Improving fast-switching free energy estimates by dynamical freezing

Paolo Nicolini and Riccardo Chelli*

Dipartimento di Chimica, Università di Firenze, Via della Lastruccia 3, I-50019 Sesto Fiorentino, Italy and European Laboratory for Non-linear Spectroscopy (LENS), Via Nello Carrara 1, I-50019 Sesto Fiorentino, Italy

(Received 18 June 2009; revised manuscript received 1 September 2009; published 22 October 2009)

An important limitation of nonequilibrium pulling experiments/simulations in recovering free energy differences is the poor convergence of path-ensemble averages. Therefore, a large number of fast-switching trajectories needs to achieve free energy estimates with acceptable accuracy. We propose a method to improve free energy estimates by drastically lowering the computational cost of steered molecular dynamics simulations employed to realize such trajectories. This is accomplished by generating trajectories where the particles not directly involved in the driven process are dynamically frozen. Such a freezing is dynamical rather than thermal because it is reached by a synchronous scaling of atomic masses and velocities keeping the kinetic energy of each particle unchanged. The forces between dynamically frozen particles can then be calculated rarely. Thus, it is possible to generate realizations of a process whose computational cost is not correlated with the size of the whole system, but only with that of the reaction site. The method is illustrated on a simple model system and its general applicability is discussed.

DOI: 10.1103/PhysRevE.80.041124 PACS number(s): 05.70.Ln, 02.70.Ns, 05.20.−y, 05.70.Ce

I. INTRODUCTION

The estimate of free energy is of basic importance in many fields of physics, chemistry, and biology. A rough classification of the methods devised for computing free energy differences can be based on the possibility of sampling a system at equilibrium or out of equilibrium. In the latter context, an interesting scenario has been disclosed by two nonequilibrium work relations, the Jarzynski equality [1] (JE) and the Crooks fluctuation theorem [2] (CFT), relating the free energy difference between two thermodynamic states to the external work performed in an ensemble of realizations switching the system between such states. A significant difference between JE and CFT is that, while the latter gives access only to free energy differences, the former can provide the whole potential of mean force (PMF) along a chosen collective coordinate. In spite of this advantage, the JE is known to furnish strongly biased free energy estimates [3–6], because the involved path-ensemble average depends crucially on a small fraction of realizations with negative dissipative work. With regard to CFT, Shirts and Pande showed [5] that the Bennett acceptance ratio [7], which may be proved using the CFT [8], can be significantly more accurate than JE. As stated above, the main limitation of the CFT is that only free energy differences can be recovered. To tackle this problem, few methods have been proposed [9–11].

From the computational point of view, JE and CFT are typically employed in numerical analysis of atomic trajectories obtained from steered molecular dynamics (MD) simulations [4,12]. For most applications, the prior target would be to improve the efficiency of path sampling, in order to reduce the computational burden which is correlated indirectly with the amount of work dissipated during the realizations. To this aim, several approaches have been developed. They include biased path sampling [13–16], generation of non-Hamiltonian equations of motion [6] and optimal protocol strategies [17]. Significant speed-up was also obtained by MD simulations with large time-steps [18]. A common limitation of all these methods lies in their dependence on the sample size, because all interparticle forces must be calculated to evolve the system in time. This means that, also the particles not involved directly in the reaction process do affect the overall computational cost.

In this article, we present a strategy for improving the efficiency of fast-switching free energy estimates via JE or CFT for large systems. The basic idea is to freeze the dynamics of a subset of particles which are supposed not to be involved in the driven process, while leaving the particles near the reaction site dynamically active. This is realized with a synchronous scaling of the masses and velocities of the involved particles by keeping their individual kinetic energy unchanged (theoretical aspects of the methodology are presented in Sec. II). Two possible implementations of the mass-velocity scaling approach implying dynamical freezing are proposed. Such procedures allow us to calculate the forces between dynamically frozen particles rarely, with strong improvement of the simulation speed-up. A third possible implementation implying dynamical warm up, instead of freezing, is also discussed.

