Strategies for the exploration of free energy landscapes: unity in diversity and challenges ahead

Fabio Pietrucci

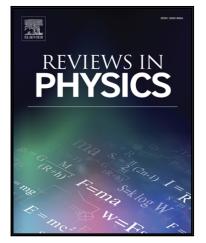# Accepted Manuscript

Strategies for the exploration of free energy landscapes: unity in diversity and challenges ahead

Fabio Pietrucci

Strategies for the exploration of free energy landscapes: unity in diversity and challenges ahead

Fabio Pietrucci

Sorbonne Universités, UPMC University Paris 6, IMPMC,
UMR 7590, CNRS, MNHN, IRD, F-75005 Paris, France*

Computer simulations play an important role in the study of transformation processes of condensed matter, including phase transitions, chemical reactions, and conformational changes of biomolecules. In principle, atomic trajectories, such as those generated using the molecular dynamics approach, contain detailed structural, thermodynamic, and kinetic information about activated processes. In practice, due to free energy barriers, there is often a wide gap between the time scale of the transformation and the time scale accessible with simulations. This review offers a practical guide to the ingenious methods aimed to accelerate the exploration and reconstruction of free energy landscapes of complex systems. The focus is on basic unifying concepts, successful strategies, and pitfalls, illustrated with examples of application to scientific problems from different disciplines. The current challenges in the field consist mainly in the cumbersome identification of optimal reaction coordinates and in the extensive recourse to expert human supervision and fine tuning of the algorithms. The full achievement of wide-spectrum formulations and easy reproducibility of results would constitute the breakthrough necessary to enter the era of routine use of enhanced sampling simulations.

Keywords: activated processes, rare events, free energy landscapes, enhanced sampling, molecular dynamics

I.   INTRODUCTION

Condensed matter at the atomic scale transforms according to a multitude of ingenious, sometimes counter-intuitive mechanisms. A few examples are the nucleation of crystals from liquids, the formation of fullerenes from carbon fragments, the folding of proteins into a precise complex geometry, or the prebiotic reactions creating amino acids and nucleotides from much simpler molecules. The study of these transformation processes - it could be called the science of change - is a fascinating challenge, where computational methods play a major role due to the experimental difficulty in capturing evanescent transition states.

The preceding examples of transformations are activated processes characterized by the following features: the system has different possible metastable structures (free energy minima separated by barriers), each one trapping the dynamics for a time that is long compared to fast bond vibrations, until a jump is eventually performed toward another metastable state. At a fundamental level, time is key to transformation phenomena: it characterizes the kinetics of the process in the form of a rate (i.e., an inverse time). It is a formidable obstacle for simulations, due to the gap between experimentally-relevant time scales and what can be currently achieved with simulations. It plays the main role in the ergodic hypothesis, granting access to ensemble averages in phase space by means of time averages:

$$\langle A \rangle = Z^{-1} \int dx dp \ A(x,p) \ e^{-H(x,p)/k_B T} = \lim_{\tau \to \infty} \tau^{-1} \int_0^\tau dt \ A(x(t),p(t)) \tag{1}$$

(with A an observable, $\{x,p\}$ the phase space coordinates and momenta, and $Z = \int dx dp \ e^{-H(x,p)/k_B T}$); and so on. Ideally, a complete understanding of an activated process would encompass all of its kinetic aspects, however such a description is often considered too difficult to achieve and one is content with reconstructing the geometric pathways and their free energy profile, i.e., formally time-less objects. With improving algorithms and computers, the reconstruction of kinetic rate networks, a large field beyond the scope of this review, will likely gain the center of the scene in the near future.

In principle, molecular dynamics (MD) simulations are ideally suited to study the mechanisms, thermodynamics, and kinetics of activated processes at a given pressure and temperature. An extremely long trajectory of the system, generated from accurate forces, would ergodically sample fluctuations within metastable states as well as all possible transitions among them. To cope with time-scale limitations, but also to get fundamental insight that is not automatically provided by a bare long trajectory, a number of different computational approaches were introduced in the last few decades, sometimes designed on purpose and sometimes borrowed from different disciplines. Beyond preliminary benchmarks on oversimplified models - popular ones being alanine dipeptide, small Lennard-Jones clusters, or the Ising magnet - these enhanced sampling methods should be critically assessed in their ability to solve complex problems of experimental interest.

The aim of the present review is to provide an overview of the main strategies to solve challenging problems of exploration and free energy reconstruction. The focus is on methods based on atomistic trajectories at finite temperature, with an emphasis on underlying concepts and on tricks of the trade. Even if extensive practice is hardly replaced, this practical guide might help readers to mitigate the steep learning curve and choose the most adapted tool for each application. There is a growing number of valuable reviews[1–25] discussing, often in full detail, the technical aspects of different methods.

