



Alternatives to conventional ensemble averages for thermodynamic properties

Andrew J Schultz and David A Kofke

Long-established molecular theories often can provide good analytical estimates of thermodynamic properties for simple molecular models. Molecular simulation, on the other hand, can provide exact properties for the same models, or others that are arbitrarily realistic, but requires significantly more computational expense. Recent developments enable molecular simulation to improve performance by exploiting information available from approximate theories. The connection is formed via alternative ensemble averages, which allow direct calculation of the correction to the theory. In favorable cases this correction is small and can be computed with low uncertainty. Alternative ensemble averages of this type can be formulated systematically using a recently introduced framework, known as mapped averaging.

Address

Department of Chemical and Biological Engineering University at Buffalo, The State University of New York, Buffalo, NY 14260-4200, USA

Current Opinion in Chemical Engineering 2018, **23**:70–76

This review comes from a themed issue on **Frontiers of Chemical Engineering**

Edited by **Randy Snurr, Jim Pfaendtner, and Veronique Van Speybroeck**

<https://doi.org/10.1016/j.coche.2019.02.002>

0959-440X/© 2018 Elsevier Ltd. All rights reserved.

Introduction: From molecular models to thermodynamic properties

Molecular theory and molecular simulation represent two paths from statistical mechanics to properties, as illustrated in Figure 1. Examples of molecular theory include integral-equation methods (e.g., Percus–Yevick [1]), lattice dynamics for crystals [2], and the SAFT equation of state [3]. These approaches yield analytical formulas for the properties, and to remain tractable they are based on a simple representation of the intermolecular interactions. Attempts to connect such methods to realistic systems rely on either ab initio calculations or experiment to parameterize the molecular model. The compromises in the model or the theory needed to make the framework tractable causes the resulting thermodynamic model to be

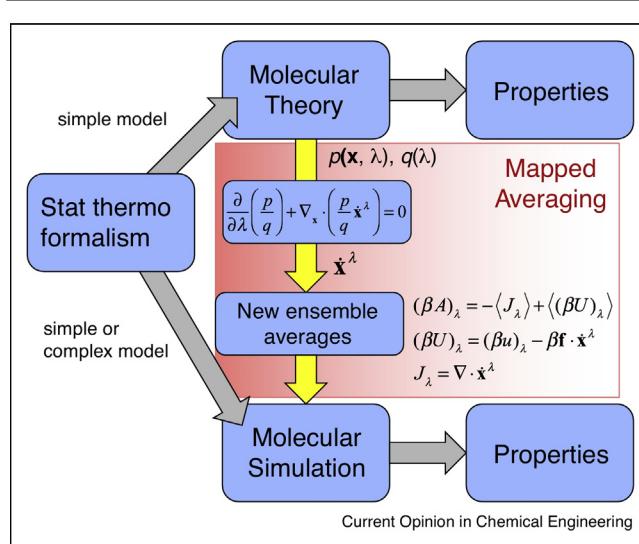
necessarily approximate. Nevertheless these approaches are extremely useful in practice.

Complementing this path is the route to properties via molecular simulation. This computational method yields properties that are in principle exact for a given molecular model. The computational (rather than analytical) nature of this approach allows it to accommodate much more complex and realistic molecular models. Although still rare, we are entering a stage where first-principles molecular models can be used with molecular dynamics or Monte Carlo sampling to obtain results that are more definitive than experiment [4], or to at least provide reliable results where experiments are not feasible (e.g., at extremes of temperature and pressure [5]).

As an aside, we note that the virial equation of state [1,6] represents a third path from statistical mechanics to properties, one exhibiting both good and bad features of the other two, as well as its own idiosyncrasies.

Historically, the primary interaction between the two paths to properties is the application of simulation to evaluate the predictions of theory, with both applied to the same molecular model. In this manner the simulation is viewed as a controlled experiment on the model, and any disagreement unambiguously points to limitations of the theory (unlike comparison to experiment, which could differ due to limits of the molecular model and/or the theory). This is indeed an important connection between theory and simulation, but apart from this there has been little synergy between them. One can formulate a theory such as lattice dynamics that provides a remarkably good—yet inexact—description of properties, and simulation generally has no way to improve its performance by making use of this knowledge.

