Variational Approach to Enhanced Sampling and Free Energy Calculations

Omar Valsson[*] and Michele Parrinello[†]
Department of Chemistry and Applied Biosciences, ETH Zurich and Facoltà di Informatica, Instituto di Scienze Computationali, Università della Svizzera italiana, Via Giuseppe Buffi 13, CH-6900 Lugano, Switzerland
(Received 1 July 2014; published 27 August 2014)

The ability of widely used sampling methods, such as molecular dynamics or Monte Carlo simulations, to explore complex free energy landscapes is severely hampered by the presence of kinetic bottlenecks. A large number of solutions have been proposed to alleviate this problem. Many are based on the introduction of a bias potential which is a function of a small number of collective variables. However constructing such a bias is not simple. Here we introduce a functional of the bias potential and an associated variational principle. The bias that minimizes the functional relates in a simple way to the free energy surface. This variational principle can be turned into a practical, efficient, and flexible sampling method. A number of numerical examples are presented which include the determination of a three-dimensional free energy surface. We argue that, beside being numerically advantageous, our variational approach provides a convenient and novel standpoint for looking at the sampling problem.

DOI: 10.1103/PhysRevLett.113.090601 PACS numbers: 05.10.-a, 02.70.Ns, 05.70.Ln, 87.15.H-

Molecular dynamics (MD) and Monte Carlo (MC) simulations have become indispensable tools in many areas of science. However, whenever there are kinetic bottlenecks that lead to the appearance of long-lived metastable states, the computational cost of sampling the systems configuration space becomes prohibitive. This has lead to an intensive search for enhanced methods capable of lifting this severe limitation. One of the oldest such methods, which is still much in use today, is umbrella sampling [1], in which an external bias is added to the system to favor transitions between states separated by kinetic barriers and to allow them to occur on the time scale of the simulation. However building such potential is very challenging; many methods have been devised to this effect [2–9].

Here we present a new and efficient approach to this problem; we propose a variational method that allows for the construction of an effective bias potential and that leads to an accurate determination of the free energy as a function of a set of chosen collective variables (CVs).

In the following we consider a system described by the microscopic coordinates $\mathbf{R} \in \mathbb{R}^{3N}$ whose dynamics (e.g., MD or MC) at temperature T evolves according to a potential energy function $U(\mathbf{R})$, and leads to a canonical equilibrium distribution $P(\mathbf{R}) = e^{-\beta U(\mathbf{R})}/Z$ where $\beta = (k_B T)^{-1}$ is the inverse temperature and $Z = \int d\mathbf{R} e^{-\beta U(\mathbf{R})}$ is the partition function of the system. We map the high-dimensional $\mathbf{R}$ space into a much smaller and smoother d-dimensional space by introducing the set of collective variables $\mathbf{s}(\mathbf{R}) = (s_1(\mathbf{R}), s_2(\mathbf{R}), \ldots, s_d(\mathbf{R}))$ that give a coarse-grained but physically cogent description of the system. The appropriate choice of these collective variables is much discussed in the literature [10] and here we assume that their selection has been wise. The free energy surface (FES) associated to the CV set $\mathbf{s}$ is defined up to constant as

$$F(\mathbf{s}) = -(1/\beta) \log \int d\mathbf{R} \delta(\mathbf{s} - \mathbf{s}(\mathbf{R})) e^{-\beta U(\mathbf{R})}. \quad (1)$$

The corresponding equilibrium distribution is $P(\mathbf{s}) = e^{-\beta F(\mathbf{s})}/Z$ and the partition function can be rewritten as $Z = \int d\mathbf{s} e^{-\beta F(\mathbf{s})}$.

We introduce now the following functional of a bias potential $V(\mathbf{s})$:

$$\Omega[V] = \frac{1}{\beta} \log \frac{\int d\mathbf{s} e^{-\beta[F(\mathbf{s})+V(\mathbf{s})]}}{\int d\mathbf{s} e^{-\beta F(\mathbf{s})}} + \int d\mathbf{s} p(\mathbf{s}) V(\mathbf{s}), \quad (2)$$

where $p(\mathbf{s})$ is an arbitrary probability distribution that is assumed to be normalized. The second term can thus be read as the expectation value of $V(\mathbf{s})$ over the distribution $p(\mathbf{s})$. As shown in the Supplemental Material [11], this functional is convex and invariant under the addition of an arbitrary constant to $V(\mathbf{s})$, $\Omega[V+k] = \Omega[V]$.

