

Well-Tempered Metadynamics Converges Asymptotically

James F. Dama,¹ Michele Parrinello,² and Gregory A. Voth¹

¹*Department of Chemistry, James Franck Institute, Institute for Biophysical Dynamics, and Computation Institute, University of Chicago, Chicago, Illinois 60637, USA*

²*Department of Chemistry and Applied Biosciences, ETH Zurich, and Facoltà di Informatica, Istituto di Scienze Computazionali, Università della Svizzera Italiana, Via G. Buffi 13, 6900 Lugano, Switzerland*

(Received 11 April 2014; published 18 June 2014)

Metadynamics is a versatile and capable enhanced sampling method for the computational study of soft matter materials and biomolecular systems. However, over a decade of application and several attempts to give this adaptive umbrella sampling method a firm theoretical grounding prove that a rigorous convergence analysis is elusive. This Letter describes such an analysis, demonstrating that well-tempered metadynamics converges to the final state it was designed to reach and, therefore, that the simple formulas currently used to interpret the final converged state of tempered metadynamics are correct and exact. The results do not rely on any assumption that the collective variable dynamics are effectively Brownian or any idealizations of the hill deposition function; instead, they suggest new, more permissive criteria for the method to be well behaved. The results apply to tempered metadynamics with or without adaptive Gaussians or boundary corrections and whether the bias is stored approximately on a grid or exactly.

DOI: 10.1103/PhysRevLett.112.240602

PACS numbers: 05.10.-a, 02.50.-r, 02.70.-c

Metadynamics [1–4] accelerates computational simulations of molecular dynamics [5] along a set of several otherwise slow collective variables (CVs) characterized by the presence of multiple metastable states, such as solid lattice order parameters [6–8], molecular dihedral angles [1,9], or protein-ligand binding coordinates [10–12], to enable the calculation of free energy surfaces that would be prohibitively expensive to calculate using unbiased simulations. It accelerates sampling by iteratively constructing and applying a bias potential to offset free energy barriers along the CVs [13]; the iterative construction rule is to add smooth hills to the bias around each successively sampled point to reduce reexploration of already-sampled CV space and thereby enhance escape from local minima [1,14,15]. It is one of a diverse set of adaptive biasing methods for efficient calculation of free energy surfaces (see, e.g., Refs. [13,15–23]) that also includes the widely applied Wang-Landau scheme [24–28]. Though conceived for the study of soft matter, like the Wang-Landau scheme its strategy applies to the efficient sampling of any multimodal distribution in which the most pernicious multimodality can be characterized in terms of just a few variables [29–32].

Though metadynamics is widely used, open questions about its convergence properties have cast doubt on the accuracy of its results and made innovations appear difficult to justify in the minds of some and have impeded its adoption in other fields [19,22,23,31–40]. The most powerful extant convergence results, due to Bussi *et al.* [34] and Dickson [39], relied on idealizing assumptions that metadynamics updates continuously in time and that the CV dynamics is effectively Brownian. Bussi *et al.* showed that metadynamics as originally formulated does not lead to a

single predictable bias under those assumptions—instead, the bias fluctuates indefinitely. Dickson showed that an analytically convenient approximation of the variant well-tempered metadynamics (WTMetaD) does converge under those assumptions, but for this approximation the bias converges to a mollified version of the desired end state—leaving roughness in the final sampling distribution. From another branch of physics, work on Wang-Landau-like methods [41–43] indirectly suggests that WTMetaD using discrete CVs, hills that bias one CV state at a time, and finite time between updates could converge as expected under fairly general CV dynamics—but again, this speaks little to practical metadynamics.

As yet there has been no rigorous demonstration that metadynamics as actually implemented [44–46], with finite-width hills deposited at finite time intervals according to a sampling dynamics with imperfect separation of time scales, is a stochastic process with a single end state that is always reached given enough time—or, if it is, that the end state reached is the intended one, as Dickson suggested might not be the case [22,39]. This Letter provides the steps in such a demonstration and the essential technical constraints that arise along the way.

Our work does not rely on assumptions of effective Brownian dynamics, continuous time updating, or idealizations of the hill function. The demonstration treats the method as implemented by applying reliable stochastic mathematics [47]. The remaining body of this Letter describes the physical argument of the proof unburdened by technical mathematics; see the references for technical details. The argument requires understanding a key auxiliary time scale, how the evolution of the bias in that time

scale approaches an ordinary differential equation (ODE), the stable state for that ODE, and why the ODE converges to that stable state. It is a precise argument under fairly general conditions that are specified later in the Letter and does not contradict practical experience: though formal convergence is robust, the actual time to convergence may remain sensitive to system details, especially the presence of slow variables orthogonal to the CVs under study [3,21].

