra computational effort needed to perform the mapped average, providing a more fair comparison with the conventional method than might be given by the uncertainties alone [15]. Each plot compares performance via the ratio of difficulties of the methods ($D_{\rm conv}/D_{\rm map}$). The *square* of this quantity gives the ratio of CPU times required by conventional versus mapped averaging to achieve a result of the same precision.

Figures 1(d) and 1(e), respectively, demonstrate the temperature and volume dependence of the performance. In each case, the system becomes more harmonic toward the zero of the abscissa and reaches melting at the opposite end. Performance improvements cluster according to whether the property is a first- or second-derivative quantity. As the conditions move away from melting, the efficiency shown by harmonic mapping becomes just extraordinary: first-derivative properties are computed with 100 to 10 000 times less effort, while speedup for the second-derivative properties exceeds a millionfold for the highest density examined.

Figure 1(f) gives the difficulty ratio for each property along the melting line, starting from the triple point ($T_{\rm tp} = 0.687$ and $\rho_{\rm tp} = 0.963$) [12]. These are the conditions where the crystal deviates most from harmonic behavior, and thus where harmonically mapped averaging is least effective. Even here, in this worst-case condition, performance of the new method is 10 to 20 times faster for first-derivative properties, and about 100 times faster for second-derivative properties.

The second-derivative quantities reported above were computed for each state point via a single simulation using fluctuation-based formulas (see Table I). An alternative procedure is to compute mapped averages of the relevant first-derivative quantity over a whole range of thermodynamic states, and then fit and differentiate these data to get the second-derivative property. Comparison of these two schemes is computationally fair only if the interest is in a single second-derivative property over that range (as we assume here), not just in a single state. In addition, it is interesting to see how the efficiency of measuring first-derivative quantities (conventional and mapped) will be reflected in deriving the second derivatives from them.

We take the isochoric heat capacity C_V as an example [the bulk modulus B (not shown here) behaved similarly]. The behavior of $U_{\rm anh}/T^2$ [see Fig. 1(b)] makes it an easier function to fit, compared to fitting $U_{\rm anh}$ directly. However, the slope of $U_{\rm anh}/T^2$ does not directly yield C_V, so additional analysis is needed to extract the uncertainty in the heat capacity from the raw data. We used second- and fourth-order polynomial fits of the conventional and mapped averages, respectively, from near-zero temperature to 1.1 $T_{\rm melt}$ (to improve fitting results up to the melting point). Figure 2 shows the error in C_V/T from fitting and differentiating, versus direct ensemble averaging. It is clear that the fitting is more efficient than the direct

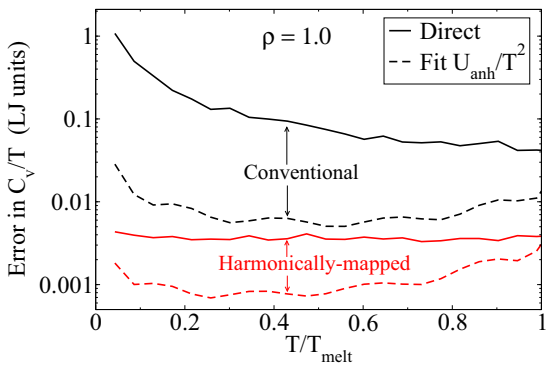

FIG. 2. (Color online) Comparison of the uncertainties from the conventional and mapped measurements of heat capacity using "direct" method (i.e., formulas in Table I) vs differentiation of a polynomial fit of $U_{\rm anh}/T^2$.

method for both conventional and mapped averages. Fitting and differentiating the conventional first-derivative average is approaching the efficiency of the direct measurement of C_V using the mapped average, although the latter remains markedly better. Fitting and differentiating the mapped first-derivative average is the best approach of all, in the present context of computing the full behavior of one second-derivative property.

To summarize: for evaluation of second-derivative properties, the choice of whether to fit first-derivative mapped averages, or perform direct measurement of the second-derivative mapped average, depends on the circumstance. If a single state point is of interest, or if multiple second-derivative properties are required, then direct averaging is preferred. If a range of state conditions is needed, and the state variable being varied is that appropriate to the necessary differentiation, then the fit-and-differentiate approach should be considered; this method is also an appealing option for cases where the configurational second derivatives [components of the Φ matrix in Eq. (7b)] are not easily determined.