New developments may be poised to remedy this situation. A given property can be expressed as an ensemble average in multiple ways, and molecular theory can be applied to yield a form that is advantageous for calculation by molecular simulation. Such alternatives have emerged only within the past five years or so, and they have yet to gain much attention. Presently we have one successful example in application to crystals [7,8*,9,10*,11], and others that have been put forth but not broadly tested or applied [12**,13]. A general framework to methodically develop new ensemble averages informed by theory has been proposed [12**], but its application is, in general, not trivial. On the one hand, its complexities may limit the

Figure 1

Parallel routes from statistical thermodynamics to properties, and the role of mapped averaging in connecting them.

ultimate effectiveness of this approach; but on the other, the framework opens up a new avenue to be explored and understood by the molecular simulation community, and perhaps used to invent new ways to improve performance. This opportunity is the focus of this Opinion.

Example: Pair and singlet density distributions

In 2013 Borgis et al. [14] presented an alternative specification of the ensemble average for the radial distribution function $g(r)$. The conventional formula is:

$$g(r) = \frac{1}{4\pi r^2 \rho} \left\langle \frac{1}{N} \sum_i^N \sum_{j < i} \delta(|\mathbf{r}_{ij}| - r) \right\rangle \quad (1)$$

where ρ is the density, N is the number of molecules, \mathbf{r}_{ij} is the separation vector for molecules i and j , and the sum is over all i, j pairs. The summand is the Dirac delta function, and its ensemble average is evaluated using histograms: values of the pair separation r are segregated into bins, with each bin giving a count of the number of times that a molecule pair is observed to have a separation represented by a particular (narrow) range of r . Averages of this type are forced to make a compromise between resolution and precision—smaller bins that yield noisy results versus larger bins that gloss over the finer features of the function.

Using the Poisson form for the delta function and an integration by parts, Borgis et al. [14] showed that the following rigorous expression can instead be used to evaluate $g(r)$:

$$g(r) = 1 - \frac{1}{\rho} \left\langle \frac{1}{N} \sum_i^N \sum_{j < i} \frac{H(r_{ij} - r)}{4\pi r_{ij}^2} \beta(\mathbf{f}_j - \mathbf{f}_i) \cdot \hat{\mathbf{r}}_{ij} \right\rangle \quad (2)$$

where $H(\cdot)$ denotes the Heaviside function, β is the reciprocal temperature, \mathbf{f}_j is the total force on molecule j , and the hat on \mathbf{r}_{ij} indicates a unit vector. Although it does not enter explicitly in the development, the “theory” underlying this formula is the ideal gas. This can be seen by its behavior when applied to the ideal-gas model—the forces there are identically zero so the average vanishes, leaving simply $g(r) = 1$; this is the exact result for an ideal gas, and is given with zero uncertainty. Contrast that with Eq. (1), which upon application to an ideal gas will yield $g(r) = 1$ only upon averaging, producing a distribution that exhibits some uncertainty, or noise.

The uncertainty reduction offered by Eq. (2) is present also when applied to interacting molecules. Even in this case, the baseline uniform contribution to $g(r)$ is handled exactly, whereas in Eq. (1) it must be obtained as part of the average. The benefit increases to the extent the system behaves ideally, so the approach is best suited for measuring small differences from the uniform baseline. Borgis et al. demonstrated advantages of the alternative formulation for conditions well away from ideal. One must again keep in mind that the uncertainty of the histogram-based approach, Eq. (1) can be dialed to almost any level by trading off against resolution. This issue does not afflict the force-based approach, Eq. (2), as it does not require histogramming and so it can be applied to arbitrary resolution without degrading the precision.

More recently, de las Heras and Schmidt [15•] derived a formula for the singlet density distribution of an inhomogeneous system that is similar to Eq. (2) for the pair distribution:

$$\rho(\mathbf{r}) = \rho_0 + \left\langle \sum_i^N \frac{\mathbf{r} - \mathbf{r}_i}{4\pi |\mathbf{r} - \mathbf{r}_i|^3} \cdot \beta \mathbf{f}_i \right\rangle, \quad (3)$$

where ρ_0 is a constant. The formula again yields an exact, noise-free result when applied to an ideal gas. In application to interacting molecules, the alternative expression shows improvement in the precision of the averages in comparison to the conventional histogram-based approach [15•].