Metadynamics iteratively constructs an umbrella potential as a sum of hill functions around successively sampled states, a process referred to as “filling the free energy surface with computational sand”[1]. These hills are smooth functions of just a few CVs and come in a variety of shapes. As a guiding example consider the typical case of Gaussian hills on a 1D periodic domain as illustrated in Fig. 1. In WMetaD, hill shape remains constant but the hill height decreases as the bias at the sampled point becomes larger. The bias, which increases monotonically, eventually changes very slowly with time [2]. The precise rule for WMetaD is

$$V_{n+1}(s) = V_n(s) + G(s, s_{n+1})e^{-V_n(s_{n+1})/\Delta T}, \quad (1)$$

letting s denote a CV state point, $G(s, s')$ an energy-valued hill function of s to be deposited when point s' is sampled, ΔT an energy-valued positive scalar parameter, s_n a sequence of CV state samples from the biased trajectory at times t_n , and $V_n(s)$ the sequence of biases at times t_n . s may be a vector [1], a discrete index [37], a function [48], or any combined set of these. The t_n are spaced evenly, and the underlying system’s dynamics must be Markovian and ergodic for any fixed bias applied during the simulation. A Gaussian hill in 1D on a periodic interval corresponds to $G(s, s') = h \exp[-(s - s')^2/\delta^2]$ where h is the hill height, δ is the hill width, and the $s - s'$ subtraction obeys the minimum image convention.

To see how metadynamics converges begin by noting that the bias has two parts with essentially different functions, illustrated in Fig. 1. The first is the average level of the bias $\bar{V}_n = \int V_n(s)ds / \int ds$. This level does not

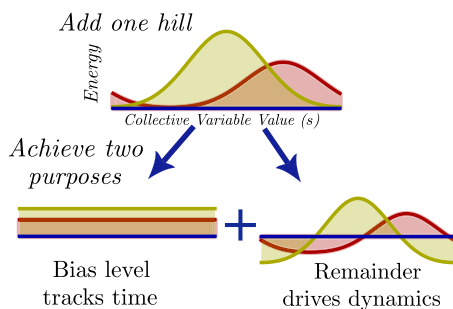


FIG. 1 (color online). Two Gaussian hills on a 1D periodic domain (top) split into hill level components (lower left) and driving bias components (lower right). The hill level exerts no forces and serves only to mark time.

affect the system dynamics. It serves only to record the total amount of bias, or “computational sand,” deposited so far, and it never stops increasing. It appears as nothing more than an intriguing way to measure the elapsed time in a simulation, as if the computational sand were emerging from an hourglass. The second is the driving bias $\tilde{V}_n(s) = V_n(s) - \bar{V}_n$. This portion accounts for all effects of the bias on dynamics and we assume that its values (and derivatives if relevant) stay within a finite range by nature—in practice, metadynamics does not give rise to infinite energy differences or forces except in pathological cases [33,34,37,40]. In technical terms, we assume the driving biases remain in a bounded compact subset of bias space.

To track the evolutions of these parts separately to reveal more about their functions, split the update into two parts:

$$\bar{V}_{n+1} = \bar{V}_n + \exp(-\bar{V}_n/\Delta T)\lambda(s_{n+1}, \tilde{V}_n), \quad (2)$$

$$\tilde{V}_{n+1}(s) = \tilde{V}_n(s) + \exp(-\bar{V}_n/\Delta T)\Gamma(s, s_{n+1}, \tilde{V}_n), \quad (3)$$

where the hill level $\lambda(s', \tilde{V}) = \int G(s, s')e^{-\tilde{V}(s')/\Delta T}ds / \int ds$ and a driving hill function $\Gamma(s, s', \tilde{V}) = G(s, s')e^{-\tilde{V}(s')/\Delta T} - \lambda(s', \tilde{V})$ account for how adding hills affects the average level of the bias and the driving bias, respectively. Written in this manner it is apparent that $\exp(-\bar{V}_n/\Delta T)$ functions as an ever-decreasing step size, and the sums of $\exp(-\bar{V}_n/\Delta T)$ make a natural auxiliary time scale. Because \tilde{V}_n is bounded by assumption and Gaussian hills are positive and bounded, $\lambda(s_{n+1}, \tilde{V}_n)$ must be bounded from above and below by some M_u and M_l ,

$$M_l \leq \lambda(s_{n+1}, \tilde{V}_n) = (\bar{V}_{n+1} - \bar{V}_n) \exp(\bar{V}_n/\Delta T) \leq M_u. \quad (4)$$