B. Difficulty of measuring free energy

One of the important applications of measuring the energy is to evaluate the FE using the well-known TI technique [11]. This application is one where direct measurement of the anharmonic contributions is particularly valuable. The TI formulation for the change in configurational FE with temperature is

$$\beta A(T) = \beta_0 A(T_0) - \int_{T_0}^{T} \frac{U(\tilde{T})}{k_B \tilde{T}^2} d\tilde{T}. \tag{8}$$

When integrating from low temperature, a practice used previously by us [8,9,16] and others [17] is to separate the configurational FE into lattice-energy, harmonic, and anharmonic contributions:

$$\beta A(T) = \beta U_{\text{lat}} + \beta A_{\text{harm}}(T) + \beta A_{\text{anh}}(T), \tag{9a}$$

where

$$\beta A_{\text{harm}}(T) = \frac{1}{2} \sum_{i}^{d(N-1)} \ln\left(\frac{\hbar\omega_i}{k_B T}\right) - \frac{d}{2} \ln N + \ln \rho, \tag{9b}$$

$$\beta A_{\text{anh}}(T) = -\int_{0}^{T} \frac{U_{\text{anh}}(\tilde{T})}{k_B \tilde{T}^2} d\tilde{T}. \tag{9c}$$

In the harmonic contribution, ω_i are the eigenvalues of the Hessian matrix of the potential energy, and the last two terms are center-of-mass contributions (which may be neglected in the thermodynamic limit).

The motivation for using the separation given in Eq. (9) rather than integrating Eq. (8) directly is to remove $1/T^2$ and $1/T$ divergences of the integrand of Eq. (8)—due to lattice and harmonic contributions, respectively—as $T \to 0$, leaving the integral of U_{anh}/T^2, which is well behaved [as seen in Fig. 1(b)]. Now, added to this, we have the benefits provided by mapped averaging for evaluation of U_{anh}/T^2. The low-temperature uncertainty in this quantity is the largest source of uncertainty in the FE when evaluated via TI from low temperature; this is also where harmonically mapped averaging is most effective in improving precision. In this manner, we see that the tremendous benefit provided by harmonic mapping at low temperature impacts the calculation of the FE even at temperatures approaching melting.

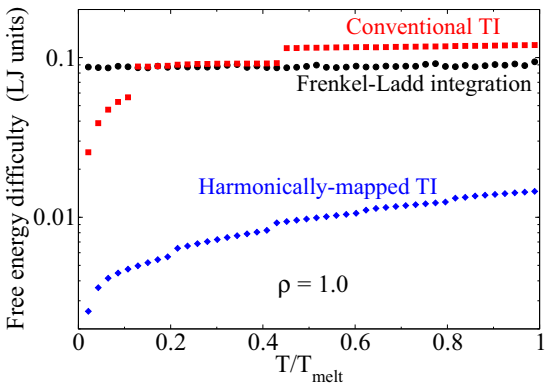

FIG. 3. (Color online) Difficulty in measuring the free energy. The small jumps in the TI results are due to changing the order of the polynomial fit (see text).

To demonstrate, we examine the stochastic uncertainty in the anharmonic FE computed according to Eq. (9c). We integrate the same functions that were fit to U_{anh}/T^2 in the previous section (which there were instead differentiated to obtain C_V). Figure 3 shows the difficulty D for measuring the FE using the TI method (both conventional and mapped), in comparison to that for the widely used Frenkel-Ladd (FL) integration scheme [18,19]. The small discontinuity in the TI uncertainty is due to changing the polynomial order with temperature. The FL difficulty is independent of T, because the integration is done over an internal coupling parameter and the process is largely unaffected by temperature. The difficulty of TI increases with temperature, which is primarily because the computational effort is cumulative with T. The computational cost of the conventional TI rises very quickly at low temperature, due to the large uncertainty in U_{anh}/T^2 at these conditions, and matches the difficulty of FL over most of the temperature range. The efficiency of harmonically mapped averaging is much better than both conventional TI and FL. Remembering that CPU time goes as the square of the difficulty, the performance speedup at low temperature approaches 1000-fold. In fact, mapped-averaging TI can yield the entire A versus T isochore up to the melting temperature, for about 50 times less computational effort than required by FL for a single temperature point.