As an example of potential application, we may cite the calculation of the binding free energy of two solvated molecules. In such a case we might "ignore" the dynamics of the solvent molecules far away from the center of mass of the two target molecules. Another example could be the calculation of the PMF related to processes of molecular traffic in transport proteins. Here, the reaction site can be localized around the protein channel. Of course, it is not always possible to distinguish *a priori* the portion of space where the driven reaction occurs. Typical examples can be found in most protein folding/unfolding processes. However, as we will see, by updating dynamically frozen and unfrozen regions at regular time intervals, it is possible to get a flexible algorithm virtually independent on the dynamical evolution

*riccardo.chelli@unifi.it

of the system. As a test case we report the calculation of the PMF of two particles solvated by a monoatomic fluid, considering the interparticle distance as collective coordinate (Sec. III). Both nonequilibrium work theorems, JE and CFT, are taken into account for estimating the free energy profile. Conclusions and perspectives of our approach are given in Sec. IV.

II. THEORY

A. Background: Jarzynski equality and Crooks fluctuation theorem

The framework in which we move is that of a classical system at the temperature β^{-1} for which we can manipulate some collective coordinate through an external parameter, λ. The Hamiltonian, $H(\mathbf{x};\lambda)$, depends on generalized coordinates $\mathbf{x}$, and, parametrically, on λ. An equilibrium state corresponding to a given value of λ, realized by constraining the collective coordinate, is described by the canonical distribution $p(\mathbf{x};\lambda)=Z_\lambda^{-1}\exp[-\beta H(\mathbf{x};\lambda)]$, with free energy $F_\lambda=-\beta^{-1}\ln Z_\lambda$. The free energy difference, $\Delta F=F_b-F_a$, between two states corresponding to $\lambda=a$ and $\lambda=b$ can be estimated using the JE [1],

$$\exp(-\beta\Delta F)=\langle\exp(-\beta W)\rangle, \tag{1}$$

where the work [W into Eq. (1)] exponential average, denoted by angular brackets, is calculated over an ensemble of realizations switching λ from a to b. A necessary condition for Eq. (1) to be valid is that the initial microstates of the realizations are sampled with canonical probability. If realizations for both directions of the process are available, then ΔF can be estimated using CFT. Various formulations of CFT can be given [19]. Here, we report a detailed form of CFT, which is valid for each pair of time-conjugate phase-space trajectories, $\vec{\Gamma}$ and $\overleftarrow{\Gamma}^*$

$$\frac{p(\vec{\Gamma})}{p(\overleftarrow{\Gamma}^*)}=\exp[\beta(W[\vec{\Gamma}]-\Delta F)], \tag{2}$$

where $p(\vec{\Gamma})$ and $p(\overleftarrow{\Gamma}^*)$ are the probabilities of observing the phase-space trajectory $\vec{\Gamma}$ and its time-conjugate $\overleftarrow{\Gamma}^*$ by sampling the initial microstates from canonical distributions. The quantity $W[\vec{\Gamma}]$ is the work performed on the system during the phase-space trajectory $\vec{\Gamma}$. It is clear that the symmetry condition on the phase-space trajectories is obtainable only if time schedules of the control parameter in the forward and backward directions [$\lambda(t)$ and $\lambda^*(t)$, respectively] are reverse in time, namely, $\lambda(t)=\lambda^*(\tau-t)$, τ being the pulling time. The work entering into JE and CFT can be calculated as $W[\vec{\Gamma}]=\int_0^\tau\partial_t H(\mathbf{x};\lambda(t))dt$. This definition, along with time symmetry condition for $\vec{\Gamma}$ and $\overleftarrow{\Gamma}^*$, implies the equality $W[\vec{\Gamma}]=-W[\overleftarrow{\Gamma}^*]$. Note that, on the basis of the ΔF definition given above, in Eq. (2) the trajectory $\vec{\Gamma}$ is the one during which the control parameter is driven from a to b. We point out that a nondetailed form of the CFT involving work distribution functions related to forward and backward directions of the pulling process can be easily derived from Eq. (2) [2,19]. Although this nondetailed form is computationally more manipulable than Eq. (2) [12], a more sound statistics is provided by the Bennett-like CFT formula [8].