The potential that renders $\Omega[V]$ stationary is, within an irrelevant constant,

$$V(\mathbf{s}) = -F(\mathbf{s}) - (1/\beta) \log p(\mathbf{s}) \quad (3)$$

for $p(\mathbf{s}) \neq 0$ and $V(\mathbf{s}) = \infty$ otherwise. This stationary point is also the global minimum of $\Omega[V]$ since the functional is convex. When the optimal bias potential [Eq. (3)] acts on the system the $\mathbf{s}$ values sampled will be only those for which $p(\mathbf{s}) \neq 0$ and $p(\mathbf{s})$ will be their resulting distribution. This offers the interesting possibility of selecting the region in CV space to be explored by appropriately choosing $p(\mathbf{s})$ (see below and Supplemental Material [11]). In more general terms the freedom of choosing $p(\mathbf{s})$ confers a high degree of flexibility to the method.

If the CVs are defined in a compact phase space of volume $\Omega_{\mathbf{s}}$, one possible and perhaps natural choice is to take $p(\mathbf{s}) = 1/\Omega_{\mathbf{s}}$, which leads to a uniform sampling in CV space as commonly done in other enhanced sampling approaches. In this case $V(\mathbf{s}) = -F(\mathbf{s})$ modulo a constant, which is the same relation as is obeyed in the asymptotic limit by the bias potential in standard metadynamics [5]. If the CVs are unbound then $p(\mathbf{s})$ can be employed to focus on the range of $\mathbf{s}$ to be explored. Another possibility is to use as $p(\mathbf{s})$

$$p(\mathbf{s}) = e^{-\beta' F(\mathbf{s})} / \left(\int d\mathbf{s} e^{-\beta' F(\mathbf{s})} \right), \tag{4}$$

where $\beta' = [k_{\mathrm{B}}(T + \Delta T)]^{-1}$ and $F(\mathbf{s})$ is our target free energy surface at inverse temperature β as defined in Eq. (1) above. This is the distribution sampled in well-tempered metadynamics with a bias factor $\gamma = \beta/\beta'$ [12]. With this choice the relation between the bias potential and free energy becomes identical to the one that asymptotically holds in well-tempered metadynamics [12,13], $V(\mathbf{s}) = -[1 - (1/\gamma)]F(\mathbf{s})$. We shall defer to a future publication the exploration of this intriguing possibility. Finally we note that it is our belief that an appropriate choice of $p(\mathbf{s})$ and smart usage of the variational flexibility of the bias potential can be of great help when considering difficult multidimensional CV spaces. We intend to explore this further in the future.

To make use of the variational property of $\Omega[V]$ we write the bias potential $V(\mathbf{s}; \boldsymbol{\alpha})$ as a function of the set of variational parameters $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \ldots, \alpha_K)$ and then minimize the function $\Omega(\boldsymbol{\alpha}) = \Omega[V(\boldsymbol{\alpha})]$ with respect to $\boldsymbol{\alpha}$. Of course the search for the minimum will be greatly facilitated by the convexity of the functional. From the converged potential $V(\mathbf{s}; \boldsymbol{\alpha})$ we can then estimate $F(\mathbf{s})$ directly from Eq. (3) if the assumed functional form has enough variational flexibility. Otherwise, we can always estimate the FES as a function of $\mathbf{s}$, or some other CVs, by employing the standard umbrella sampling relation

$$P(\mathbf{R}) \propto e^{\beta V(\mathbf{s}(\mathbf{R}))} P_V(\mathbf{R}), \tag{5}$$

where $P_V(\mathbf{R})$ is the distribution biased by $V(\mathbf{s}(\mathbf{R}))$ (see the Supplemental Material for further discussion on this equation [11]). The reweighting can also be performed before the potential has fully converged, or even on the fly during the optimization if the biasing potential converges quickly to a quasistationary state during the optimization process.