Before closing this section, we should address issues relating to quantum effects. All of the foregoing FE methods have assumed classical behavior, even though we are invoking the limit of zero temperature in the calculations. If quantum effects are unimportant at the ultimate temperature of interest, then it is acceptable to use an entirely classical treatment in the TI process, as long as we are consistent and use a classical harmonic reference [i.e., Eq. (9b)] for the low-temperature behavior. Conversely, if the temperature of interest for the calculations is one where quantum effects are significant, then they must be introduced at some point. However, even then, this need not necessarily be done during the TI process. One can perform the TI assuming classical mechanics, and then perform additional FE calculations (at the temperature

of interest) to introduce the quantum effects [20,21]. Alternatively, it is possible, if desired, to include quantum effects in $A_{\rm harm}$ and throughout the TI calculation; details of how this is accomplished with mapped averaging will be considered in future work. The preferred approach taken to these issues will depend on the larger context in which the FE calculation is performed.

C. Thermodynamic limit

Extrapolation to the thermodynamic limit is used to evaluate properties accurately at the macroscopic level. The usual way to do this entails simulating the system at progressively larger values of N and/or r_c, and extrapolating the data to the limits $1/N \to 0$ and $1/r_c \to 0$. This procedure can be computationally expensive, inasmuch as it requires at least two simulation sizes at each state point (or more, to ensure the extrapolation is linear) and necessarily involves larger systems and longer-ranged interactions than the minimum that one might prefer to use. In this section, we examine both the finite-size and long-range corrections (FSC and LRC, correcting for $N < \infty$ and $r_c < \infty$, respectively), exploiting knowledge of the harmonic behavior in that limit. Harmonically mapped averaging plays a role as it estimates, directly and precisely, the anharmonic contribution. We find that the anharmonic contribution is much less sensitive to N and r_c than are the lattice and harmonic contributions. This suggests a strategy in which the harmonic behavior is extrapolated to the thermodynamic limit separately from the anharmonic contributions. The pressure is taken as an example to demonstrate these considerations.

Figure 4 presents the system-size dependence of the pressure. The main figure shows both the pressure $P(N)$ as measured in simulation (direct), and the same value but with a correction [16] added based on extrapolation of the harmonic-system contribution:

$$P_{\rm HarmFSC}(N) = P(N) + (P_{\rm harm}(\infty) - P_{\rm harm}(N)). \qquad (10)$$

$P_{\rm harm}(\infty)$ is obtained by extrapolation of $P_{\rm harm}(N)$ versus $1/N$; this is much easier to accomplish than extrapolation of $P(N)$ itself, as it involves no molecular simulation. Thus, the finite-size dependence of the corrected pressure is due only to the anharmonic contribution. It can be noticed that the slope of the corrected pressure is much smaller than the direct one, and in fact it is almost flat. Hence, a quick (but accurate) estimate of the pressure in the $1/N \to 0$ limit can be obtained using only a single simulation of a relatively small system. Similar behavior has been observed [8,9,16] in measuring the FSC of the FE.

The effect of temperature on the size dependence of the corrected pressure is depicted in the inset of Fig. 4. The slope increases with temperature, showing that the finite-size effects of the corrected pressure grow as the system deviates from harmonicity. However, the slope of the corrected pressure at the highest temperature visited ($T = 0.8$) is still much smaller ($\approx$6 times smaller) than that of the direct one.

By way of explanation, one should consider that the effect of increasing the system size N (or equivalently, V) at constant r_c is to admit a finer spectrum of short-wavelength fluctuations while also introducing new fluctuations at larger length scales. The new short-wavelength fluctuations will not differ qualitatively from those present at smaller system sizes, and so one might expect their effect on the properties to be small. The longer-wavelength fluctuations are qualitatively different, but typically they have small amplitudes, and hence they are adequately characterized by a harmonic treatment. The consequence is that most of the finite-size effects are captured by the extrapolation of the harmonic contribution alone, as demonstrated here.