B. Dynamical freezing via mass-velocity scaling

In our approach, the initial microstates of the realizations are obtained by means of standard equilibrium MD simulations using, e.g., constant-volume constant-temperature (NVT) equations of motion. This stage does not differ from the one usually employed in pulling MD simulations. The innovation stems exclusively from the strategy devised to carry out pulling trajectories, that, in addition to mechanical external work, also accounts for mass-velocity scaling. Specifically, once the particle positions and momenta are obtained by NVT sampling, we pass to Hamiltonian (constant-volume constant-energy or NVE) dynamics for performing pulling trajectories, during which instantaneous mass-velocity scaling is applied to selected particles (keeping their individual kinetic energy unchanged). From the physical point of view, the change in dynamical laws from NVT to NVE-type corresponds to a break of energy exchange between system and thermostat. This situation is not new in the context of pulling simulations. It is known since the work by Jarzynski [20], who showed that the instantaneous detachment of the thermal bath from a system subject to pulling processes does not affect free energy estimates obtained by Eq. (1). At contrary, care must be taken when mass-velocity scaling is applied. We deal with a sort of NVE$\rightarrow$NVK$\rightarrow$NVE dynamical swapping, where NVK denotes isokinetic dynamics. During NVK dynamics, masses and momenta/velocities of selected particles change instantaneously, while all other dynamical variables do not vary. In such a case, JE and CFT are still valid (indeed, initial microstates are sampled with canonical probability), though external work must be determined using the proper form of the phase-space compressibility. As NVE dynamics takes place, say, in the time interval $[t_1,t_2]$, the work performed on the system corresponds to $H(t_2)-H(t_1)$, where $H(t)$ is the total energy of the system at time t. If isokinetic mass-velocity scaling occurs at a given time t_0 in the interval $[t_1,t_2]$, then additional work must be taken into account. It is given by [21]

$$\begin{aligned} W_{\mathrm{nvk}} &= \lim_{\Delta t\to 0}\left\{U(t_0+\Delta t)-U(t_0-\Delta t)-\beta^{-1}\int_{t_0-\Delta t}^{t_0+\Delta t}\Lambda_{\mathrm{nvk}}dt\right\} \\ &= -\beta^{-1}\lim_{\Delta t\to 0}\left\{\int_{t_0-\Delta t}^{t_0+\Delta t}\Lambda_{\mathrm{nvk}}dt\right\}, \end{aligned} \tag{3}$$

where $U(t)$ is the potential energy of the system and Λ_{nvk} is the phase-space compressibility associated to the NVK equations of motion producing the scaling event. In Eq. (3), the limit indicates that the scaling event is instantaneous.

Let us consider, for simplicity, a system made of N identical particles of mass m. Suppose that at time t_0 (during a pulling trajectory) we want to change the mass of L labeled particles from the "natural" value m to the value m', by keeping their kinetic energy unchanged. This implies that the coordinates of all N particles and the momenta of the $N-L$ unlabeled particles do not change during NVK dynamics,

while momenta of the L labeled particles change according to the isokinetic constraint. After mass-velocity scaling, NVE equations of motion are restored with the new masses and momenta. The equations of motion employed for mass-velocity scaling are [22]

$$\left.\begin{aligned}\dot{p}_i &= F_i - \alpha_i p_i \\ \dot{r}_i &= p_i/m(t)\end{aligned}\right\} \quad i = 1, \ldots, 3L$$

$$\left.\begin{aligned}\dot{p}_i &= F_i \\ \dot{r}_i &= p_i/m\end{aligned}\right\} \quad i = 3L+1, \ldots, 3N, \tag{4}$$

where the index i refers to a single Cartesian component. According to Ref. [22], the parameter α_i is introduced to keep the kinetic energy related to the ith momentum constant. The mass of the labeled particles, $m(t)$, is an explicit function of time, though its time-dependence does not matter. In fact, as we will see, the relevant quantities entering into play are only the initial and final values of the mass, m and m'. In Eq. (4), the last two formula are reported for completeness and represent the (NVE) equations of motion of the other particles. By definition, they do not enter in the dynamics as mass-velocity scaling occurs.

To apply JE and CFT, the end microstates must be consistent with the "native" equilibrium state of the system, i.e., the state in which all masses are equal to m. This means that, if we perform a mass-velocity scaling at the time t_0, then we need to perform the opposite scaling before ending the trajectory, say at time t_0'. In formula, the mass of the labeled particles must vary in time as $m(t)=m+(m'-m)\theta(t-t_0)\theta(t_0'-t)$ during the forward trajectories and as $m(t)=m+(m'-m)\theta(t+t_0-\tau)\theta(\tau-t_0'-t)$ during the backward ones, where τ is the pulling simulation time and $\theta(x)$ is the Heaviside step function. The time-dependence of $m(t)$ in forward and backward directions of the process takes into account the time-symmetry of conjugate trajectories that must be satisfied for CFT to be valid. We now focus on the estimate of W_{nvk} by means of Eq. (3). The phase-space compressibility arising from Eq. (4) is

$$\Lambda_{\mathrm{nvk}} = \sum_{i=1}^{3N}\left[\frac{\partial \dot{r}_i}{\partial r_i} + \frac{\partial \dot{p}_i}{\partial p_i}\right] = -\sum_{i=1}^{3L}\left[\alpha_i + p_i\frac{\partial \alpha_i}{\partial p_i}\right]. \tag{5}$$