Mapped averaging: A framework for deriving ensemble averages

It is possible to generalize the development summarized above to encompass other thermodynamic properties, using molecular theories other than the ideal-gas model as an implicit starting point for the alternative ensemble average. The framework for doing this begins with the idea

of targeted perturbation, proposed by Jarzynski in 2002 as a means to improve free-energy calculations [16]. The key equation presented by Jarzynski can be written as:

$$(\beta A)(\lambda') - (\beta A)(\lambda) = -\ln \langle J \exp[-(\beta U)(\mathbf{x}; \lambda') + (\beta U)(\mathbf{X}; \lambda)] \rangle_{\lambda}. \quad (4)$$

Here, $A(\lambda)$ is the free energy as a function of a system parameter λ , which might be a thermodynamic state variable, a force-field parameter, a constraint, or something else that affects the free energy (we lump β with A and U to allow for the case where λ depends on temperature); the desired free energy difference is for the perturbation $\lambda \rightarrow \lambda'$. On the right is an ensemble average over the system defined by λ . The quantity being averaged is the Boltzmann factor of the energy change accompanying the λ perturbation, but it is defined in conjunction with a mapping $\mathbf{X} \rightarrow \mathbf{x}$ of the system coordinates; J is the Jacobian for this mapping.

The aim of the mapping is to put the system in a configuration \mathbf{x} that is more representative of a configuration at λ' than \mathbf{X} is. This is illustrated in Figure 2. More specifically, one would like the relative weight of configurations in the neighborhood of \mathbf{x} at λ' to equal those around \mathbf{X} at λ ; the mapping may cause these two neighborhoods to differ in size, and that is accounted for by the Jacobian. The relative weight at \mathbf{X} may change with $\lambda \rightarrow \lambda'$ because the absolute weight at \mathbf{X} changes, and/or because the normalizing constant (i.e., the partition function) changes. Regardless, the mapping aims to move to a new configuration \mathbf{x} that preserves this relative

weight. In the original context of Jarzynski's targeted perturbation, this prescription ensures that the important phase space at λ' is adequately sampled even though the system at λ governs sampling—it is as if the two important phase spaces were fully overlapping. The result is a more precise and accurate free-energy difference via Eq. (4).

Jarzynski demonstrated the approach to evaluate the free-energy change associated with the formation of a cavity in a dense fluid [16]. Subsequently, Tan et al. [17,18] proposed "harmonically-targeted temperature perturbation" as an application to improve calculation of free energies of crystals, and Paliwal and Shirts [19] demonstrated targeted perturbation for several applications in the context of multistate free-energy calculations.

A broader range of applications of targeted perturbation can be found by considering a *differential* perturbation of λ , which then can be used to develop expressions for the first [7,16] and second [7,12^{••}] derivatives of the free energy. In this manner, targeted perturbation can be used to evaluate thermodynamic properties. Thus, for the perturbation $\lambda \rightarrow \lambda + d\lambda$ we define the mapping according to $\mathbf{X} \rightarrow \mathbf{X} + \dot{\mathbf{x}}^\lambda d\lambda$. The key quantity here is the *mapping velocity* $\dot{\mathbf{x}}^\lambda$, which defines the rate of change of the coordinates as λ is varied from its base value (which typically is the one that governs the simulation). Straightforward analysis yields expressions for the free-energy derivatives: (with μ representing a second parameter along with λ in the second derivative):

$$(\beta A)_\lambda \equiv \frac{\partial(\beta A)}{\partial \lambda} = -\langle J_\lambda \rangle + \langle (\beta U)_\lambda \rangle \quad (5a)$$

$$\begin{aligned} (\beta A)_{\lambda\mu} &\equiv \frac{\partial^2(\beta A)}{\partial \lambda \partial \mu} \\ &= -\langle J_{\lambda\mu} - J_\lambda J_\mu \rangle + \langle (\beta U)_{\lambda\mu} \rangle \\ &\quad - \text{Cov} [J_\lambda - (\beta U)_\lambda, J_\mu - (\beta U)_\mu] \end{aligned} \quad (5b)$$

The Jacobian and energy derivatives are expressed in terms of the mapping velocity [12^{••}]:

$$J_\lambda = \nabla_{\mathbf{x}} \cdot \dot{\mathbf{x}}^\lambda \quad (6a)$$

$$J_{\lambda\mu} - J_\lambda J_\mu = \nabla_{\mathbf{x}} \cdot \dot{\mathbf{x}}_\mu^\lambda + \dot{\mathbf{x}}^\mu \cdot \nabla_{\mathbf{x}} (\nabla_{\mathbf{x}} \cdot \dot{\mathbf{x}}^\lambda) \quad (6b)$$

$$(\beta U)_\lambda = (\beta u)_\lambda - \beta \mathbf{f} \cdot \dot{\mathbf{x}}^\lambda \quad (6c)$$

$$\begin{aligned} (\beta U)_{\lambda\mu} &= (\beta u)_{\lambda\mu} - \left(\dot{\mathbf{x}}_\mu^\lambda + \dot{\mathbf{x}}^\mu \cdot \nabla_{\mathbf{x}} \dot{\mathbf{x}}^\lambda \right) \cdot \beta \mathbf{f} + \dot{\mathbf{x}}^\mu \cdot \beta \mathbf{f} \cdot \dot{\mathbf{x}}^\lambda \\ &\quad - \left(\dot{\mathbf{x}}^\lambda \cdot (\beta \mathbf{f})_\mu + \dot{\mathbf{x}}^\mu \cdot (\beta \mathbf{f})_\lambda \right), \end{aligned} \quad (6d)$$