Notwithstanding strengths and peculiarities of distinct approaches, unity in diversity will be highlighted, as enhanced sampling techniques mostly draw from a rather small common toolbox. It is a well appreciated fact that creativity within a domain is largely a matter of new compositions of existing building blocks.[26] Clearly, not only success but also failure stories need to be critically examined, in the attempt to identify key challenges and evaluate promising strategies for the future. An important point is the distinction between two problems: accelerated exploration and precise sampling. The former deals with the need to escape rapidly from local minima that would normally trap the system for long times. The latter consists in the extensive accumulation of samples in each relevant region of the configuration space, to achieve a precise estimation of the equilibrium probability distribution (and possibly of the kinetic properties). Some techniques are especially targeted to either of the issues, while others try addressing both at once. In general, the distinct nature of the two problems should be kept in mind, and, depending on specific system features, it can be more effective to tackle each task separately.

II.   AN OVERVIEW OF EXPLORATION STRATEGIES: ESCAPING METASTABLE STATES

Activated processes are characterized by metastable states, i.e., regions of configuration space of high probability (low free energy), separated by barriers, i.e., interface regions of very low probability (high free energy). At moderate temperature, a trajectory started in a local minimum will remain trapped into it for a very long time, before eventually

crossing a barrier and reaching another metastable state (note that the passage over the barrier is generally very rapid in itself). To give an idea, within the approximations of transition state theory[27] a free energy barrier ΔF^* affects exponentially the transition rate k according to Eyring formula:

$$k = \frac{k_B T}{h} e^{-\Delta F^*/k_B T} \tag{2}$$

At room temperature, this yields rates of the order of 10^6, 10^3 and 1 s^{-1} for barriers of the order of 15, 22, and 29 $k_B T$, respectively (or about 9, 13, 17 kcal/mol). In comparison, ab initio molecular dynamics simulations (based on density functional theory) nowadays are typically limited to the sub-nanosecond time scale, whereas classical all-atom force fields are typically limited to the microsecond time scale. Hence, barriers due to breaking of one or a few interatomic bonds slow down the kinetics of activated processes to the point that standard computer simulations become of limited use.

To allow the exploration of multiple metastable states, enhanced sampling techniques introduce smart tricks able to accelerate the escape from the local free energy minimum by several orders of magnitude at once. Most existing methods exploit the following ideas:

1. couple the trajectory at the target temperature with trajectories at higher temperatures: the latter may diffuse more easily across barriers, since, according to Boltzmann distribution, the population of energy barriers relative to minima increases exponentially with temperature.

2. add to the natural forces in the system an artificial, external biasing force, designed in such a way as to enhance the population of high free energy regions (especially barriers) with respect to its negligible equilibrium value.

3. starting from a suitable configuration, breed successive generations of trajectories applying a sort of "natural selection" where only the offspring fulfilling some requirements is kept and propagated.

While a number of related zero temperature techniques exist, aimed to explore local minima and the connecting saddle points in the potential energy surface, in this review the focus is on algorithms working with dynamical trajectories at finite temperature. In the following sections the three main strategies are discussed in more detail.

III. EXPLOITING HIGH TEMPERATURES

Since the relative probability of two microstates differing in energy by ΔE is given by Boltzmann's expression $e^{-\Delta E/k_B T}$, increasing the temperature of a simulation is a simple way to enhance the sampling of high energy regions: the latter include barrier tops, thus high temperatures can be exploited to escape from local minima in a short time.

In the simulated annealing algorithm,[28] a Monte Carlo or MD trajectory is generated starting at high temperature. In these conditions, the system is able to explore a large portion of its configuration space without being hampered by energy barriers. Then the temperature is progressively reduced, and if the cooling speed is low enough the system has a chance to relax into the global energy minimum without remaining trapped into local minima (see Fig. 1). This simple and powerful approach has been applied to a wide range of physical problems, including the determination of low energy structures of atomic clusters[29] and biological macromolecules[30], but also to many non-physical problems including computer circuit design and training artificial neural networks, as it can be seen as a general-purpose optimization tool. In this sense, it is a nice example of a physics-inspired approach borrowed to other disciplines. The efficiency of simulated annealing depends on the cooling protocol, which unfortunately needs to be tailored on the detailed topography of the energy landscape of a particular system: in practical cases, the system may remain trapped into local minima (whose number can be immense in large systems), especially in presence of high barriers.[31]

In a similar spirit, parallel tempering[32,33] connects trajectories at high temperature, expected to explore quickly, with trajectories at the lower, physically-interesting temperature. In this case, however, multiple replicas of the same system are simulated, each one at a different constant temperature. To couple together the otherwise independent simulations and allow diffusion in temperature space, at regular time intervals a swap between two replicas i and j is attempted (i.e., a swap between their temperatures) according to the following acceptance probability:

$$P = \min\left\{ 1 \ , \ \exp\left[\left(\frac{1}{k_B T_j} - \frac{1}{k_B T_i} \right) (U(x_i) - U(x_j)) \right] \right\} \tag{3}$$