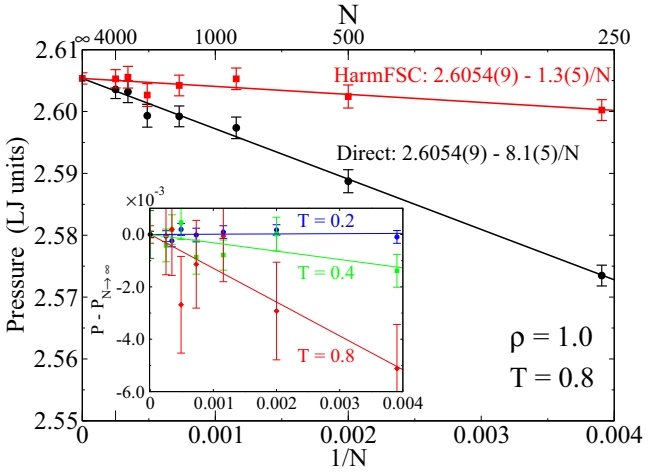

FIG. 4. (Color online) Finite-size effects of measuring the pressure directly (no correction) and using the HarmFSC correction $[P_{\rm harm}(\infty) - P_{\rm harm}(N)]$. All data were computed via harmonically mapped averaging with $r_c = 3.0$. Inset: Effect of temperature on the finite-size effects of the corrected pressure. Each line is the HarmFSC-corrected pressure for the indicated temperature, minus the value for $N \to \infty$ (i.e., the y intercept), which is subtracted to allow a common scale for all temperatures. Lines are linear fits weighted by uncertainties.

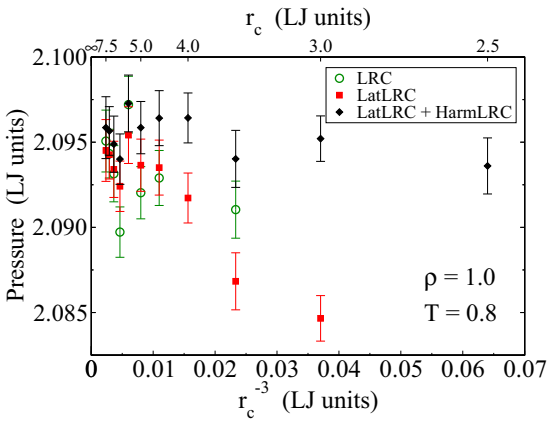

FIG. 5. (Color online) Infinite-range pressure as estimated by three choices for the correction to the finite-range simulation average: the standard LRC (assuming homogeneous medium) [11]; LatLRC $\equiv P_{\rm lat}(\infty) - P_{\rm lat}(r_c)$; and HarmLRC $\equiv P_{\rm harm}(\infty) - P_{\rm harm}(r_c)$. Pressure estimates are plotted as a function of the potential-cutoff radius used in the simulation. Some points from LRC and LatLRC with short r_c are out of the scale (as indicated by their absence where other points are plotted at a given r_c). All data were obtained via harmonically mapped averaging of a system of $N = 4000$ LJ atoms.

Assuming that the model of interest is not one with a truncated potential, full characterization of the thermodynamic limit should consider also extrapolation to an infinite cutoff. The effect of truncation on the pressure for fixed N is depicted in Fig. 5, showing different choices for the LRC (the uncorrected pressure is not presented as it is outside the range of the plot). First, we examine application of the standard LRC [11]—which assumes a homogeneous, isotropic medium beyond r_c—to the measured pressure. This treatment provides an unsatisfactory result: for the smaller r_c values (2.5 and 3.0), the corrected pressure is off the scale of the figure, and, even worse, no systematic trend is exhibited with increasing r_c, precluding a meaningful extrapolation. Second, we use a lattice correction (LatLRC), which is defined as $P_{\rm lat}(\infty) - P_{\rm lat}(r_c)$. A truncation of $r_c = 300$ is used to estimate the infinite-lattice pressure, which is converged to the eighth digit to the infinite-r_c limit. This approach greatly improves upon the standard LRC, in that it exhibits a linear behavior with $1/r_c^3$. However, its slope is still noticeable, and the $r_c = 2.5$ corrected pressure falls out of the scale of the figure. To improve this correction even more, a harmonic correction (HarmLRC), defined as $P_{\rm harm}(\infty) - P_{\rm harm}(r_c)$, is added. Thus, the cutoff dependence of this "lattice + harmonic" correction is due only to the anharmonic contribution. The figure shows that the slope is markedly reduced, and in fact it is statistically flat. Unlike the behavior of the FSC, for LRC the temperature has a negligible role on the anharmonic pressure versus r_c relation; running MC at different temperatures (including melting) finds the same statistically flat dependence (not shown). In contrast, we find that the sensitivity of the LRC and LatLRC corrections increases with increasing temperature.