In order to implement the optimization procedure, we shall need to estimate the gradient $\Omega'(\boldsymbol{\alpha})$,

$$\frac{\partial \Omega(\boldsymbol{\alpha})}{\partial \alpha_i} = -\left\langle \frac{\partial V(\mathbf{s}; \boldsymbol{\alpha})}{\partial \alpha_i} \right\rangle_{V(\boldsymbol{\alpha})} + \left\langle \frac{\partial V(\mathbf{s}; \boldsymbol{\alpha})}{\partial \alpha_i} \right\rangle_p, \tag{6}$$

and the Hessian $\Omega''(\boldsymbol{\alpha})$,

$$\frac{\partial^2 \Omega(\boldsymbol{\alpha})}{\partial \alpha_j \partial \alpha_i} = \beta \operatorname{Cov}\left[\frac{\partial V(\mathbf{s}; \boldsymbol{\alpha})}{\partial \alpha_j}, \frac{\partial V(\mathbf{s}; \boldsymbol{\alpha})}{\partial \alpha_i} \right]_{V(\boldsymbol{\alpha})} - \left\langle \frac{\partial^2 V(\mathbf{s}; \boldsymbol{\alpha})}{\partial \alpha_j \partial \alpha_i} \right\rangle_{V(\boldsymbol{\alpha})} + \left\langle \frac{\partial^2 V(\mathbf{s}; \boldsymbol{\alpha})}{\partial \alpha_j \partial \alpha_i} \right\rangle_p, \tag{7}$$

where $\langle \cdots \rangle_{V(\boldsymbol{\alpha})}$ and $\operatorname{Cov}[\cdots]_{V(\boldsymbol{\alpha})}$ are the expectation value and the covariance, respectively, obtained in a biased simulation employing the potential $V(\mathbf{s}; \boldsymbol{\alpha})$, and $\langle \cdots \rangle_p$ is an expectation value in the distribution $p(\mathbf{s})$. A natural approach is to expand $V(\mathbf{s}; \boldsymbol{\alpha})$ in a linear basis set and use the coefficient of this expansion as variational parameters,

$$V(\mathbf{s}; \boldsymbol{\alpha}) = \sum_{\mathbf{k}} \alpha_{\mathbf{k}} G_{\mathbf{k}}(\mathbf{s}). \tag{8}$$

Given the fact that in general the FES is a rather smooth function of the CVs, a small number of terms in this expansion will usually suffice. This is to be contrasted with metadynamics where a large number of Gaussians are used to represent $V(\mathbf{s})$. As we shall see below this leads to great efficiency. In this case the gradient and the Hessian simplify,

$$\frac{\partial \Omega(\boldsymbol{\alpha})}{\partial \alpha_{\mathbf{i}}} = -\langle G_{\mathbf{i}}(\mathbf{s}) \rangle_{V(\boldsymbol{\alpha})} + \langle G_{\mathbf{i}}(\mathbf{s}) \rangle_p, \tag{9}$$

$$\frac{\partial^2 \Omega(\boldsymbol{\alpha})}{\partial \alpha_{\mathbf{j}} \partial \alpha_{\mathbf{i}}} = \beta \operatorname{Cov}[G_{\mathbf{j}}(\mathbf{s}), G_{\mathbf{i}}(\mathbf{s})]_{V(\boldsymbol{\alpha})}. \tag{10}$$

The gradient and Hessian terms for the constant term $\alpha_{\mathbf{0}}$ are zero for any given $p(\mathbf{s})$ so we can naturally take the constant term as zero and drop it from the linear expansion.