where $U(x)$ is the potential energy. This algorithm, simple to implement and well-suited to parallel computers, has found many applications especially in the biophysical community (see Ref.[23] for a recent review). However, for a given physical system, it is not straightforward to predict how much the sampling can be accelerated and what is the
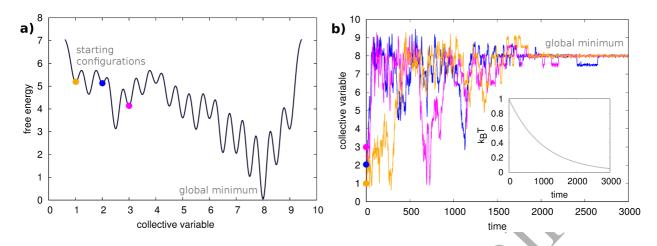
FIG. 1. Langevin simulation on a simple model illustrating the simulated annealing algorithm. a) The free energy landscape is traced as a function of a collective variable. b) Three different trajectories (orange, blue, and purple lines), started from different initial coordinates, evolve until finding the global minimum. The temperature (inset) decreases during each trajectory.

optimal setup[34]. Several investigations (see, e.g., Ref.[35,36] and references therein) led to rules of broad validity, like choosing the number of replicas so as to keep a sizable overlap between potential energy distributions, exchanging replicas very often, and going as high in temperature as possible. Note however that there is a physical limit to the highest temperature that a system can endure: e.g., barriers for bulk diffusion of atoms in crystals are easily larger than 1 eV, or equivalently $k_B \cdot 11,605$ K, much beyond the melting point of any material. In these and other cases, techniques accelerating a few relevant degrees of freedom, described in the next section, can be more useful than injecting thermal energy in all degrees of freedom at once. Many variations of parallel tempering (also called replica exchange) have been introduced to improve efficiency, too many to review here (see, e.g., Ref.[37,38]).

Other schemes have also been proposed, for instance in an interesting single-replica approach (related to the simulated tempering idea of Marinari and Parisi[39]) the simulation performs a random walk in temperature space, driven by a Langevin equation that employs as drift the deviations of the energy from its average value.[40] The algorithm was demonstrated by folding several small proteins in explicit solvent at a moderate cost, and was followed by related schemes[41,42]. In the temperature-accelerated dynamics method[43] a high-temperature simulation confined to a local minimum is exploited to identify the possible escape channels, further characterized in terms of zero-temperature energy barrier, and the information collected is exploited - under the assumptions of harmonic transition state theory - to extrapolate escape rates at a lower temperature of interest. The method has been applied especially to diffusion processes in solids.[3] Instead, in the methods called (driven) adiabatic free energy dynamics / temperature-accelerated MD[44–46] the temperature of up to several collective variables, capturing the slow degrees of freedom of the activated process, is kept higher than the temperature of the other coordinates of the system, allowing the system to overcome free energy barriers. These approaches, somehow related to those of the next section, where selected collective variables are biased with suitable artificial forces, have been applied mostly to study protein conformational dynamics.

IV. ADDING A BIAS TO THE INTERATOMIC FORCES

In metastable systems, the free energy landscape $F(s)$, defined as a function of a collective variable $s(x)$ ($x \in \mathbb{R}^{3N}$) suitably tracking transformations, features valleys much deeper than $k_B T$ separated by barriers. The very shape of the landscape suggests an ideal, simple way to enhance the sampling of the barrier regions, weakly populated at equilibrium: adding to the potential energy of the system a bias $V_B(s)$, such as to compensate free energy variations, $F(s) + V_B(s) \approx 0$. In other words, adding a biasing force $-dV_B(s)/ds$ that systematically counterbalances the average force $-dF(s)/ds$ would result in a trajectory $s(t)$ similar to a random walk with uniform probability along s (in the language of Langevin equations, the drift field of deterministic forces would be suppressed).

Of course, since in actual problems the free energy landscape is not known in advance, this procedure remains a gedankenexperiment until some prescription is given on how to build $V_B(s)$. To this aim, a number of techniques have been introduced and exploited with considerable success. Some schemes are formulated in terms of a bias potential, that can be readily applied not only to MD but also to Monte Carlo simulations, others in terms of a bias force, but

clearly the underlying idea is the same. A more important distinction is between algorithms exploiting a constant bias $V_B(s)$ or a time-dependent $V_B(s,t)$: analysing the results of the latter is typically more cumbersome since simulations are out of equilibrium. The first class includes umbrella sampling[47], hyperdynamics[48], accelerated MD[49], boxed MD[50], and more. The second class includes local elevation[51], conformational flooding[52], metadynamics[53], self-healing umbrella sampling[54], adaptive biasing force[55,56], targeted MD[57], steered MD[58], adiabatic bias MD[59,60], and more. It is important to stress that the efficiency of all these techniques depends critically upon the choice of the collective variables, as discussed in the following (sections VI and VII).