Figure 2

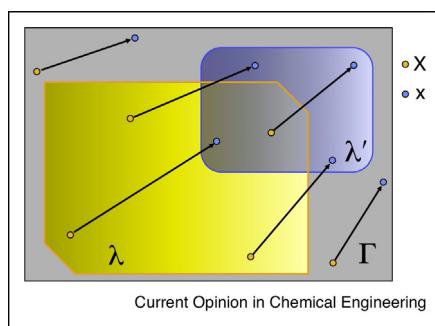


Illustration of the objective of the mapping. Large gray square represents the full $3N$ -dimensional phase space Γ , and shaded regions represent the configurations important to the system at states λ (yellow shades) and λ' (blue shades), respectively. Lines show mapping $\mathbf{X} \rightarrow \mathbf{x}$ of particular points in phase space. The aim of the procedure is to map points in the λ -important configurations into new points having the same relative weight in the set important at λ' , while configurations not important to λ map to others not important to λ' . (The mapping is bijective, so the lines could be drawn with arrows on both ends).

where $\mathbf{f} \equiv -\nabla_{\mathbf{x}}u$ is the phase-space force vector and $\phi \equiv \nabla_{\mathbf{x}}\nabla_{\mathbf{x}}u$ is the force-constant matrix (Hessian) for a given configuration. Note that we use uppercase U to denote a function that varies with λ due to its direct dependence on λ as well as the effect of the mapping (Lagrangian frame); lowercase u depends on λ only through its direct effect (if any), and does not include the effects of mapping (Eulerian frame).

These equations are rigorous, and will yield correct formulas for the free-energy derivatives—and thus alternative ensemble averages for properties—given any mapping velocity. The choice of $\dot{\mathbf{x}}^\lambda$ is important however in generating new expressions that are useful in practice, that is, that yield simulation averages having low uncertainties for a given amount of sampling. A perfect mapping is one that will leave the averaged quantity in Eq. (4) unchanged with the perturbation in λ . We have shown that such a mapping will satisfy this conservation equation for the relative weight p/q [12^{••}]:

$$\frac{\partial}{\partial \lambda} \left(\frac{p}{q} \right) + \nabla_{\mathbf{x}} \cdot \left(\frac{p}{q} \dot{\mathbf{x}}^\lambda \right) = 0 \quad (7)$$

where $p(\mathbf{x}; \lambda) = \exp(-\beta u(\mathbf{x}))$, and $q(\lambda) = \int p(\mathbf{x}; \lambda) d\mathbf{x}$. This is a high-dimensional partial differential equation, and it cannot be solved for any realistic molecular system (we cannot even evaluate $q(\lambda)$). However, we are permitted to use an approximate $p(\mathbf{x}; \lambda)$, one which allows an expression for $q(\lambda)$ and for which a solution to Eq. (7) is possible; the $\dot{\mathbf{x}}^\lambda$ so obtained can be used to generate still-rigorous ensemble averages via Eqs. (5) and (6). This is the way that approximate theory can be used to generate alternative ensemble averages that lead to molecular simulations with improved performance. This route is indicated in Figure 1 with the yellow arrows. We refer to these alternative ensemble averages as “mapped averages.”

It is amusing to find that a very familiar formula in statistical mechanics results from a simple application of this framework, showing that we already know of an example of mapped averaging. Considering a volume derivative ($\lambda \equiv V$) and employing an ideal-gas reference, we have $p(\mathbf{x}) = 1$ and $q(V) = V^N$, and a solution of Eq. (7) is (denoting \mathbf{r}_i as the element of \mathbf{x} relating to molecule i):

$$\dot{\mathbf{r}}_i^V = \frac{\mathbf{r}_i}{3V}. \quad (8)$$

Thus, the mapping velocity corresponds to a simple homogeneous affine expansion of positions with volume. The volume derivative yields the pressure, which according to Eq. (5a), (6a), and (6c) is:

$$\frac{\beta P}{\rho} = -\frac{1}{\rho}(\beta A)_V = 1 + \frac{\beta}{3} \left\langle \frac{1}{N} \sum_i^N \mathbf{r}_i \cdot \mathbf{f}_i \right\rangle, \quad (9)$$

which is the well-established virial formula for the pressure [1,20,21].