Therefore, for the same reasons that $\dot{x}e^x \sim 1$ implies $x \sim \ln t$ and $e^{-x} \sim 1/t$,

$$b_l/n \leq \exp(-\bar{V}_n/\Delta T) \leq b_u/n \quad (5)$$

for all large n for some positive b_l and b_u . An infinite sum of such terms always diverges, which implies there exists a time scale τ that is one to one with the original simulation time scale t and satisfies

$$\tau(t_n) \equiv \sum_{i=0}^{n-1} \exp(-\bar{V}_i/\Delta T). \quad (6)$$

Examining the updates in this time scale will reveal a simple structure to asymptotics of the overall method.

At early times, $\exp(-\bar{V}/\Delta T)$ is large, so few hills must be added to accumulate a given increment of τ time. As time goes on, $\exp(-\bar{V}/\Delta T)$ becomes smaller, so more hills must be added to accumulate the same increment. This is illustrated in Fig. 2. Because each hill is proportional to $\exp(-\bar{V}/\Delta T)$, the bias changes overall by an amount




FIG. 2 (color online). Typical cumulative bias updates from a flat sampling distribution over short intervals of equal τ durations ($d\tau$) at early (top) and late (bottom) times. Yellow lines denote the bias after each successive update during the interval, with the top line in each figure the final updated bias. Red dashed lines denote the change in average bias level. As time goes on, the total bias deposited in each interval of fixed τ duration stays roughly constant while the number of hills and the quality of sampling both increase.

roughly proportional to the length of the increment. As time goes on, even the shortest increments $d\tau$ will see the addition of many hills—corresponding to sampling for long t times—yet because they are short the change in the bias will be small. Because the change in bias over each of these short increments is small, the update over each is approximately as if all of the sampling were done on the starting bias surface; because the sampling goes for a long time and many hills are added, the update over each is approximately an integral over hill functions chosen from the biased equilibrium distribution $p_b(s) \sim \exp(-[F(s) + V(s)]/T)$. Together, the observations in the previous sentence imply that the cumulative update rule in the limit of large τ and small $d\tau$ is well approximated by

$$\tilde{V}(s, \tau + d\tau) = \tilde{V}(s, \tau) + d\tau \int \Gamma(s, s', \tilde{V}(\tau)) p_b(s') ds', \quad (7)$$

suggesting the asymptotics follow the long-time differential equation

$$d\tilde{V}(s, \tau)/d\tau = \int \Gamma(s, s', \tilde{V}(\tau)) p_b(s') ds'. \quad (8)$$

where $r(\tau)$ is a positive parameter involving T , ΔT , and normalization factors for $p_w(s, \tau)$ (see the algebra in the Supplemental Material [50]). When $G(s, s') = h e^{-(s-s')^2/\delta^2}$, this derivative is always less than zero unless $p_w(s) = p_c(s)$, in which case it is exactly zero. Thus the

A related equation can be found in the appendix of Branduardi *et al.* [49]. Though this argument may seem heuristic, it is precise given conditions defined later in this Letter [47].

Next, given this asymptotic ODE one must characterize its asymptotic behavior in turn. One viable strategy is to determine the fixed points of the equation and then characterize the stability of those fixed points. The fixed points are those biases for which $\int \Gamma(s, s', \tilde{V}) p_b(s') ds' = 0$, $\int G(s, s') \exp[-\tilde{V}(s')/\Delta T] p_b(s') ds' = C$ for some constant C , a very appealing result because it corresponds to the heuristic used to motivate metadynamics when it was first introduced: assuming perfect sampling [1]. For the 1D Gaussians being considered, this amounts to

$$e^{-\tilde{V}(s)/\Delta T} e^{-[F(s)+\tilde{V}(s)]/T} \sim C, \quad (9)$$

or

$$\tilde{V}(s) = -(1 + T/\Delta T)^{-1} F(s) + C' \quad (10)$$

for a different constant C' —exactly the end state originally suggested for WTMetaD [2].