Without entering into full details, the bias potential can be constructed according to two main strategies. In the first case, a well-defined function $V_B(s)$ is chosen before the start, aimed to focus the sampling in a predefined region of configuration space. A common example is an umbrella sampling simulation where a parabolic $V_B(d_{ij}) = \frac{k}{2}(d_{ij} - d_0)^2$ forces a trajectory to keep the interatomic distance d_{ij} within the neighborhood of a fixed value d_0. Repeating similar simulations for different d_0 values allows to sample thoroughly a wide range of distances, including the breaking of chemical bonds (in a similar spirit, the "blue moon" method[61] constrains the collective variable to a fixed value). Note that reasonable starting configurations, describing a continuous pathway, need to be provided for each window: this problem is in general non-trivial, thus umbrella sampling is often used in tandem with other techniques for the discovery of transition pathways like metadynamics or steered MD. A different group of constant bias methods reduces the depth of free energy minima, thus enhancing escape rates, by systematically reducing the magnitude of interatomic forces. One can either soften selected components of a classical force field[49] (e.g., the torsional potential for protein backbones), or add "hills" of repulsive energy at pre-established minima locations[48] (e.g., adatom adsorption sites on a regular metallic surface).

In the second group of methods, the form of the bias is not guessed in advance, rather $V_B(s,t)$ is constructed on-the-fly during the simulation according to some trajectory-dependent ("history-dependent") recipe. The basic idea is that an unbiased trajectory spends more time in low free energy regions, thus the local shape of the landscape can be approximately inferred from short-time probability distributions. Iteratively, a bias is constructed that tends to flatten the free energy landscape, eventually filling the local minima until the trajectory can diffuse without barriers. This simple and powerful idea is shared by several related methods[51-55,62-65]. Among them, metadynamics[53] enjoyed a broad diffusion thanks to the simplicity of the algorithm and to the considerable amount of work that clarified theoretical foundations and convergence properties[66-70], with several effective variants introduced (notably by hybridization with replica exchange MD[71,72] and in the well-tempered form[73]). Due to the complex nature of free energy calculations, efforts to apply a technique to a broad range of different problems are extremely important to improve a method and bring it to maturity.

FIG. 2. Langevin simulation on a simple model illustrating the metadynamics algorithm. a) The free energy landscape (black line) is traced as a function of a collective variable, together with the sum of the landscape plus the bias potential at different times (colored lines). b) The trajectory of the system in the space of the collective variable is traced as a function of time.

In metadynamics, $V_B(s,t)$ is constructed as a sum of Gaussian contributions ("hills") added along the trajectory of the collective variable $s(t)$ at regular intervals.[5,13,25] The Gaussians, small both in height and width compared to the local minima, tend to fill the latter bringing the system to overcome large barriers in a short simulation time (see Fig. 2). The method is usually employed in the most effective way on a space of collective variables of low dimensionality (mostly 1D or 2D), so that the bias potential exerts its action in a precise direction. In some cases however it has been also employed in high dimensional spaces of collective variables[74,75], somehow more in the spirit of the local elevation

approach[51]. Two purposes can be achieved: accelerating the transitions between metastable states by many orders of magnitude, and reconstructing the equilibrium free energy landscape. The latter result could appear puzzling given the out-of-equilibrium character of metadynamics simulations, however the bias potential can be analyzed after a transient time (i.e., after having filled the relevant local minima) to provide an unbiased estimator of the free energy landscape, similarly to a cast shaped from a mould.[66–70] The basic rationale is the following: if stationary conditions are observed, with a uniform probability in the collective variable space, this means that the system is free to diffuse without being trapped by minima and barriers, hence the free energy landscape is compensated by the bias potential. This argument is pertinent to the original metadynamics formulation, with fixed Gaussian height W_0: in an alternative formulation, well-tempered metadynamics[73], the height is progressively reduced along the simulation according to $W = W_0 \exp[-V_B(s,t)/k_B \Delta T]$, and at convergence the bias potential stops fluctuating and provides an estimate of the free energy profile (multiplied by a factor $\Delta T/(T + \Delta T)$). Here ΔT is a parameter controlling the extent of the exploration, with $\Delta T \to 0$ corresponding to unbiased MD and $\Delta T \to \infty$ corresponding to fixed-height metadynamics. Some respective merits of the two formulations will be addressed in Section VI.

V. GENERATING ENSEMBLES OF SHORT REACTIVE TRAJECTORIES

Instead of exploiting high temperatures or deformations of the free energy landscape, a third class of algorithms focus on ways to construct and collect a large number of transition pathways connecting two predefined initial (A) and final (B) states.[7,8,21] The most widespread of these approaches is transition path sampling[2]: given an initial reactive trajectory going from A to B, a configuration along it is selected, momenta and/or positions are perturbed in a random way, and two half-trajectories are propagated (e.g., by MD), one forward and one backward in time. If the new trajectory (the union of the two halves) connects A and B, it is accepted, otherwise it is rejected, and the procedure is iterated. Starting from this simple Monte Carlo scheme in path space, a transition path ensemble can be generated that eventually contains information on the mechanism of the activated process. Each transition path is a relatively short trajectory between A and B that does not spend a long time fluctuating around minima. Note that the initial reactive pathway must be obtained by other means (e.g., applying one of the techniques in the previous sections or by interpolation), and it is not necessary to know in advance an approximate reaction coordinate (see Section VII), even if in practice an order parameter, easier to obtain, is needed to distinguish states A and B.[2] Transition path sampling was applied with success to a number of problems from materials science to chemistry to biology: examples include autoionization of liquid water[76], acid deprotonation in solution[77], nucleation and growth of nanocrystals[78], folding of small proteins[79]. This technique is often regarded as a rigorous reference approach to obtain transformation mechanisms, even if its computational cost may be conspicuous.