Leveraging theory to enhance simulation

The framework outlined in the previous section opens up a vast space of possibilities to improve the performance of molecular simulations, and relatively little has been done so far to explore these opportunities. For example, recently we have shown [22] that the alternative averages for the density functions, Eqs. (2) and (3), can be obtained from the mapped-averaging framework by using $p(\mathbf{x}) \equiv 1$ —in these two cases the relevant free-energy derivative is a functional derivative with respect to the pair or singlet potential, respectively. In fact, one drawback in these formulas is that they do not handle well the parts of the density function that are practically zero, such as inside the core of the pair potential. They correctly give a zero density, but only on average, and hence they are noisy where conventional histogram methods give the zero density with no uncertainty. This performance is a consequence of the use of an ideal-gas implicit reference, which assumes uniform density everywhere. We can improve this by instead selecting $p(\mathbf{x})$ to account for interactions pair-by-pair. The mapped-averaging framework then yields the following new formula for the pair distribution:

$$g(r) = e^{-\beta u_2(r)} \left[1 - \frac{1}{\rho} \left\langle \frac{1}{N} \sum_i^N \sum_{j < i} \frac{H(r_{ij} - r)}{4\pi r_{ij}^2} e^{+\beta u_2(r_{ij})} \beta(\mathbf{f}_j - \mathbf{f}_i)^* \cdot \hat{\mathbf{r}}_{ij} \right\rangle \right] \quad (10)$$

where $u_2(r)$ is the (spherically-symmetric) pair potential and the $*$ on the total-force difference indicates that both forces are exclusive of the direct \mathbf{f}_{ij} contribution.

In a similar manner, the virial formula for the pressure, Eq. (9), can be extended [12^{••}]:

$$\frac{\beta P}{\rho} = 1 + B_2 \rho + \frac{1}{2} \left\langle \frac{1}{N} \sum_i^N \sum_{j < i} \dot{r}_{ij}^V(r_{ij}) \beta(\mathbf{f}_j - \mathbf{f}_i)^* \cdot \hat{\mathbf{r}}_{ij} \right\rangle, \quad (11a)$$

where B_2 is the second virial coefficient [1], and the pair-distance mapping velocity is:

$$\dot{r}_{ij}^V(r_{ij}) = \int_0^{r_{ij}} \left(\frac{\tilde{r}}{r_{ij}} \right)^2 \left(e^{-\beta(u_2(\tilde{r}) - u_2(r_{ij}))} - 1 \right) d\tilde{r} \quad (11b)$$

Whereas Eq. (9) can be viewed as providing an ensemble average that corrects the ideal-gas law, we have in Eq. (11)

a formula that provides an ensemble average that corrects the second-order virial series.

Mapped averaging has given its most impressive performance in the application where it was originally conceived—crystalline systems [7–11,12[•],23,24]. For monoatomic crystals, the energy can be approximated by a simple isotropic harmonic form, corresponding to a Gaussian density about each lattice site:

$$p(\mathbf{x}, \beta) = \prod_{i=1}^N \exp(-\beta c |\Delta \mathbf{r}_i|^2) \quad (12)$$

where $\Delta \mathbf{r}_i$ is the displacement of atom i from its lattice-site position, and c is a constant (it has no effect on the result, so it may be left undetermined). Taking $\lambda \equiv \beta$, from Eq. (7) the mapping velocity is:

$$\dot{\mathbf{r}}_i^\beta = -\frac{\Delta \mathbf{r}_i}{2\beta}, \quad (13)$$

and we arrive at an expression for the thermodynamic energy:

$$\beta U_{\text{avg}} = (\beta A)_\beta = \frac{3}{2}(N-1) + \beta \left\langle U + \frac{1}{2} \sum_i^N \Delta \mathbf{r}_i \cdot \mathbf{f}_i \right\rangle \quad (14)$$

We recognize the first term on the right-hand side to be the harmonic contribution, showing that the average provides a direct measurement of U^{ah} , the anharmonic contribution to the energy. We refer to averages such as this, which are formulated in reference to the behavior of a harmonic system, as *harmonically mapped averages* (HMA). Note that we could not formulate an alternative average via the more obvious approach of directly subtracting a configuration-dependent harmonic energy:

$$\beta U_{\text{avg}} \neq \frac{3}{2}(N-1) + \beta \left\langle U - c \sum_i^N |\Delta \mathbf{r}_i|^2 \right\rangle, \quad (15)$$

because $\beta \langle c \sum_i^N |\Delta \mathbf{r}_i|^2 \rangle \neq 3(N-1)/2$ (the equality holds only if the average is taken while sampling a true harmonic system). The ability to generate mapped-averaging expressions methodically, rather than just intuitively, is important to the formulation of alternative ensemble averages that are correct and effective in practice. This is true *a fortiori* when considering properties based on higher-order free-energy derivatives, such as the heat capacity [7].