We see then that using a small cutoff (as small as $r_c = 2.5$ for this LJ prototype) gives a quick, yet accurate, measure of the $1/r_c \to 0$ limit while, in contrast, a cutoff of at least $r_c = 6.5$ is required for other corrections to get statistically similar results. The cost of this larger cutoff is considerable, as a simulation with $r_c = 6.5$ requires about 17 times more effort than one using $r_c = 2.5$. This extra cost to obtain comparable accuracy adds to the cost of obtaining comparable precision (Sec. III A) for the conventional versus mapped averaging.

D. Effect of mapping on the trajectory properties

While Fig. 1(a) shows how small the standard deviation of the harmonically mapped energy is relative to the conventional one, the uncertainty of the average depends also on the correlation in the data. Figure 6 presents the autocorrelation of the quantities sampled to compute the energy via mapped and conventional averaging. The figure clearly shows the faster decay of the mapped-average ($U + \mathbf{F} \cdot \Delta\mathbf{r}$) correlation, which reaches zero (on the scale of the figure) after less than half the sampling required for the correlation of the conventional average (U) to approach zero. In practical terms, the faster decay of correlations translates into more independent samples for a given number of MC or MD steps. This is demonstrated in Fig. 7, which presents the fraction of MC cycles that yield an effectively independent sample. The measure is defined as $\mathrm{Var}/(\sigma^2 M)$, where Var is the sample variance of the measured quantity and M is the total number of MC trials (summed over independent runs) performed to yield

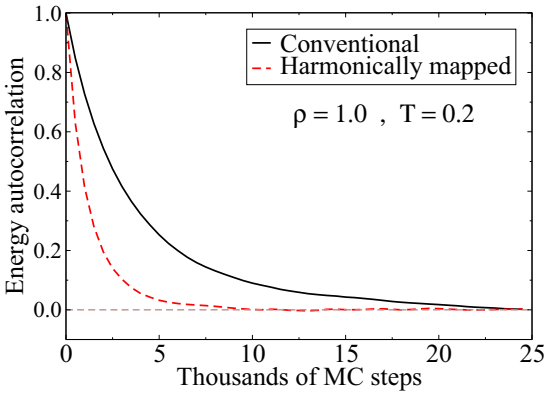

FIG. 6. (Color online) Normalized autocorrelation function of the conventional and mapped measurements of the energy. A trajectory of 2×10^5 cycles is used.

an average with uncertainty σ. Consistent with Fig. 6, the figure shows that two to four times more independent samples can be obtained using mapped averaging versus conventional averaging, even though both are taken over the same set of configurations. Similar conclusions (albeit with less effect) are made from study of correlations in MD simulation (not shown).

Study of the LRC and FSC in Sec. III C showed that the anharmonic contributions are dominated by short-wavelength fluctuations, whereas harmonic contributions are affected also by long-wavelength fluctuations. The latter will take longer to decorrelate, and consequently any quantities they contribute to will require longer periods for decay of their autocorrelation, as seen here.

A related but distinct consideration is the amount of sampling required for the system to equilibrate after starting from a nonequilibrated initial configuration. This point is examined in Fig. 8, which shows the process of equilibration for both averaging methods. The instantaneous energy difference (from the average) is normalized by the initial energy difference of each method in order to distinguish between the equilibration speed and accuracy ($U_0 - U_{\rm avg}$ is much larger for conventional averaging). As a marker for

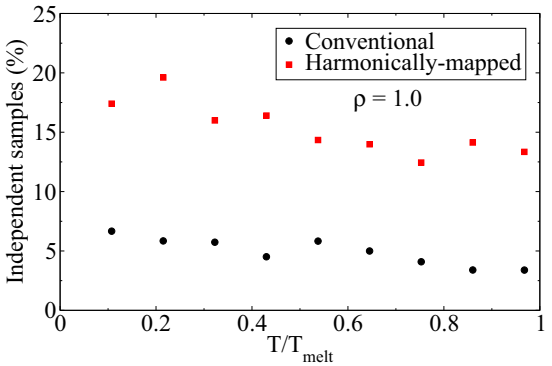

FIG. 7. (Color online) Number of independent energy samples out of 2×10^5 measured samples. The number of independent samples is estimated as Var/σ^2 with Var and σ the sample variance and uncertainty, respectively.