Several related methods, including transition interface sampling[80], forward flux sampling[81], non-equilibrium umbrella sampling[82], and milestoning[83], exploit the main ideas of transition path theory, decomposing however the region between A and B into a series of slices by means of suitable surfaces. Depending on the detailed recipe, short trajectories are typically started from one surface and stopped once they reach another surface. Often the surfaces are defined with the help of low-dimensional order parameters (collective variables). Interesting features of these approaches are the direct estimation of kinetic rates, less straightforward in transition path sampling, the access to precious information on the committor / reaction coordinate (see section VII), and the possibility to tackle also non-equilibrium processes.

VI. RECONSTRUCTING A CONVERGED FREE ENERGY LANDSCAPE

Once the relevant metastable states and transition mechanisms are explored, using for instance the approaches of the previous sections, a conceptually distinct problem is to sample extensively each explored region of the configuration space until obtaining precisely converged (ideally within $k_B T$ uncertainty) equilibrium probabilities. The latter, expressed as a function of a collective variable s, are equivalent to the free energy landscape:

$$F(s) = -k_B T \log P(s) = -k_B T \log \left(\frac{\int dx \ e^{-U(x)/k_B T} \delta(s - s(x))}{\int dx \ e^{-U(x)/k_B T}} \right) \tag{4}$$

$$\Delta F_{AB} = -k_B T \log \left(\frac{P_B}{P_A} \right) = -k_B T \log \left(\frac{\int_B dx \ e^{-U(x)/k_B T}}{\int_A dx \ e^{-U(x)/k_B T}} \right) \tag{5}$$

where $x \in \mathbb{R}^{3N}$ are the Cartesian coordinates of all atoms. Note that once the free energy landscape is known, estimating potential energy differences between selected regions yields entropy differences, otherwise difficult to access:

$$\Delta F_{AB} = \Delta E_{AB} - T\Delta S_{AB} = \langle U \rangle_B - \langle U \rangle_A - T\Delta S_{AB} \qquad (6)$$

Here, the focus is not necessarily on fast escape from local minima, but on the controlled accumulation of samples covering the relevant space. In umbrella sampling, for instance, each region of s space is sampled separately. Eventually, the biased probability distributions are first corrected mathematically to remove the bias, and then combined together in a global estimate of the free energy profile using the weighted histogram analysis method[84–89] (in a similar fashion, thermodynamic integration[90] provides a free energy profile starting from average forces estimated from independent simulations constrained at fixed values of s[61]). Umbrella sampling is a rather robust approach, widely used in different fields: it has the advantage that additional simulation can be easily added where needed, in selected windows, to systematically improve the statistics, and it is ideally suited for exploiting parallel supercomputers. On the other hand, applications are mostly limited to one- or two-dimensional free energy landscapes, since the computational cost may grow quickly with dimensionality. Moreover, the window positions and sizes initially chosen need often to be adjusted after a first round of simulations: overlap between the biased probability distributions is needed, but free energy gradients, not known in advance, can displace the distributions from the center of the windows creating gaps.
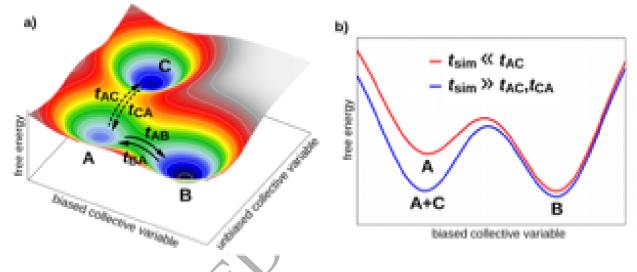

FIG. 3. a) Example of free energy landscape including three local minima A, B, and C. The biased collective variable describes well the transitions between A and B, whereas transitions between A and C are described by an orthogonal collective variable, that is not biased. t_{AB}, t_{BA}, t_{AC}, and t_{CA} indicate the mean first passage time, within a standard (unbiased) MD simulation, for the corresponding transitions. b) Free energy profiles reconstructed using a biased sampling technique (umbrella sampling, metadynamics, etc., with simulations started in A or B) in two limiting cases. If the total simulation time t_{sim} is much shorter than t_{AC}, the reconstructed profile converges to the red curve, where each minimum include microstates belonging either to A or to B (C is not explored). Conversely, if $t_{\text{sim}} \gg t_{AC}$ the simulation will explore also C, and if $t_{\text{sim}} \gg t_{CA}$ (implying multiple transitions between A and C) the reconstructed free energy profile will converge to the blue curve, where the first minimum includes the probability of both A and C. Both reconstructed profiles (red and blue) are correct, provided the correct interpretation (A, A+C, or B) is given to each minimum. Instead, in the intermediate case where t_{sim} is similar to t_{AC} the reconstruction of a truly converged free energy profile may be seriously hampered, requiring great care in the analysis of the simulation and a detailed inspection of the atomic trajectories.