Already, the fact that there is only one unique fixed point in this case is a good sign. In fact, the fixed point is both unique and globally attracting; i.e., the ODE will converge to the desired end state whatever its physically reasonable initial state, as can readily be verified with an appropriate choice of Lyapunov function. One choice is the Kullback-Leibler divergence (relative entropy) of the current tempering-reweighted state $p_w(s, \tau) \sim \exp[-\tilde{V}(s, \tau)/\Delta T] \times \exp(-[F(s) + \tilde{V}(s, \tau)]/T)$ from the target sampling end state $p_w(s, \tau) \sim 1$ named $p_c(s)$, given by

$$D(p_c || p_w) = \int p_c(s) \ln \frac{p_w(s, \tau)}{p_c(s)} ds. \quad (11)$$

After a series of algebraic manipulations, the τ time derivative is given by

$$dD(p_c || p_w)/d\tau = -r(\tau) \int \int [p_w(s, \tau) - p_c(s)] G(s, s') [p_w(s', \tau) - p_c(s')] ds ds', \quad (12)$$

relative entropy is a Lyapunov function for the ODE and the driving bias satisfying Eq. (10) is a global attractor. The method converges asymptotically.

With that guiding example followed to its conclusion, now consider adapting the argument for more general

cases. The most practically important generalization is across hill functions, e.g., geometrically adapted hills [49], hills designed for metadynamics on finite domains [37,40,51], and hills interpolated over discrete grids [44]. Considering the update Eq. (1) once more but leaving the hill function unspecified, one may follow the same intuitions given design constraints on $G(s, s')$.

First, to find the time scale τ well behaved, the quantity $\int G(s, s') ds' / \int ds$ must exist and be positive for any s' : hills must add bias, enough to eventually fill the region, so the region must be finite in size. Also, ΔT must be finite and positive. Second, for the ODE to have meaningful fixed points, the equation

$$\int G(s, s') p_c(s') ds' = C \quad (13)$$

must have at least one positive, normalized, physically possible solution $p_c(s)$ for some constant C : there must be a possible sampling distribution from which the bias update will not change that distribution. Finally, for relative entropies to be Lyapunov functions, $G(s, s')$ must be positive semidefinite: the hills must be deposited so that they encourage escape from local traps, never tightening of traps.

Given all of these conditions on $G(s, s')$, the reasoning above can be extended to show that the instantaneous tempering-reweighted distribution for any such method asymptotically approaches the set of normalized solutions of Eq. (13), possibly converging to one or possibly fluctuating in the set indefinitely: the Lyapunov function is the infimum of the relative entropies above over all physically realizable choices of p_c . Using smooth, boundary-consistent Gaussian hills satisfies all these conditions and guarantees a unique solution. The nonuniqueness and possible fluctuation are confined to the details that the hill function cannot resolve: the null space of $\int G(s, s') ds'$. For instance, if the hills are chosen to be step functions on certain regions of CV space, the detail within each region will not be flattened but all coarser aspects of the free energy surface will be—as in Wang-Landau sampling [24].

The analysis above also suggests more general tempering schemes. A new type of metadynamics, introduced here and inspired by density of states methods [24,27,35] and self-healing umbrella sampling [19], is “globally tempered metadynamics,” (GTMetaD), where hill size is independent of location and adaptively set based on global properties of the bias; the update rule here is

$$V_{n+1}(s) = V_n(s) + w[V_n]G(s, s_{n+1}), \quad (14)$$

where $w[V]$ is a unitless positive scalar functional of the bias and the other variables are as before. If $w[V]$ can be factored into bias-level and driving-bias components as $\omega[\tilde{V}] \exp(-\tilde{V}/\Delta T)$, then all of the previous convergence reasoning for WTMetaD also holds for GTMetaD. Simple

GTMetaD schemes of this form appear to offer substantial advantages over WTMetaD [27,35] and, furthermore, one could use both temperings together in hybrid schemes.