By enabling its direct calculation, U^{ah} can be computed to a given precision tens to thousands of times (depending

on temperature and density) more quickly than via conventional averaging [7,8[•],9,10[•]]. The greatest speedup is obtained at low temperature or high density, where the harmonic-reference starting point is most applicable. Even though the anharmonic contribution is smallest at these conditions, its precise measurement is important for computing the free energy [7,8[•],9]: if employing thermodynamic integration from zero temperature, the integrand requires U^{ah}/T^2 , so any noise in the simulation average is greatly amplified at $T \rightarrow 0$ while the ratio remains finite in this limit. This low-temperature part to the integral is by far the largest contributor to the uncertainty in the free energy at, say, the melting temperature when performing temperature integration.

We mention in passing that we recently presented an HMA formula for the temperature for use in microcanonical (NVE) simulations of crystals [11]:

$$T = \frac{1}{3(N-1)k_B} \left\langle \sum_{i=1}^N \left(\frac{\mathbf{p}_i^2}{2m} - \frac{1}{2} \Delta \mathbf{r}_i \cdot \mathbf{f}_i \right) \right\rangle, \quad (16)$$

where \mathbf{p}_i is the momentum of molecule i . This formulation exhibits a very small variance, but its advantage is not as much as it might be, due to strong negative correlations that reduce the uncertainty of the conventional form for the temperature [11]. Still, this representation might provide an interesting basis for a new thermostat.

We are currently engaged in a project to implement HMA formulas in some well-used molecular simulation codes (mapped averaging in all its forms is implemented routinely in Eatomica, our in-house molecular simulation code base [25]). Specifically, we have written scripts to process output from VASP [26], and have developed capabilities in LAMMPS [27] to compute crystalline properties this way. We plan to perform similar implementations for HOOMD-blue [28,29] and Cassandra [30]. We hope that these developments will make the HMA advances accessible to the community of molecular simulation practitioners. The usefulness of these methods might induce developers to implement options for computing coordinate second derivatives (Hessian matrix) in addition to first derivatives (forces), as a standard development practice.

Additional considerations

We collect here a few worthwhile observations and points that have not been mentioned or highlighted above. Mapped averaging . . .

- . . . does not provide something for nothing. It gains its advantages by input from the underlying molecular theory, and by information obtained from the configurational derivatives (forces and Hessian) that appear in

the averages. Generally we find that the improved precision of the mapped averages is worth any added expense incurred by evaluating these derivatives.

- . . . does not affect sampling of configurations. It can be used with any molecular dynamics or Monte Carlo method. The only thing changed is the quantity that is averaged. Consequently, any number of properties can be evaluated by mapped averaging in a single simulation—its use for one property does not interfere with a mapped average for any other.
- . . . is not (as of yet) suitable for transport or kinetic properties, nor related quantities such as time-correlation functions. As currently constituted, mapped averaging applies only to properties expressible as equilibrium free-energy derivatives. A route to application to transport properties might be found through the formalism of maximum caliber [31–33], but no development has yet been attempted in this direction.
- . . . can help offset deficiencies in sampling, and thereby provide enhancements in accuracy as well as precision. This is suggested by applications which have shown that HMA averages [7]: (a) are less susceptible to finite-size effects; (b) are less affected by truncation of the potential; (c) de-correlate more quickly, and thereby provide more statistically independent configurations for a given number of samples; (d) are faster to equilibrate from a non-equilibrium initial condition; (e) are less affected by small errors introduced by having a larger time step in MD simulations [11].
- . . . is not the only way to derive a given alternative ensemble average. The advantage of the mapped-averaging framework reviewed here is that it provides a methodical approach to do this, given a clear formulation for the molecular structure.
- . . . is efficient only to the extent that the theory used to approximate the molecular structure is accurate. Although the alternative averages derived from a theory do not become inaccurate where the theory fails, their effectiveness in reducing the uncertainty is diminished. We have found for example that mapped-averaging formulas based on gas-phase structure are not particularly efficient when computing properties of liquid phases [12^{••},13].