Metadynamics, as discussed in Section IV, is a very successful and widespread technique that provides at the same time a fast exploration of metastable states and an estimate of the free energy landscape as function of a few collective variables. Notwithstanding a large and growing literature of successful applications, however, the first result is often easier to achieve than the second. The reason is that it is relatively easy to guess collective variables that perform well as order parameters, i.e., that distinguish the different free energy minima, whereas it is much more difficult, in many problems, to obtain good reaction coordinates, i.e., coordinates capturing the details of the transition mechanism and of the transition states (see Section VII for a precise definition). A collective variable that is a poor approximation of the true reaction coordinate may be conveniently spot out in standard metadynamics (with fixed Gaussian height) since it typically displays hysteresis in the growth of the bias potential: after having filled the main minima, instead of growing parallel to itself $V_B(s,t)$ accumulates in different regions of the landscape in an uneven fashion. Well-tempered metadynamics is less immediate from the diagnostic point of view, demanding a careful joint analysis of

the bias profile and of the collective variable trajectory [20], at the same time it presents the advantage of a simpler post-processing of bias profiles and of preventing excessive bias forces possibly occurring in long simulations when the Gaussian height is fixed. Another typical signature of non-optimal variables is the observation of a forward transition only, without a reverse one. A possibility in these cases is to stop the simulation after a first transition to estimate the height of the barrier from the sum of the Gaussians, at the risk of overestimating the barrier since the system will not necessarily pass through the optimal transition state (see Ref.[5] for more details and examples).

A specific problem that is able to affect the reconstruction of free energy landscapes in all biased sampling methods (including umbrella sampling, all variants of metadynamics, etc.) is illustrated in Fig. 3. In many practical applications, the system may overcome barriers and undergo transitions not only in the direction of the chosen biased variable s, but also along another direction. Since no bias is applied along the latter, such transitions occur just because of random fluctuations, and are likely to be observed when the corresponding mean first passage time is smaller than the total simulation time t_{sim}. Sometimes, this corresponds to undesired transformations (that can be prevented imposing a suitable repulsive potential, or wall), and sometimes to an interesting unforeseen process that leads to rethink the study. When the extra transitions are fast compared to t_{sim}, equilibrium is achieved and the corresponding extra metastable states are incorporated into the reconstructed free energy profile (see Fig. 3). However, when the transition time is similar to t_{sim}, the behavior of the simulation is erratic and out of control, and it may result impossible to reconstruct a truly converged free energy landscape within the available t_{sim} (irrespective of the observation of an apparent convergence of the bias profile, as may happen, e.g., in well-tempered metadynamics)[5,13,20,22,91–93]. The last statement applies to all free energy reconstruction techniques, it is a logical consequence of hidden orthogonal barriers combined with finite simulation time (a well-known fact to practitioners in the field), and it does not contradict mathematical proofs of convergence like in Ref.[67,69,94,95], which have a formal rather than practical scope.

It is important to realize that the diagnosis of this kind of problems is not always a trivial matter. As a general rule, it is of paramount importance to perform a very careful inspection of the atomistic trajectories to check if an ergodic behavior is attained within the ensemble of configurations explored.

Different strategies can be adopted to cure these problems, as discussed in Section VIII: heuristically improving the design of the collective variables, extending their number employing bias exchange metadynamics[15,72], combining metadynamics with parallel tempering[71], extracting the relevant degrees of freedom from a large set by means of path collective variables[96–99] or sketch maps[100,101], employing on-the-fly clustering of configurations[74], resorting to transition path sampling and related techniques to gain insight on the transformation mechanism, and splitting the exploration problem from the free energy calculation one.

Clearly, the choice of the collective variables is critical in all techniques that aim to reconstruct free energy landscapes: Given the number of different available techniques, and given the pitfalls so common in this field, it is insightful to compare different methods on the same problem, and it is important to analyse the landscapes and transition mechanisms with diagnostic tools, e.g., committor analysis (Section VII). Considering the complexity of the structure and dynamics of many condensed matter systems, it is reasonable to accept that the convergence of any free energy calculation method cannot be guaranteed in all circumstances but critically depends on the choice of parameters and collective variables, and on the amount of computational resources.

VII. A CENTRAL CONCEPT: THE REACTION COORDINATE

A key ingredient in the study of activated processes is the reaction coordinate. In a one-step process passing from an initial metastable state A to a final one B (and vice-versa), it can be defined as the committor function $P_B(x)$ (or equivalently $P_A(x)$) that associates to each configuration x in the full $3N$-dimensional space of the system the probability to reach B (or equivalently A).[2,19,102–105]. Such probability can be computed generating a large number of trajectories starting from x and counting how many reach B before reaching A. Typically, trajectories differ from each other because of different initial velocities, randomly drawn from a canonical distribution, or, in the case of overdamped Langevin dynamics, because of a different sequence of random numbers in the noise term.