The application of the isokinetic constraint, i.e., $d_t[p_i^2/m(t)]=0$, sets $3L$ conditions, one for each $\dot{p}_i$ component

$$\dot{p}_i = \frac{p_i \dot{m}(t)}{2m(t)}, \quad i = 1, \ldots, 3L. \tag{6}$$

Combining the first line of Eq. (4) with Eq. (6), we obtain an expression for α_i

$$\alpha_i = \frac{F_i}{p_i} - \frac{\dot{m}(t)}{2m(t)}. \tag{7}$$

Upon substituting Eq. (7) into Eq. (5), we get the phase-space compressibility $\Lambda_{\mathrm{nvk}}=3L\dot{m}(t)/[2m(t)]$. Finally, exploiting Λ_{nvk} into Eq. (3), we obtain

$$W_{\mathrm{nvk}} = \frac{3L}{2\beta}\ln\left(\frac{m}{m'}\right). \tag{8}$$

We remark that Eq. (8) could also be derived by derivation of Eq. (6) with respect to p_i and using the ensuing equation into Eq. (5). It is worth noting that, by scaling back the masses of the L labeled particles from m' to m, the opposite work, $-W_{\mathrm{nvk}}$, is obtained. This follows immediately from Eq. (8) upon swapping the masses. Considering the time-dependence of $m(t)$ discussed above, the *global work arising from mass-velocity scaling is null*, regardless of the number of scaling events.

It is now interesting to see how momenta and velocities of the labeled particles change because of the scaling event. For the case of mass-velocity scaling occurring at time t_0, by integrating Eq. (6) side by side, we get

$$p_i(t_0^+) = \left(\frac{m'}{m}\right)^{1/2} p_i(t_0^-), \tag{9}$$

where $p_i(t_0^-)$ and $p_i(t_0^+)$ is the ith component of the momentum before and after the scaling, respectively. The relation existing between the velocity component before and after the scaling event is obtained by exploiting the velocity definition [second line of Eq. (4)] into Eq. (9),

$$\dot{r}_i(t_0^+) = \left(\frac{m}{m'}\right)^{1/2} \dot{r}_i(t_0^-). \tag{10}$$

For instance, if the mass quadruplicates upon scaling, i.e., $m'=4m$, then the momentum doubles, while velocity halves. No other changes occur in the dynamics.

III. RESULTS: IMPLEMENTATIONS OF MASS-VELOCITY SCALING

In this section, we present three possible implementations of the mass-velocity scaling method in order of increasing efficiency and generality. In the simplest implementation [23], we proceed by scaling the masses m and the velocities $\dot{r}$ of *all* atoms at the beginning of the realizations. As seen in the previous section, for a given atom, this can be accomplished by setting $m'=sm$ and $\dot{r}'=s^{-1/2}\dot{r}$, where s is an arbitrary positive parameter that, in principle, can be different for each atom. The realizations are then performed with the new masses m' and with starting velocities $\dot{r}'$ till reaching the final value of λ. In this context, s can simply be viewed as an additional external parameter acting on the system together λ, but, at variance with λ, with any net effect on the total work. In fact, at the end of the realizations, we must restore the natural value of the mass, m, for all particles, thus performing the opposite work (see discussion in Sec. II B). During pulling realizations, the work is only from λ changes and can be computed without specific care about the initial "shock." It is interesting to note that, a uniform mass-velocity scaling (same s for all particles) is dynamically equivalent to an instantaneous change of Δt (the time-step) by the quantity $\Delta(\Delta t)=\Delta t(s^{-1/2}-1)$. Therefore, a scaling down of the masses, $s<1$, corresponds to a time-step increase and vice versa. This means that our approach, limited