The part of the argument above requiring the most careful technical analysis is the derivation of the long-time ODE Eq. (8). Precise, permissive sufficient conditions for this intuition to hold constitute an important result for the wider field of adaptive biasing enhanced sampling methodology. It can be proven using the method of Theorem 6.6.1 from Kushner and Yin [47] given two additional conditions on the underlying system dynamics. Roughly put, this theorem proves that a stochastic iteration of the form $\theta_{n+1} = \theta_n + \epsilon_n Y_n(\theta_n, x_n)$, where θ_n are the desired iterates, ϵ_n are step sizes, and $Y_n(\theta_n, x_n)$ are updates that depend on the current iterate θ_n and the randomly sampled state x_n of a Markov model whose transitions may depend on θ_n , converges like solutions of the differential equation $d\theta/d\tau = \langle Y(\theta, x) \rangle$ with probability 1, where $\tau = \sum \epsilon_n$ and the expectation is taken over the equilibrium distribution of x given θ . It holds if the sum of the ϵ_n diverges while the sum of ϵ_n^2 converges, the Y_n are bounded, and a few other technical conditions hold. In our case, θ is the driving bias and the theorem applies to Eq. (3) with s_{n+1} a projection of a fine-grained system state x_n sampled from a fixed-bias-Markovian dynamics.

The first technical condition not already established matches the claim that one can approximate cumulative bias updates over fixed τ increments as adding infinite hills chosen from the biased equilibrium distribution. A precise requirement is that however many hills are added, the expected difference between updating from the actual sampling distribution and from the quasiequilibrium sampling distribution should always be at least as small as a single update step, i.e.,

$$\begin{aligned} & \sum_{m=n+1}^{\infty} e^{-\tilde{V}_m/\Delta T} (E_{\text{ne}, \tilde{V}_m, x_n} - E_{e, \tilde{V}_m}) \Gamma(s, s_m, \tilde{V}_m) \\ & = O(e^{-\tilde{V}_n/\Delta T}), \end{aligned} \quad (15)$$

for every s, x_n visited, and V_n visited. Here $E_{\text{ne}, \tilde{V}_n, x_n}$ is the expectation over system states given nonequilibrium dynamics under bias \tilde{V}_n given that the (fully specified) system state is x_n at t_n , and E_{e, \tilde{V}_n} is the equilibrium expectation under bias \tilde{V}_n . Since $\Gamma(s, s_m, \tilde{V}_n)$ is independent of t and the ϵ_n decay as $1/n$, if the system is Markovian and time invariant this reduces to

$$\sum_{m=n+1}^{\infty} \frac{1}{m} [p_{V_n}(s, x_n, t_m - t_n) - p_b(s, V_n)] = O(1/n). \quad (16)$$

Here $p_V(s, x, t)$ is the probability of going from a system state x to any system state that maps to CV state s in time t with the dynamics subject to fixed bias V , while $p_b(s, V)$ is

the equilibrium CV distribution under fixed bias V . Note that it is each fixed-bias dynamics, not the CV dynamics or adaptively biased dynamics, that is Markovian here—a standard case for molecular dynamics and Monte Carlo simulations. This sum measures how quickly a system approaches equilibrium, so the condition is a more precise formulation of the intuition that metadynamics works only when the system equilibrates reasonably quickly. If non-equilibrium fluctuations in the system decay exponentially or stretched exponentially, the usual case for nondriven ergodic systems away from criticality, then the differences in the sum decay properly and the condition is satisfied.

The second matches the claim that one can approximate the cumulative update over a τ time increment as if all of the sampling were done on the starting bias surface. This can only be true if the sum of many hills does not change too rapidly with small changes in the bias. In more precise terms, the sum in the previous paragraph must have derivatives with respect to V_n bounded by the step size,

$$\sum_{m=n}^{\infty} \frac{1}{m} \frac{\delta(p_{V_n}(s, x_n, t_m - t_n) - p_b(s, V_n))}{\delta V_n} = O(1/n) \quad (17)$$

for all s , x_n visited, and V_n visited. A sufficient condition for this is that $p_V(s, x, t)$ should have bounded functional derivatives with respect to V that relax to their long-time value with any algebraic dependence faster than $1/t$. This amounts to a precise formulation of the intuition that metadynamics only works if the system does not jam or become unstable in response to the bias, and it is also a permissive condition for nondriven ergodic systems away from a phase transition. It will typically be satisfied whenever the system is finite and the short-time generator of dynamics is a smooth function of the bias, for instance when the bias is implemented simply as an added force in Langevin molecular dynamics [1] or as an added energy in Metropolis Monte Carlo calculations [37].