The committor function varies smoothly from the value zero in the initial state to the value one in the final state, assuming a value of 0.5 at the transition state. Therefore it is a suitable indicator of the evolution of the transformation, including kinetic information about the fate of configurations, hence a good definition of reaction coordinate. The applicability is general, going from $T > 0$ down to $T = 0$ (where it reduces to the intrinsic reaction coordinate along the gradient of potential energy[106,107]), encompassing barriers of energetic or of entropic origin, purely diffusive processes, etc.

Unfortunately, besides its conceptual relevance, the committor cannot in general be explicitly computed for all possible configurations x, as their number is huge in non-trivial systems. Moreover, in complex systems it may be hard to visualize and some more approximate definition of reaction coordinate, with a simple analytical formulation, can be helpful for human insight.[105] However, a committor analysis can still be a very useful a-posteriori tool, after

having sampled some transition pathways, restricting the initial configurations x to those already sampled. In this way, for instance, a configuration in-between the transition pathway can be identified as a transition state if trajectories generated from it are split in two halves, 50% falling to A and 50% to B. Note that the computational cost is here limited whenever the number of putative transition state structures is limited: in a practical algorithm, from each of the latter structures a few tens of trajectories can be initially generated, each one relatively short as it needs only to relax to a nearby free energy minimum. Only if trajectories do not all fall in the same minimum, the initial structure is worth a more extended investigation shooting hundreds more trajectories until a precise estimate of $P_B(x)$.

Another application of committor analysis is the validation of a putative definition of reaction coordinate $s(x)$: after reconstructing the free energy profile $F(s) = -k_B T \log(P(s))$ for transitions between A and B, configurations can be drawn from the vicinity of the saddle point, to be tested as putative transition states. A lack of halfway splitting of the committor probability should be interpreted as an indication that $s(x)$ is a poor approximation of the ideal reaction coordinate.[2] A paradigmatic example is the ion-pair distance as coordinate to study NaCl dissociation in aqueous solution.[102,108–110] Committor analysis is increasingly employed not only in classical but also in ab initio MD simulations, for transition state- and reaction coordinate validation[97,98,111,112].

Several algorithms have been designed to identify optimal definitions of reaction coordinates for specific processes, usually requiring extensive simulations[19,104,113–122]. Knowledge, even approximate, of the reaction coordinate is extremely useful to enhance the sampling of the activated process, using the techniques described in previous sections. Despite the differences among biasing algorithms like umbrella sampling, metadynamics, adaptive biasing force, temperature accelerated MD, steered MD etc., they all are likely to be effective if the reaction coordinate is known. Conversely, a poor definition of reaction coordinate is likely to seriously hamper all of those algorithms, with the risk of observing non-optimal transition mechanisms (see Section VI). Identifying the reaction coordinate for a given process is not only a technical step in enhanced sampling, rather it corresponds with a deep understanding of the nature and physical driving forces of a transformation process, and as such it is generally desirable even when enhanced sampling is not needed. Considering increasing computer power, growing system sizes and thus growing amounts of simulation data, efficient and automated algorithms distilling reaction coordinates are needed, even if the result may not always easily translate into human-readable insight.

Given the central role of reaction coordinates both for understanding and effective simulation of activated processes, the next section is devoted to discuss some insightful case studies.

VIII. CASE STUDIES

In practice, detailed information on transition pathways is usually unavailable at the early stage of a study, therefore approximate definitions of reaction coordinates are often guessed based on heuristic arguments, drawing on literature, subjective experience and intuition. Albeit non-rigorous, such approach is often successful in practice when the results of test simulations are employed to critically revise the original guess in an iterative optimization process, with a sizable cost in both computer and human resources. Collective variables are selected based on the particular features of the system under study. Some widespread basic ingredients, used alone or in simple algebraic combinations, are: distances, angles, and coordination numbers in chemistry and nanoscience; local-density or local-symmetry indicators in phase transitions; dihedral angles, gyration radii, contact maps, normal modes, root mean square deviations, and principal-component-analysis eigenvectors in biomolecules.