to a uniform scaling down of the masses, is complementary to the technique proposed by Lechner *et al.* [18], which is based on large time-steps for improving the MD simulation speed-up. However, it is known that an increase in time-step, while increasing the speed-up of a simulation from one side, on the other side increases dissipation correlated with the energy drift coming from inaccurate integration of the equations of motion [18]. The same effect is observable as the masses are reduced, because of the globally faster dynamics. The advantage of the mass-velocity scaling approach stems from the fact that a faster dynamics accelerates the atomic rearrangements, eventually increasing the reversibility of the nonequilibrium process. Furthermore, we should consider that, in a generic system, the atomic masses can be very different. In such a case, we could take advantage of the fact that the time-step is usually tuned to integrate accurately the fastest motions, which are often correlated with the lightest atoms. This implies that a scaling down of the masses limited to heavy atoms may not affect significantly the dissipation arising from energy-drift [23]. We point out that in the mass-velocity scaling method, invertible phase-space mapping, which is a necessary condition for the validity of both JE and CFT, is always generated.

Although the efficiency of the approach illustrated above is supported by the successful application of the (complementary) large-time-step method [18], few technical aspects may severely limit its performances. The most crucial one is certainly correlated with the system size, whose increase has two basic drawbacks: the increase in energy-drift-related dissipation and the computational overhead due to the evaluation of the interatomic forces. To tackle the problems arising from the system size, we could try to scale the masses up of the atoms that, in the initial microstate, are far away from the reaction site (second implementation). Such a scaling must practically stop the motion of these unimportant atoms so that the forces between them do not need to be calculated during the following dynamics. Care must be taken when selecting the particles that undergo dynamical freezing. For instance, if we consider a process where dissipation is due to interactions between the escorted solute and the solvent (indeed, a common situation in pulling experiments or calculations), then we have to keep the solute and few solvent particles dynamically active, while keeping all other solvent particles dynamically frozen. In principle, given the distinguishability of the particles in MD simulations, we could label a given number of solvent particles, chosen randomly, and then keep them dynamically active in each pulling realization. This procedure would preserve the symmetry of mass-velocity scaling in forward and backward directions of the process and hence the validity of CFT even in its nondetailed form. However, using this protocol, the price we need to pay is the loss of control on the dissipated work. As stated above, an alternative approach could be to select the solvent particles on the basis of a geometric criterion, for example, their distance from the center of mass of the solute particles. Such a solution would allow to reduce the dissipated work significantly. Unluckily, it does not guarantee time-symmetry of mass-velocity scaling in forward and backward directions of the process. In spite of this, it is evident that the distance-based criterion could be a valid approximation as dynamically active particles are localized in a region large enough to embed the space where most dissipation occurs. Therefore, the reliability of this approach (in the context of CFT) is based on the reasonable assumption that dissipation is a local phenomenon in nonequilibrium processes. As we will see, this is just the case of our model-system simulations. As the JE is concerned, the issue of time-symmetry of mass-velocity scaling is not a problem at all. In fact, using JE, we do not need to satisfy the above requirement in the practice. If we use the described geometric criterion, then we shall obtain an ensemble of realizations with different time schedules for mass-velocity scaling. However, if from the one side JE is strictly valid for path-ensemble averages realized with well-established time schedule of the control parameters (including mass-velocity scaling), on the other side we may always calculate the path-ensemble average over an arbitrary collection of time schedules. In other words, the path-ensemble average on the right hand side of Eq. (1) can be performed simultaneously over pulling trajectories and time schedules as well. This makes our approach theoretically correct for JE.

We illustrate the method on the binding process of two particles immersed into a Lennard-Jones fluid. This model system contains much of condensed-phase physics and may be viewed as an elementary example of molecular docking. Specifically, by using steered MD simulations, we calculate the PMF along the distance r between two particles (the solute-particles) interacting through a double-well potential of the form $V(r)=[3(r-1)^2-0.3](r-3)^2$. All quantities are in reduced units. The two solute-particles and the 6798 solvent-particles have the same initial mass m and evolve in a cubic box with standard periodic boundary conditions. The Lennard-Jones potentials for the solvent-solvent and solute-solvent interactions are identical and vanish in the distance range 3.0–3.5 through a cubic switching function. The equations of motion are integrated with a time-step of 5×10^{-3}. Steered MD simulations are carried out using the external potential $5000(r-\lambda)^2$ to switch the system between the end states $\lambda=0.5$ and $\lambda=3.5$. The initial microstates for 5000 forward and 5000 backward realizations are sampled from equilibrium NVT simulations, setting the temperature to 0.8 and the density to 0.85. All realizations are performed at constant pulling speed of 0.3 with NVE equations of motion. The criterion adopted for dynamical freezing, also described above, is based on the distance of the solvent-particles from the centroid of the solute-particles at the initial time. One mobility sphere including the solute-particles and several solvent-particles is thus created. Various cutoff radii for the mobility sphere are tried ($R_m=4$, 6.9, and 10).