These two conditions, together with the stated design constraints on the hill function and the assumption that the driving bias stays properly bounded throughout the simulation, are enough to allow for a detailed proof that WTMetaD as actually implemented converges asymptotically and converges to simple, designable end states [47]. However, it gives no guarantees that it will converge quickly, especially when used with poorly chosen CVs. The end states are exactly those suggested by the simple quasiequilibrium heuristic that Laio and Parrinello used to motivate metadynamics when it was first introduced [1], with no mollification artifacts [39], and Bussi *et al.*'s gauge free energy concept [34] proves key. This analysis is flexible enough to cover metadynamics with or without adaptive Gaussians [49] or boundary corrections [37,40,51], and with or without an approximating grid [44]. Furthermore it provides design constraints for such methods and demonstrates a powerful yet little-known tool for rigorously

characterizing nonequilibrium, non-Markovian systems. As sampling methods become ever more sophisticated and diverse, ensuring that sufficiently powerful analysis strategies are brought in to complement them becomes ever more important.

This material is based upon work supported by the National Science Foundation (NSF Grant No. CHE-1214087) and the European Union (Grant No. ERC-2009-AdG-247075). The authors would like to thank Professor Jonathan Weare for thoughtful advice, a close reading, and recommending K&Y, and Professor Iliia Horenko for his close reading and advice as well.

-
- [1] A. Laio and M. Parrinello, *Proc. Natl. Acad. Sci. U.S.A.* **99**, 12562 (2002).
 - [2] A. Barducci, G. Bussi, and M. Parrinello, *Phys. Rev. Lett.* **100**, 020603 (2008).
 - [3] A. Barducci, M. Bonomi, and M. Parrinello, *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **1**, 826 (2011).
 - [4] L. Sutto, S. Marsili, and F.L. Gervasio, *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2**, 771 (2012).
 - [5] M. Karplus and J. A. McCammon, *Nat. Struct. Biol.* **9**, 646 (2002).
 - [6] R. Martoňák, A. Laio, and M. Parrinello, *Phys. Rev. Lett.* **90**, 075503 (2003).
 - [7] A. R. Oganov, R. Martoňák, A. Laio, P. Raiteri, and M. Parrinello, *Nature (London)* **438**, 1142 (2005).
 - [8] J. Sun, D. D. Klug, R. Martoňák, J. A. Montoya, M.-S. Lee, S. Scandolo, and E. Tosatti, *Proc. Natl. Acad. Sci. U.S.A.* **106**, 6077 (2009).
 - [9] J. P. M. Jämbeck and A. P. Lyubartsev, *J. Phys. Chem. Lett.* **4**, 1781 (2013).
 - [10] F.L. Gervasio, A. Laio, and M. Parrinello, *J. Am. Chem. Soc.* **127**, 2600 (2005).
 - [11] J. Pfandner, D. Branduardi, M. Parrinello, T.D. Pollard, and G. A. Voth, *Proc. Natl. Acad. Sci. U.S.A.* **106**, 12723 (2009).
 - [12] V. Limongelli, M. Bonomi, L. Marinelli, F.L. Gervasio, A. Cavalli, E. Novellino, and M. Parrinello, *Proc. Natl. Acad. Sci. U.S.A.* **107**, 5411 (2010).
 - [13] M. Mezei, *J. Comput. Phys.* **68**, 237 (1987).
 - [14] F. Glover, *Computers & Operations Research* **13**, 533 (1986).
 - [15] T. Huber, A. E. Torda, and W. F. van Gunsteren, *J. Comput.-Aided Mol. Des.* **8**, 695 (1994).
 - [16] B. A. Berg and T. Neuhaus, *Phys. Rev. Lett.* **68**, 9 (1992).
 - [17] C. Bartels and M. Karplus, *J. Comput. Chem.* **18**, 1450 (1997).
 - [18] E. Darve and A. Pohorille, *J. Chem. Phys.* **115**, 9169 (2001).
 - [19] S. Marsili, A. Barducci, R. Chelli, P. Procacci, and V. Schettino, *J. Phys. Chem. B* **110**, 14011 (2006).
 - [20] E. Darve, D. Rodríguez-Gómez, and A. Pohorille, *J. Chem. Phys.* **128**, 144120 (2008).
 - [21] L. Zheng, M. Chen, and W. Yang, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 20227 (2008).
 - [22] B. M. Dickson, F. Legoll, T. Lelièvre, G. Stoltz, and P. Fleurat-Lessard, *J. Phys. Chem. B* **114**, 5823 (2010).