Since the gradients and the Hessian are computed statistically they are intrinsically noisy and one would need very long sampling times if we were to use them in conventional deterministic optimization algorithms. Thus we turn to the vast literature on stochastic optimization methods [14] and use a recent stochastic gradient descent–based algorithm [15]. In this algorithm, we consider at iteration n both the instantaneous iterate $\boldsymbol{\alpha}^{(n)}$ and the averaged iterates $\bar{\boldsymbol{\alpha}}^{(n)} = (n+1)^{-1} \sum_{k=0}^{n} \boldsymbol{\alpha}^{(k)}$. The instantaneous iterate is then updated using the recursion equation

$$\boldsymbol{\alpha}^{(n+1)} = \boldsymbol{\alpha}^{(n)} - \mu[\Omega'(\bar{\boldsymbol{\alpha}}^{(n)}) + \Omega''(\bar{\boldsymbol{\alpha}}^{(n)})[\boldsymbol{\alpha}^{(n)} - \bar{\boldsymbol{\alpha}}^{(n)}]], \tag{11}$$

where μ is a fixed step size and the gradient and Hessian are always obtained by using the averaged iterates $\bar{\boldsymbol{\alpha}}^{(n)}$, which amounts to taking a first-order Taylor expansion of the gradient $\Omega'(\boldsymbol{\alpha}^{(n)})$ around $\bar{\boldsymbol{\alpha}}^{(n)}$. As we show below, the instantaneous iterates $\boldsymbol{\alpha}^{(n)}$ fluctuate considerably while their averages $\bar{\boldsymbol{\alpha}}^{(n)}$ vary smoothly. This leads to a well-behaved biasing potential $V(\mathbf{s};\bar{\boldsymbol{\alpha}}^{(n)})$ and to a smoothly converging estimate of $F(\mathbf{s})$, either directly from Eq. (3) or through reweighting using Eq. (5). The averaging of the iterates also allows for a rather short sampling time at each iteration (~1 ps in the cases examined here). The choice of the step size is at present still a matter of trial and error; we expect it to depend on the system and on the functional form of $V(\mathbf{s};\boldsymbol{\alpha})$. We note that in many cases it may be too costly to obtain the complete Hessian $\Omega''(\bar{\boldsymbol{\alpha}}^{(n)})$ so for practical reasons one can consider only its diagonal part as done here. In our experience so far this does not seem to cause any ill effect.

We now turn to exemplifying how the new method works in practice. In the main text we consider only angular CVs, but in fact any variable can be treated in a similar way (see Supplemental Material [11]). In this case the natural choice is to take $p(\mathbf{s}) = 1/(2\pi)^d$ where d is the number of biased CVs. We expand $V(\mathbf{s})$ in a Fourier series, $V(\mathbf{s}) = \sum_{\mathbf{k}} \alpha_{\mathbf{k}} e^{i\mathbf{k}\mathbf{s}}$, and use the expansion coefficients as variational parameters (see Supplemental Material for details [11]). With the chosen constant $p(\mathbf{s})$ one always has $\langle V(\mathbf{s})\rangle_p = 0$ which fixes the zero of $V(\mathbf{s})$ during minimization and facilitates judging the convergence of the simulation. Each calculation is started with all variational parameters set to zero, that is $V(\mathbf{s},\bar{\boldsymbol{\alpha}}^{(0)}) = 0$.

It has become customary to test any new free energy method on alanine dipeptide in vacuum, and we shall adhere to this tradition. Conventionally the FES of alanine dipeptide is described in terms of the two backbone dihedral angles Φ and Ψ (see Supplemental Material [11]), but in vacuum only the Φ angle is a slow degree of freedom while Ψ can be considered as a fast degree of freedom. Therefore, by biasing only Φ one can still obtain a proper sampling of phase space. In Fig. 1 we show the evolution of the minimization process when using only the backbone dihedral angle Φ as a CV. It is seen that in this case the bias potential $V(\Phi)$ evolves in a manner resembling that of standard metadynamics, filling progressively all the different minima and smoothly converging to the reference free energy profile $F(\Phi)$ obtained with metadynamics. In the same figure we also show two randomly chosen coefficients in the expansion of $V(\Phi)$, where we observe that while their instantaneous values oscillates greatly their averages converge smoothly. The same convergent behavior is observed in the value of $\Omega[V]$ in Fig. 1(c). We have also performed a conventional calculation using both backbone dihedral angles Φ and Ψ as CVs with similar satisfactory results (see Supplemental Material [11]). Solvating the alanine dipeptide in explicit water and using the two traditional CVs also leads to gratifying results (see Supplemental Material [11]).