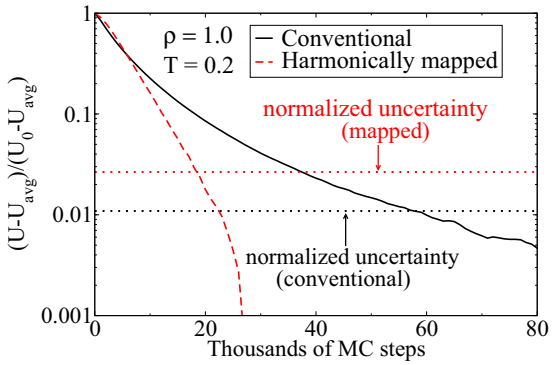

FIG. 8. (Color online) Equilibration rate of conventional and harmonically mapped averages. All data are normalized to the initial energy U_0 (off the average U_{avg}). The dotted lines are the normalized uncertainties from a production run of 1×10^4 MC steps.

equilibration, we used the point where that normalized energy difference is equal to the likewise-normalized uncertainty corresponding to a short (1×10^4 trials) production run (this uncertainty was given by scaling it up from a much longer run, where it could be determined more precisely). We observe about three times smaller equilibration time with the mapped average.

Finally, we studied the effect on harmonic mapping of inaccuracies in the sampling of configurations. Specifically, in Fig. 9 we examine conventional and mapped averages as a function of the step size used in a MD simulation, fixing the total physical time simulated. A larger step size will yield configurations that deviate more from the theoretically correct trajectory, for the given potential, initial conditions, etc. The figure shows that the mapped average is much less sensitive to these inaccuracies. Therefore, a larger time step can be used in the MD simulation without losing as much accuracy in the property measurements. This is especially important with models like those based on density functional theory (DFT), where the cost of each MD step is expensive; allowance for larger time steps permits the simulation based on mapped averaging to visit much longer dynamics than the conventional method.

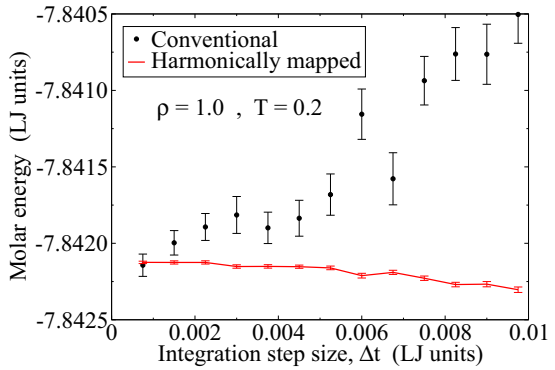

FIG. 9. (Color online) Convergence of the average energy with respect to MD time-step size. Red line simply joins the points. All points are generated using the same integration time, $t = 5000$ (LJ units).

IV. CONCLUSIONS

We have introduced and detailed the performance of harmonically mapped averaging for the calculation of properties of monatomic crystals. The method entails a reformulation of the ensemble averages taken in an otherwise standard molecular simulation. The idea has been applied here to calculation of first- and second-derivative properties (with respect to both temperature and volume), and it shows extraordinary efficiency in comparison to standard approaches. Harmonically mapped averaging yields properties to a given precision with tens to millions of times less computational effort, depending on property and state condition. We have demonstrated that the technique shows significant advantage in other respects as well: the quantities it measures are less susceptible to finite-size effects and truncation of the potential, they equilibrate and decorrelate more quickly, and they are less affected by small inaccuracies in sampling. The methodology requires no change to the sampling algorithm, and the averages it specifies are straightforward to implement. They do, however, require evaluation of first and perhaps second derivatives of the potential, and the difficulty and expense in implementing such a calculation has to be weighed against the efficiencies brought by the mapped average. First derivatives are usually available, as they are needed to compute forces for MD simulations; if second derivatives of the potential are too difficult to access, one may turn to numerical derivatives of ensemble averages to obtain second-derivative thermodynamic properties, as demonstrated in Sec. III A.