Outlook

In principle, limits on the effectiveness of mapped averaging might be remedied by invoking more sophisticated theories for the structure. However, theories that are too complex will preclude finding a practical solution of the mapping equation, Eq. (7). This challenge, coupled with the promise demonstrated by application of HMA to crystals, can spur efforts to cleverly synthesize molecular theory, the mapping equation, and molecular simulation toward more effective outcomes.

Indeed, the formulation of alternative ensemble averages represents a strategy to enhance molecular simulation

that is qualitatively distinct from other well studied families of methods, such as non-Boltzmann sampling, biased selection of configurations, or other techniques to accelerate sampling or convergence. Arguably, mapped averaging is a concept that has merits independent of its role in molecular simulation, and can be more properly viewed as occupying the foundations of statistical mechanics itself. Regardless, it is not often that one is presented with a qualitatively new direction to improve molecular simulations, so there is much room to invent, explore and be creative. It remains to be seen though whether mapped averaging will play out as an important new path to advance the field, or whether it is a curiosity with a few useful applications but limited general effectiveness.

Acknowledgment

This work is support by the U.S. National Science Foundation, grants CHE-1464581 and OAC-1739145.

References

1. Hansen J-P, McDonald I: *Theory of Simple Liquids*. 3rd Edition. London: Academic Press; 2006.
2. Dove MT: *Introduction to Lattice Dynamics*. New York: Cambridge University Press; 2005.
3. Chapman W, Gubbins K, Jackson G, Radosz M: **Saft: Equation-of-state solution model for associating fluids**. *Fluid Phase Equilib* 1989, **52**:31–38.
4. Shaul KRS, Schultz AJ, Kofke DA, Moldover MR: **Semiclassical fifth virial coefficients for improved ab initio helium-4 standards**. *Chem Phys Lett* 2012, **531**:11–17.
5. Sun T, Brodholt JP, Li Y, Vočadlo L: **Melting properties from ab initio free energy calculations: Iron at the earth's inner-core boundary**. *Phys Rev B* 2018, **98**:224301 <http://dx.doi.org/10.1103/PhysRevB.98.224301>.
6. A.J. Masters, Virial expansions, *J Phys: Condens Matter* 20 (28) (2008) 283102.
7. Moustafa SG, Schultz AJ, Kofke DA: **Very fast averaging of thermal properties of crystals by molecular simulation**. *Phys Rev E* 2015, **92**(4):043303.
8. Moustafa SG, Schultz AJ, Zurek E, Kofke DA: **Accurate and precise ab initio anharmonic free-energy calculations for metallic crystals: Application to hcp fe at high temperature and pressure**. *Phys Rev B* 2017, **96**:014117.
9. Moustafa SG, Schultz AJ, Kofke DA: **Harmonically Assisted Methods for Computing the Free Energy of Classical Crystals by Molecular Simulation: A Comparative Study**. *J Chem Theory Comput* 2017, **13**(2):825–834.
10. Schultz AJ, Kofke DA: **Comprehensive high-precision high-accuracy equation of state and coexistence properties for classical Lennard-Jones crystals and low-temperature fluid phases**. *J Chem Phys* 2018, **149**(20):204508.
11. Moustafa SG, Schultz AJ, Kofke DA: **Effects of thermostating in molecular dynamics on anharmonic properties of crystals**:

Here it is shown how harmonically mapped averaging can enable calculation of properties of a system modeled using high-level ab initio methods. A variety of techniques are brought together to make possible calculation of the free energy, with sufficient precision to support (in future work) evaluation of phase stability at conditions too extreme to be readily measured by experiment.

The free energy of the Lennard-Jones crystal is calculated using harmonically mapped averaging. The equation of state and coexistence lines are reported with very high precision and accuracy.

Application to fcc al at high pressure and temperature. *J Chem Phys* 2018, **149**(12):124109.

12. Schultz AJ, Moustafa SG, Lin W, Weinstein SJ, Kofke DA: **•• Reformulation of Ensemble Averages via Coordinate Mapping.** *J Chem Theory Comput* 2016, **12**(4):1491-1498.
This is the paper where the general framework for mapped averaging is introduced. It provides three applications to demonstrate its use and capabilities.

13. Lin WS, Schultz AJ, Kofke DA: **Electric-field mapped averaging for the dielectric constant.** *Fluid Phase Equilib* 2018, **470**:17-24.

14. Borgis D, Assaraf R, Rotenberg B, Vuilleumier R: **Computation of pair distribution functions and three-dimensional densities with a reduced variance principle.** *Mol Phys* 2013, **111**(22-23):3486-3492.