In Fig. 1(A) we report representative results of our numerical tests, namely, the free energy profile estimated by the JE [Eq. (1)] in the forward ($\Phi_{\mathrm{JEf}};\lambda=a\equiv0.5\rightarrow\lambda=b\equiv3.5$) and backward ($\Phi_{\mathrm{JEb}};\lambda=b\rightarrow\lambda=a$) directions of the process, and by a CFT-based estimator [9]

$$\Phi_{\mathrm{CFT}}(\lambda)=-\beta^{-1}\ln\left[\left\langle\frac{n_F e^{-\beta W_F(\lambda)}}{n_F+n_B e^{\beta[-W_F(b)+\Delta F]}}\right\rangle_F+\left\langle\frac{n_B e^{\beta[W_B(a)-W_B(\lambda)]}}{n_F+n_B e^{\beta[W_B(a)+\Delta F]}}\right\rangle_B\right], \tag{11}$$

where $W_F(\lambda)$ is the work performed on the system to switch

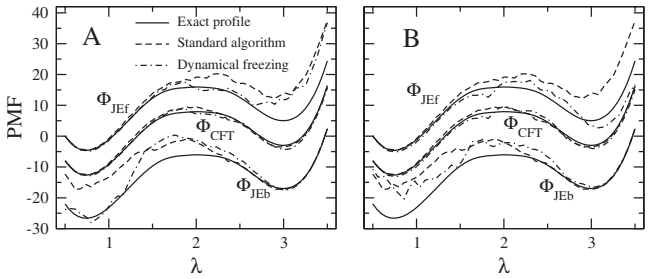

FIG. 1. Potential of mean force as a function of λ, calculated using the 2nd and 3rd implementations of dynamical freezing method (panels A and B, respectively). For the sake of clarity, Φ_{JEf}, Φ_{JEb}, and Φ_{CFT} are shifted. Dashed lines: standard algorithm. Dot-dashed lines: dynamical freezing. Solid lines: exact potential of mean force. The potential of mean force from dynamical freezing in panels A and B is obtained using $R_m=4$ and $R_m=3$, respectively.

the control parameter from a to λ in the forward realizations and $W_B(\lambda)$ is the work performed on the system to switch the control parameter from b to λ in the backward realizations. The angular brackets indicate path-ensemble averages in the forward and backward directions calculated over n_F and n_B realizations, respectively. Finally, $\Delta F=\Phi_{CFT}(b)$ is the free energy difference between the end states and is calculated via Bennett-like formula [8]. In the figure [panel (a)], all free energy profiles estimated by dynamical freezing method are calculated by setting $R_m=4$. The PMFs calculated with standard algorithms, i.e., without dynamical freezing, are also drawn for comparison. The exact PMF, Φ_{TI}, is from thermodynamic integration. The error on Φ_{CFT} is calculated by fitting the free energy offset [11] with a least-squares procedure aimed at minimizing the quantity

$$\sigma^2 = |b-a|^{-1} \int_a^b [\Phi_{CFT}(\lambda) - \Phi_{TI}(\lambda)]^2 d\lambda. \qquad (12)$$

For the Φ_{JEf} and Φ_{JEb} profiles, the offset is taken to satisfy the equalities $\Phi_{JEf}(a)=\Phi_{TI}(a)$ and $\Phi_{JEb}(b)=\Phi_{TI}(b)$. In Table I we report σ for all test cases. The basic outcome gained from Fig. 1(A) and from the table is that the accuracy of dynamical freezing technique is comparable to, or even better than, the standard algorithm. The most relevant difference lies in the simulation speed-up, which is also reported in the table in terms of the average ratio, N_{sa}/N_{df}, between the number of interparticle contacts accounted for in the calculation with standard algorithm (N_{sa}) and with dynamical freezing (N_{df}). The gain in speed-up is even larger than 20 and clearly depends on the ratio between dynamically frozen and unfrozen particles. It is remarkable that a further increase in the system size, while increasing N_{sa}, does not affect N_{df} with obvious improvement of computational efficiency.