Methods developed here can be extended to handle crystalline systems in other contexts: (1) In preliminary studies we find mapped averaging to be effective for more complex models, including multibody potentials (viz., the extended Finnis-Sinclair model of tungsten, and DFT-based simulations of aluminum and of iron). (2) Extension to molecular crystals should be possible, and effective to the extent that the crystal still has harmonic character. The appropriate mapping for rotation is not obvious, but there is work to build on when attempting this [9]. (3) The treatment presented here has been entirely classical. Nuclear quantum effects can be handled using semiclassical or path-integral methods [22]. It is likely that harmonically mapped averaging can be used to some benefit when applied to such simulations. However, the temperature dependence exhibited by quantum harmonic systems differs markedly from the classical form, so the mapping given by Eq. (5) will require reformulation. (4) Harmonically mapped averages of the energy and heat capacity can be performed in NPT simulations, in the same manner as described here for NVT. There is of course no need to compute the pressure then, but the underlying concept can be applied to enhance sampling of the volume [10]. (5) In terms of properties, we have focused here on an important but not comprehensive set of solid-phase behaviors, and the method could be extended to handle others. In particular, anisotropic mappings can be formulated to improve calculation of other elastic constants. A noteworthy feature of the method is that all mapped averages can be measured simultaneously, because the enhancement is not targeted at modifying sampling.

Finally, we should point out that, apart from their usefulness to molecular simulation, the alternative formulation of the

ensemble averages given here are, in and of themselves, intrinsically interesting: they isolate anharmonic effects in a way that is novel, nontrivial, and rigorous. Indeed, they may form a basis for new understanding or new theories for crystalline behavior. Such considerations are as of now completely unexplored.

ACKNOWLEDGMENTS

Acknowledgment is made to the donors of the American Chemical Society Petroleum Research Fund for partial support of this research. Funding was also provided by an IMPACT grant from the University at Buffalo.

[1] M. Born and K. Huang, *Dynamical Theory of Crystal Lattices*, International Series of Monographs on Physics (Clarendon Press, Oxford, 1954).
[2] M. T. Dove, *Introduction to Lattice Dynamics* (Cambridge University Press, New York, 2005).
[3] C. Jarzynski, Phys. Rev. E **65**, 046122 (2002).
[4] D. A. Kofke, Mol. Phys. **102**, 405 (2004).
[5] D. A. Kofke, Fluid Phase Equilib. **228–229**, 41 (2005).
[6] M. J. Louwerse and E. J. Baerends, Chem. Phys. Lett. **421**, 138 (2006).
[7] A. P. Thompson, S. J. Plimpton, and W. Mattson, J. Chem. Phys. **131**, 154107 (2009).
[8] T. B. Tan, A. J. Schultz, and D. A. Kofke, J. Chem. Phys. **133**, 134104 (2010).
[9] T. B. Tan, A. J. Schultz, and D. A. Kofke, J. Chem. Phys. **135**, 044125 (2011).
[10] A. J. Schultz and D. A. Kofke, Phys. Rev. E **84**, 046712 (2011).
[11] D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications*, 2nd ed. (Academic Press, San Diego, CA, 2002).
[12] R. Agrawal and D. Kofke, Mol. Phys. **85**, 43 (1995).
[13] H. C. Andersen, J. Chem. Phys. **72**, 2384 (1980).
[14] A. J. Schultz and D. A. Kofke, J. Comput. Chem. **36**, 573 (2015).
[15] A. J. Schultz and D. A. Kofke, J. Chem. Theory Comput. **10**, 5229 (2014).
[16] T. B. Tan, A. J. Schultz, and D. A. Kofke, J. Chem. Phys. **132**, 214103 (2010).
[17] M. A. van der Hoef, J. Chem. Phys. **113**, 8142 (2000).
[18] D. Frenkel and A. J. C. Ladd, J. Chem. Phys. **81**, 3188 (1984).
[19] C. Vega, E. Sanz, J. L. F. Abascal, and E. G. Noya, J. Phys.: Condens. Matter **20**, 153101 (2008).
[20] J. J. Morales and K. Singer, Mol. Phys. **73**, 873 (1991).
[21] S. Habershon and D. E. Manolopoulos, J. Chem. Phys. **135**, 224111 (2011).
[22] C. P. Herrero and R. Ramírez, J. Phys.: Condens. Matter **26**, 233201 (2014).