- [23] S. Singh, C.-c. Chiu, and J. J. Pablo, *J. Stat. Phys.* **145**, 932 (2011).
- [24] F. Wang and D. P. Landau, *Phys. Rev. Lett.* **86**, 2050 (2001).
- [25] K. M. Aberg, A. P. Lyubartsev, S. P. Jacobsson, and A. Laaksonen, *J. Chem. Phys.* **120**, 3770 (2004).
- [26] S. Trebst, D. A. Huse, and M. Troyer, *Phys. Rev. E* **70**, 046701 (2004).
- [27] R. E. Belardinelli and V. D. Pereyra, *Phys. Rev. E* **75**, 046701 (2007).
- [28] F. A. Escobedo and F. J. Martinez-Veracoechea, *J. Chem. Phys.* **129**, 154107 (2008).
- [29] F. Liang, C. Liu, and R. J. Carroll, *J. Am. Stat. Assoc.* **102**, 305 (2007).
- [30] Y. F. Atchadé and J. S. Liu, *Statistica Sinica* **20**, 209 (2010).
- [31] N. Chopin, T. Lelièvre, and G. Stoltz, *Stat. Comput.* **22**, 897 (2012).
- [32] I. Bilonis and P. Koutsourelakis, *J. Comput. Phys.* **231**, 3849 (2012).
- [33] A. Laio, A. Rodriguez-Fortea, F. L. Gervasio, M. Ceccarelli, and M. Parrinello, *J. Phys. Chem. B* **109**, 6714 (2005).
- [34] G. Bussi, A. Laio, and M. Parrinello, *Phys. Rev. Lett.* **96**, 090601 (2006).
- [35] D. Min, Y. Liu, I. Carbone, and W. Yang, *J. Chem. Phys.* **126**, 194104 (2007).
- [36] T. Lelièvre, M. Rousset, and G. Stoltz, *Nonlinearity* **21**, 1155 (2008).
- [37] Y. Crespo, F. Marinelli, F. Pietrucci, and A. Laio, *Phys. Rev. E* **81**, 055701 (2010).
- [38] G. A. Tribello, M. Ceriotti, and M. Parrinello, *Proc. Natl. Acad. Sci. U.S.A.* **107**, 17509 (2010).
- [39] B. M. Dickson, *Phys. Rev. E* **84**, 037701 (2011).
- [40] M. McGovern and J. de Pablo, *J. Chem. Phys.* **139**, 084102 (2013).
- [41] R. E. Belardinelli and V. D. Pereyra, *J. Chem. Phys.* **127**, 184105 (2007).
- [42] R. E. Belardinelli, S. Manzi, and V. D. Pereyra, *Phys. Rev. E* **78**, 067701 (2008).
- [43] G. Fort, B. Jourdain, E. Kuhn, T. Lelièvre, and G. Stoltz, *arXiv:1207.6880*.
- [44] M. Bonomi, D. Branduardi, G. Bussi, C. Camilloni, D. Provasi, P. Raiteri, D. Donadio, F. Marinelli, F. Pietrucci, R. A. Broglia, and M. Parrinello, *Comput. Phys. Commun.* **180**, 1961 (2009).
- [45] G. Fiorin, M. L. Klein, and J. Héning, *Mol. Phys.* **111**, 3345 (2013).
- [46] G. A. Tribello, M. Bonomi, D. Branduardi, C. Camilloni, and G. Bussi, *Comput. Phys. Commun.* **185**, 604 (2014).
- [47] H. J. Kushner and G. G. Yin, *Stochastic Approximation and Recursive Algorithms and Applications* (Springer, New York, 1997), 2nd ed.
- [48] G. A. Tribello, M. Ceriotti, and M. Parrinello, *Proc. Natl. Acad. Sci. U.S.A.* **109**, 5196 (2012).
- [49] D. Branduardi, G. Bussi, and M. Parrinello, *J. Chem. Theory Comput.* **8**, 2247 (2012).
- [50] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevLett.112.240602> for detailed discussion of the Lyapunov function Eq. (11), especially algebraic manipulations used to derive Eq. (12) and its analog in the case p_c is not unique.
- [51] F. Baftizadeh, X. Biarnes, F. Pietrucci, F. Affinito, and A. Laio, *J. Am. Chem. Soc.* **134**, 3886 (2012).