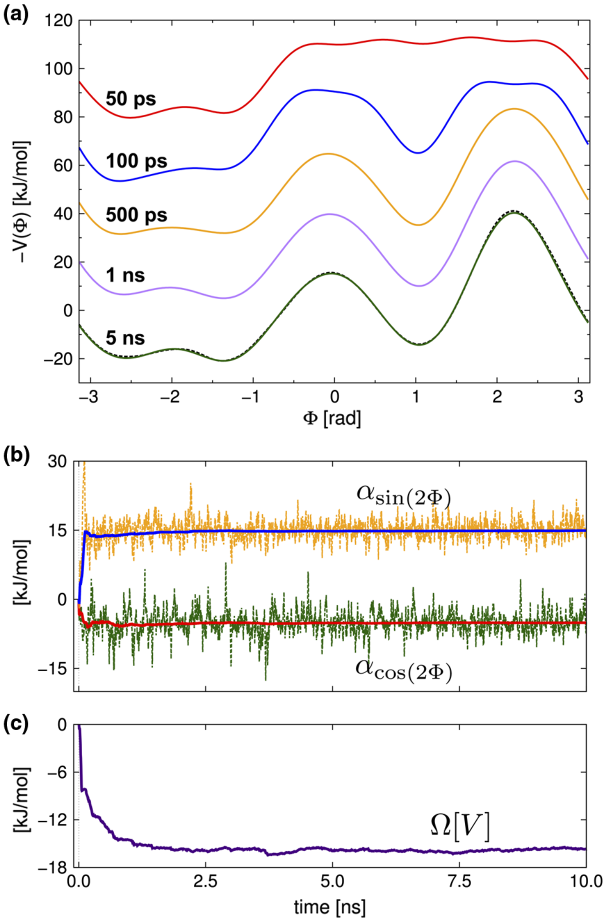

FIG. 1 (color online). Time evolution of the bias potential during minimization for alanine dipeptide in vacuum at 300 K. Only the backbone dihedral angle Φ acts as a CV, and we include 12 basis functions in the expansion of $V(\Phi)$. (a) The $F(\Phi)$ as estimated by the negative of the bias potential (for clarity, the potentials are shifted relative to one another by 25 kJ/mol). At 5 ns we also compare our variational result with a fully converged estimate of the FES from well-tempered metadynamics (black dashed line). The two curves are almost indistinguishable. (b) Time evolution of the instantaneous (dashed lines) and average (solid lines) expansion coefficients for the basis functions $\cos(2\Phi)$ and $\sin(2\Phi)$. (c) Time evolution $\Omega[V]$ estimated by a running average from the start of the simulation (see Supplemental Material [11]).

As noted earlier, in many cases the FESs are rather smooth functions so one can obtain a good representation of $V(\mathbf{s})$ with only a minimal basis set. In alanine dipeptide, both in vacuum and in water, we obtain a rather good description of the FES with only seven basis functions per CV. We make use of this ability to represent the FES with a minimal basis set in our next example.

This is the more challenging case of a Ala$_3$ peptide in vacuum. While its conformations are described by its six backbone dihedral angles Φ_1, Ψ_1, Φ_2, Ψ_2, Φ_3, Ψ_3 (see Supplemental Material [11]), only the three Φ angles suffice as CVs. As we increase the dimensionality of the CV space the number of variational parameters increases exponentially with d. To keep the number of variational parameters small we shall only use a minimal basis set. In the Ala$_3$ case this leads to use only 342 basis functions.