TABLE I. Error on the potential of mean force estimates Φ_{JEf}, Φ_{JEb}, and Φ_{CFT} using standard algorithm and dynamical freezing (2nd and 3rd implementations). Last column: speed-up by dynamical freezing in terms of the ratio N_{sa}/N_{df}. The quantity R_m stands for the radius of the mobility spheres (see text for details).

	σ (Φ_{JEf})	σ (Φ_{JEb})	σ (Φ_{CFT})	N_{sa}/N_{df}
Standard algorithm	5.7	5.8	0.7	
2nd impl. $R_m=4$	4.3	3.3	0.6	20.2
2nd impl. $R_m=6.9$	4.2	5.1	0.7	4.5
2nd impl. $R_m=10$	3.4	5.3	0.9	1.6
3rd impl. $R_m=3$	2.3	5.0	0.8	31.5
3rd impl. $R_m=4$	3.8	6.4	0.6	15.5
3rd impl. $R_m=6.9$	3.2	3.8	0.7	3.8
3rd impl. $R_m=10$	4.1	5.1	0.7	1.5

Although the fixed-mobility-sphere implementation gives satisfactory results for our model, a more general approach that does not need the prior knowledge of the dynamical behavior of the system would be preferable. To this aim, we resort to a third implementation that can be summarized as follows: (1) centers for mobility spheres are identified on the basis of the collective coordinate, and in particular, by considering the atoms that are mostly (directly or indirectly) involved in the escorted dynamics. It is evident that also this implementation requires some, however smaller, degree of chemical intuition. For instance, in the complex process of protein folding/unfolding, the choice is straightforward: the mobility spheres can be centered on each protein atom. In our model, two mobility spheres, centered on the solute-particles, are defined. (2) Once the atoms bearing the mobility spheres are established, the masses of the particles out of the mobility spheres are scaled up to a very large, but finite, value (in our tests $s=10^{10}$). (3) Then a period of λ-guided evolution is performed for all particles, including the frozen ones. The aim of this procedure is to guarantee the correct time-evolution of the velocities of the frozen particles, a condition required to apply correctly next mass-velocity scaling (see point 4 below). However, during this period, the forces between frozen particles are not calculated, their initial values being used into equations of motion. (4) At a given arbitrary time, the list of frozen and unfrozen particles is updated on the basis of the new positions of the atoms bearing mobility spheres. Therefore, particles may pass from the dynamically frozen region to some mobility sphere (with ensuing mass-velocity scaling), and vice versa. At this step the forces between frozen particles are calculated again for the following dynamics. Steps 2, 3, and 4 are then repeated, till the end of the realization is reached. Finally, we restore the natural mass m of the particles, thus vanishing the work arising from mass-velocity scaling. Of course, this last operation does not require additional calculations because the work is given by an analytic expression [Eq. (8)]. We point out that this implementation, as the second one (see discussion above), is applicable, in principle, only combined to the JE. On the contrary, due to loss of time-reversal in mass-velocity scaling protocol, CFT does not hold. However, for the same reasons supporting the validity of the second implementation, the CFT is valid in the limit of very large mobility spheres, because the altered dynamics does not affect dissipation. Practically, we find that this limit is reachable using mobility spheres with a radius comparable to the distance to which interparticle interactions vanish. We remark that other similar geometric criteria could be used. For instance, by

updating mobility spheres in order to keep the number of dynamically active solvent-particles constant.

In Fig. 1(B) we report the results obtained using mobility spheres centered on the solute-particles and updated every 2.5 time units (three times during a pulling trajectory). The possibility of updating the mobility spheres allows us to adopt cutoff radii smaller than those used in the fixed-mobility-sphere technique. In the figure, we draw the data obtained with $R_m=3$. Φ_{JEf}, Φ_{JEb}, and Φ_{CFT} from standard algorithm and Φ_{TI} are also drawn. The usual error and speed-up are given in the table for all performed tests. Also in this case, the accuracy of dynamical freezing compares to that of standard algorithm. The slightly lower speed-up for a given R_m with respect to the fixed-mobility-sphere implementation arises from the not complete overlap of the two mobility spheres during the realizations. The speed-up obtained with $R_m=3$ is however impressive, considering the almost unchanged accuracy.