15. de las Heras D, Schmidt M: **Better than counting: Density profiles from force sampling.** *Phys Rev Lett* 2018, **120**:218001.
The authors show how quantities normally evaluated using histograms can instead be evaluated using an alternative ensemble average that does not involve binning in histograms. In Ref. [22] we show that their formula can be derived as an instance of mapped averaging.

16. Jarzynski C: **Targeted free energy perturbation.** *Phys Rev E* 2002, **65**:046122.

17. Tan TB, Schultz AJ, Kofke DA: **Efficient calculation of temperature dependence of solid-phase free energies by overlap sampling coupled with harmonically targeted perturbation.** *J Chem Phys* 2010, **133**:134104.

18. Tan TB, Schultz AJ, Kofke DA: **Efficient calculation of α - and β -nitrogen free energies and coexistence conditions via overlap sampling with targeted perturbation.** *J Chem Phys* 2011, **135**(4):044125.

19. Paliwal H, Shirts MR: **Multistate reweighting and configuration mapping together accelerate the efficiency of thermodynamic calculations as a function of molecular geometry by orders of magnitude.** *J Chem Phys* 2013, **138**(15):154108.

20. Louwense MJ, Baerends EJ: **Calculation of pressure in case of periodic boundary conditions.** *Chem Phys Lett* 2006, **421**(1-3):138-141.

21. Thompson AP, Plimpton SJ, Mattson W: **General formulation of pressure and stress tensor for arbitrary many-body interaction potentials under periodic boundary conditions.** *J Chem Phys* 2009, **131**(15):154107 doi:10.1063/1.3245303.

22. Purohit A, Schultz AJ, Kofke DA: **Force-sampling methods for density distributions as instances of mapped averaging.** *Mol Phys.* 2019 <http://dx.doi.org/10.1080/00268976.2019.1572243>.

23. Purohit A, Schultz AJ, Moustafa SG, Errington JR, Kofke DA: **Free energy and concentration of crystalline vacancies by molecular simulation.** *Mol Phys* 2018, **116**(21-22):3027-3041 <http://dx.doi.org/10.1080/00268976.2018.1481542>.

24. Cheng B, Ceriotti M: **Computing the absolute Gibbs free energy in atomistic simulations: Applications to defects in solids.** *Phys Rev B* 2018, **97**:054102 <http://dx.doi.org/10.1103/PhysRevB.97.054102>.

25. Schultz AJ, Kofke DA: **Etomica: An object-oriented framework for molecular simulation.** *J Comput Chem* 2015, **36**:573-583 <http://dx.doi.org/10.1002/jcc.23823>.

26. Kresse G, Furthmüller J: **Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set.** *Phys Rev B* 1996, **54**:11169-11186.

27. Plimpton S: **Fast parallel algorithms for short-range molecular dynamics.** *J Comp Phys* 1995, **117**:1-19.

28. Anderson JA, Lorenz CD, Travesset A: **General purpose molecular dynamics simulations fully implemented on graphics processing units.** *J Comput Phys* 2008, **227**(10):5342-5359.

29. Glaser J, Nguyen TD, Anderson JA, Lui P, Spiga F, Millan JA, Morse DC, Glotzer SC: **Strong scaling of general-purpose molecular dynamics simulations on gpus.** *Comput Phys Commun* 2015, **192**:97-107.

30. Shah JK, Marin-Rimoldi E, Mullen RG, Keene BP, Khan S, Paluch AS, Rai N, Romanuelo LL, Rosch TW, Yoo B, Maginn EJ: **Cassandra: An open source Monte Carlo package for molecular simulation.** *J Comput Chem* 2017, **38**(19):1727-1739.

31. Stock G, Ghosh K, Dill KA: **Maximum caliber: A variational approach applied to two-state dynamics.** *J Chem Phys* 2008, **128**(19):194102 <http://dx.doi.org/10.1063/1.2918345>.

32. Dixit PD, Wagoner J, Weistuch C, Pressé S, Ghosh K, Dill KA: **Perspective: Maximum caliber is a general variational principle for dynamical systems.** *J Chem Phys* 2018, **148**(1):010901 <http://dx.doi.org/10.1063/1.5012990>.

33. Pressé S, Ghosh K, Lee J, Dill KA: **Principles of maximum entropy and maximum caliber in statistical physics.** *Rev Mod Phys* 2013, **85**:1115-1141 <http://dx.doi.org/10.1103/RevModPhys.85.1115>.