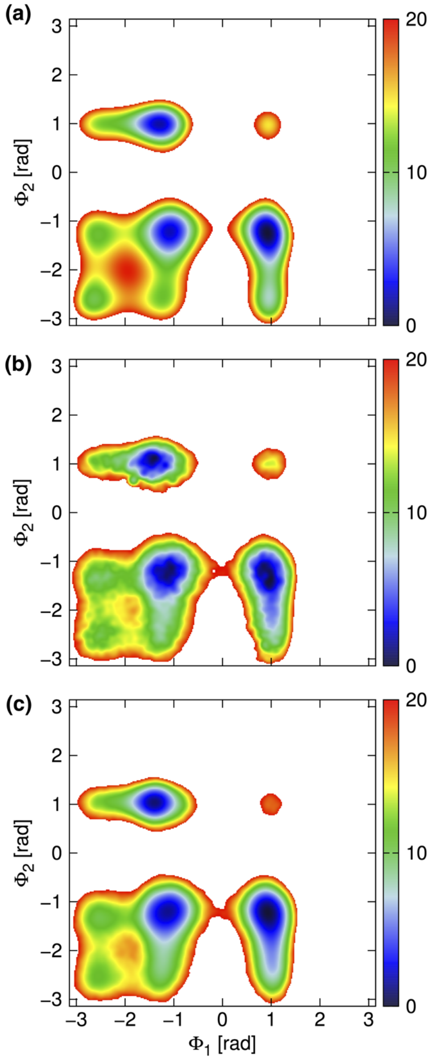

FIG. 2 (color online). The two-dimensional FESs $F(\Phi_1, \Phi_2)$ obtained with the variational approach for Ala_3 in vacuum at 300 K using the backbone dihedral angles Φ_1, Φ_2, and Φ_3 as CVs and 342 basis functions in the expansion of $V(\Phi_1, \Phi_2, \Phi_3)$. (a) FES obtained from a projection of $F(\Phi_1, \Phi_2, \Phi_3)$. (b) FES obtained with on-the-fly reweighting. (c) Reference results from a 500 ns parallel tempering simulation (eight replicas with an aggregated simulation time of 4 μs). The color scale of the FESs is given in units of kJ/mol. All FESs have their minimum value set to zero and are cut such that regions higher than 8 $k_{\mathrm{B}}T$ ($\approx$ 20 kJ/mol) are not shown. See Supplemental Material for further details [11].

With this choice during optimization all three CVs quickly become diffusive and the bias potential $V(\Phi_1, \Phi_2, \Phi_3)$ converges after 50–100 ns of simulation time.

Despite employing a minimal basis we get a rather good representation of the FES, as shown in Fig. 2 where we present a two-dimensional projection of $F(\Phi_1, \Phi_2, \Phi_3)$ on Φ_1 and Φ_2 (projections for the other CVs can be seen in the Supplemental Material [11]). Any lack of full variational flexibility in the bias potential can furthermore be fully corrected by performing an on-the-fly reweighting during the optimization process. As observed in Fig. 2 the FES obtained in this manner is in excellent agreement with reference results from an extensive 500 ns parallel-tempering simulation. In the Supplemental Material we show additional reweighted FESs for other CVs not biased during the simulation that also are in excellent agreement with the reference results [11].

While this approach offers a significant improvement over metadynamics and other similar methods its usefulness still depends on an appropriate choice of the CVs. As in metadynamics, a poor choice of the CVs will manifest itself in a hysteretic behavior during the optimization process (see Supplemental Material [11]). However our variational approach, with its potential for handling many CVs, can greatly alleviate the problem. A further help in this direction is the possibility of adding variational flexibility in the definition of the CVs.

To improve upon the method we can borrow all the ideas that have been applied to metadynamics, like parallel tempering [16], multiple walkers [17], or bias exchange [18]. Furthermore, metadynamics itself can be used to sample the averages needed in Eqs. (6) and (7) by employing a recent improved reweighting scheme [19]. The sampling power of the variational approach can thus be further enhanced by biasing the metadynamics with CVs different from $\mathbf{s}$.