IV. CONCLUSIONS

In this article, we present a computational approach for improving the efficiency of fast switching free energy estimates obtained by the Jarzynski equality or the Crooks fluctuation theorem. The method we propose is distinct from other computational strategies devised to sample or to create trajectories with low dissipation [6,13–16]. Rather, we modify the dynamical behavior of "unimportant" particles by changing, with a sort of instantaneous alchemical transformation, their masses and velocities without affecting the total energy. Speed-up arises from the fact that no mutual interactions need to be calculated among dynamically frozen particles. There is a strict conceptual analogy between dynamical freezing and quantum mechanics/molecular mechanics (QM/MM) method. In the latter case, in fact, unimportant particles, typically those far from the reaction site, are treated with classical (and hence computationally cheap) potentials, while the important particles undergo the most accurate quantum mechanical treatment. The flexibility of dynamical freezing allows for the treatment of a variety of problems in addition to mechanical pulling simulations. We expect it is appropriate to various thermodynamic processes like particle insertion in a fluid or alchemical transformations in general. Moreover, since dynamical freezing does not alter the algorithms usually employed in steered MD simulations, it is prone to be combined with other approaches proposed for improving the efficiency of free energy estimates. For the same reason, implementation of dynamical freezing in MD simulation codes is rather simple.

ACKNOWLEDGMENTS

We thank Piero Procacci and Simone Marsili (Department of Chemistry, University of Florence, Italy) for stimulating discussions and insightful suggestions and Riccardo Petraglia for preliminary tests on the mass-velocity scaling approach (Ref. [23]). This work has been supported by the European Union under Contract No. RII3-CT-2003-506350 and by the Italian Ministry of University.

[1] C. Jarzynski, Phys. Rev. Lett. **78**, 2690 (1997).
[2] G. E. Crooks, J. Stat. Phys. **90**, 1481 (1998).
[3] H. Oberhofer, C. Dellago, and P. L. Geissler, J. Phys. Chem. B **109**, 6902 (2005).
[4] S. Park and K. Schulten, J. Chem. Phys. **120**, 5946 (2004).
[5] M. R. Shirts and V. S. Pande, J. Chem. Phys. **122**, 144107 (2005).
[6] S. Vaikuntanathan and C. Jarzynski, Phys. Rev. Lett. **100**, 190601 (2008).
[7] C. H. Bennett, J. Comput. Phys. **22**, 245 (1976).
[8] M. R. Shirts, E. Bair, G. Hooker, and V. S. Pande, Phys. Rev. Lett. **91**, 140601 (2003).
[9] D. D. L. Minh and A. B. Adib, Phys. Rev. Lett. **100**, 180602 (2008).
[10] R. Chelli, S. Marsili, and P. Procacci, Phys. Rev. E **77**, 031104 (2008).
[11] R. Chelli and P. Procacci, Phys. Chem. Chem. Phys. **11**, 1152 (2009).
[12] P. Procacci, S. Marsili, A. Barducci, G. F. Signorini, and R. Chelli, J. Chem. Phys. **125**, 164101 (2006).
[13] F. M. Ytreberg and D. M. Zuckerman, J. Chem. Phys. **120**, 10876 (2004).
[14] S. X. Sun, J. Chem. Phys. **118**, 5769 (2003).
[15] P. L. Geissler and C. Dellago, J. Phys. Chem. B **108**, 6667 (2004).
[16] D. Wu and D. A. Kofke, J. Chem. Phys. **122**, 204104 (2005).
[17] T. Schmiedl and U. Seifert, Phys. Rev. Lett. **98**, 108301 (2007).
[18] W. Lechner, H. Oberhofer, C. Dellago, and P. L. Geissler, J. Chem. Phys. **124**, 044113 (2006).
[19] G. E. Crooks, Phys. Rev. E **61**, 2361 (2000).
[20] C. Jarzynski, Phys. Rev. E **56**, 5018 (1997).
[21] R. Chelli, J. Chem. Phys. **130**, 054102 (2009).
[22] D. J. Evans and G. P. Morriss, *Statistical Mechanics of Nonequilibrium Liquids* (Academic Press, London, United Kingdom, 1990).
[23] R. Petraglia, Diploma thesis, University of Firenze, Italy, 2007.