The main result of this Letter is the introduction of the functional $\Omega[V]$ and the practical demonstration of its usefulness. We believe that there is ample room for improvement. The optimization procedure presented here is not necessarily optimal and different systems and CVs might require different optimization strategies and different basis sets. We plan to explore a number of alternative procedures. For instance one could think of setting up an iterative procedure in which an approximate calculation is made for $F_0(\mathbf{s})$ using Eq. (3) at the early stages of the calculation. One can then insert into Eq. (2) a new $p_0(\mathbf{s}) = e^{-\beta F_0(\mathbf{s})}/\int d\mathbf{s} e^{-\beta F_0(\mathbf{s})}$. The resulting functional is then optimized and the procedure iterated until at convergence, after k steps, $p_k(\mathbf{s}) = e^{-\beta F_k(\mathbf{s})}/\int d\mathbf{s} e^{-\beta F_k(\mathbf{s})}$ and $V_k(\mathbf{s}) \approx 0$ to the desired accuracy.

These brief discussion on the potential for improvements and modifications of the scheme is by far not exhaustive but is meant to indicate some of the future lines of investigation. In this very first application we are only using a small fraction of the potentialities of this method; much more exciting developments are to be expected. We would also like to point out the potential of our method in the development of a more rigorous coarse-graining procedure.

Finally we note that the systems considered here are by necessity simple, as conventionally done when introducing a completely new method. The strengths and limitations of our approach will become clearer as it is further developed.

The method has been implemented in a development version of the PLUMED 2 [20] plug-in and will be made publicly available in the coming future.

The authors would like to thank David Chandler for insightful discussions. All calculations were performed on the Brutus HPC cluster at ETH Zurich. We acknowledge the European Union Grant No. ERC-2009-AdG-247075 for funding.

---

[*] omar.valsson@phys.chem.ethz.ch
[†] parrinello@phys.chem.ethz.ch

[1] G. Torrie and J. Valleau, J. Comput. Phys. **23**, 187 (1977).

[2] T. Huber, A. E. Torda, and W. F. Gunsteren, J. Comput.-Aided Mol. Des. **8**, 695 (1994).

[3] E. Darve and A. Pohorille, J. Chem. Phys. **115**, 9169 (2001).

[4] F. Wang and D. P. Landau, Phys. Rev. Lett. **86**, 2050 (2001).

[5] A. Laio and M. Parrinello, Proc. Natl. Acad. Sci. U.S.A. **99**, 12562 (2002).

[6] U. Hansmann and L. Wille, Phys. Rev. Lett. **88**, 068105 (2002).

[7] L. Maragliano and E. Vanden-Eijnden, Chem. Phys. Lett. **426**, 168 (2006).

[8] C. F. Abrams and E. Vanden-Eijnden, Proc. Natl. Acad. Sci. U.S.A. **107**, 4961 (2010).

[9] P. Maragakis, A. van der Vaart, and M. Karplus, J. Phys. Chem. B **113**, 4664 (2009).

[10] A. Barducci, M. Bonomi, and M. Parrinello, WIREs: Comput. Mol. Sci. **1**, 826 (2011).

[11] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.113.090601 for further theoretical details, computational details, and additional results and free energy surfaces.

[12] A. Barducci, G. Bussi, and M. Parrinello, Phys. Rev. Lett. **100**, 020603 (2008).

[13] J. F. Dama, M. Parrinello, and G. A. Voth, Phys. Rev. Lett. **112**, 240602 (2014).

[14] H. J. Kushner and G. G. Yin, *Stochastic Approximation and Recursive Algorithms and Applications* (Springer-Verlag, New York, 2003).

[15] F. Bach and E. Moulines, in *Advances in Neural Information Processing Systems 26*, edited by C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger (Curran Associates, Inc., Red Hook, NY, 2013), pp. 773–781.

[16] G. Bussi, F. L. Gervasio, A. Laio, and M. Parrinello, J. Am. Chem. Soc. **128**, 13435 (2006).

[17] P. Raiteri, A. Laio, F. L. Gervasio, C. Micheletti, and M. Parrinello, J. Phys. Chem. B **110**, 3533 (2006).

[18] S. Piana and A. Laio, J. Phys. Chem. B **111**, 4553 (2007).

[19] P. Tiwary and M. Parrinello, J. Phys. Chem. B, doi:10.1021/jp504920s, (2014).

[20] G. A. Tribello, M. Bonomi, D. Branduardi, C. Camilloni, and G. Bussi, Comput. Phys. Commun. **185**, 604 (2014).