A. Multidimensional spaces of collective variables

In some complicated cases, notably those involving the conformational dynamics of proteins or the search of low-energy structures in atomic nanoclusters, the existence of multiple ($\gg 2$) metastable states of interest suggests the opportunity to explore a space of several collective variables, in the attempt to capture the slow degrees of freedom contributing to reaction coordinates. For instance, it has been recently demonstrated how already a small 9-residue peptide in solution exhibits a complex multistate dynamics with a hierarchy of time scales, requiring five collective variables to be adequately described.[123] Among enhanced sampling methods able to explore high-dimensional collective variable spaces, it was recently shown[124] that the temperature-accelerated techniques of Ref.[44–46] can be effectively exploited to estimate free energy differences between landmark points (minima and saddles) in complex landscapes. In this context, it is important to realize that dynamical trajectories of complicated systems tend to fill a tiny fraction (the low-free energy part) of the mathematically conceivable configuration space, namely a network formed by the bottom of minima and their interconnections[125]. Therefore efficient enhanced sampling strategies must focus on the exploration of the relevant part only of high-dimensional landscapes. Two more examples of successful strategies are the following.
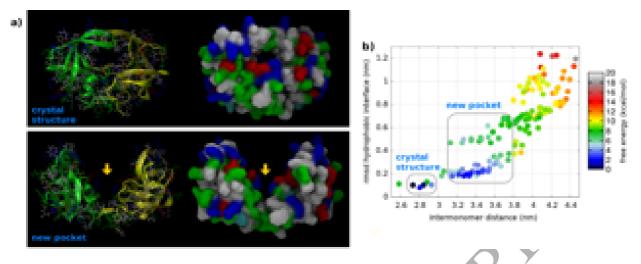
FIG. 4. Bias exchange metadynamics simulations (all atoms, explicit solvent) of HIV-1 protease dimerization, performed with four replicas biased along different collective variables, led to the discovery of a new druggable binding pocket (arrow in panel a)). The pocket opens transiently at the interface of the dimer and corresponds to a metastable state. In panel b), the free energy landscape projected onto two collective variables is shown in the form of structural clusters. See Ref.[126] for details.

In bias exchange simulations[15,72,127], several metadynamics simulations are run in parallel, on the same system at the same temperature. Each replica is biased along a different collective variable, and at fixed time intervals the replicas are allowed to exchange coordinates based on a Metropolis criterion with probability

$$P_{ab} = \min \left\{ 1 \ , \ \exp \left[\left(V_B^a(x^a, t) + V_B^b(x^b, t) - V_B^a(x^b, t) - V_B^b(x^a, t) \right) / k_B T \right] \right\} \tag{7}$$

where a, b label replicas and V_B denotes the bias potential of metadynamics. In this simple scheme each one-dimensional metadynamics simulation is able to quickly overcome barriers, while frequent exchanges (typically every ~ 1 ps) allow exploring in an efficient way a space of several (typically less than 10, but also more) collective variables. When stationary conditions are achieved, e.g., parallel growth of each bias potential profile, the one-dimensional free energy profiles are obtained along the different collective variables. Remarkably, application of the weighted-histogram analysis technique in Ref.[128] delivers also the multidimensional free energy landscape as a function of all the collective variables.

These approaches, in combination with accurate all-atom explicit-solvent force fields, have been applied to a number of challenging problems, including the folding of globular proteins,[128–131] the dynamics of intrinsically disordered proteins,[132–134] or protein-ligand and protein-protein association processes[135–138], and so on. Recently a study of the dimerization mechanism of HIV protease, a key enzyme in the viral life-cycle and target of anti-AIDS therapies, led to the discovery of a new druggable binding pocket appearing in correspondence of a previously unknown metastable state along the dimerization pathway, opening new possibilities for drug design[126] (see Figure 4). Effective definitions of collective variables include the radius of gyration, the number of contacts among polar or hydrophobic residues, the amount of alpha or beta (this latter being the most challenging to produce) secondary structure[139], the root mean square deviation of protein regions from reference structures[126], or, increasingly acknowledged as an important slow degree of freedom, the number of water molecules in the core of the protein or at the interface between biomolecules.[130,135] Assuming that the dynamics of the system in the multidimensional space of the collective variables can be approximated by a Smoluchowski-type equation, discretization of the space on a regular grid provides an explicit expression for the rates between microstates[140,141]. In this way, a Markovian kinetic model can be built from bias exchange simulations[128], allowing to estimate the network of kinetic rates among metastable states, including folding rates of proteins and on/off rates of ligand binding, to be directly compared with experiments[128,135].

A different strategy to sample a high-dimensional collective variables space is given by social permutation-invariant topological (SPRINT) coordinates exploited in combination with standard metadynamics[75]. The main idea is that the atoms of a nanocluster (or any other condensed matter system) form a network of bonds (covalent, ionic, hydrogen bonds, etc., depending on the system). The exploration of different structures can be seen as an exploration of different graph topologies, tracked using tools borrowed from graph theory. Starting from the $N \times N$ adjacency matrix a_{ij} of the graph formed by interatomic bonds, defined by means of smooth functions of the distance R_{ij} (e.g., $a_{ij} = 1/[1 + \exp((R_{ij} - \rho)/\lambda)])$, a space of N collective variables is defined as shown in Fig. 5, i.e., multiplying together the largest eigenvalue $\lambda^{\max}$ (isolated for connected graphs by virtue of Perron-Frobenius theorem) and the

corresponding eigenvector $\mathbf{v}^{\max}$. The former increases with the number of bonds in the system, while the latter can be expressed as a sum over all possible walks in the graph (it is interesting to note the analogies with the PageRank algorithm behind the search engine of Google[142]). The resulting collective variables are sensitive to changes of bond topology well beyond the first coordination shell. After enforcing permutation invariance upon exchange of identical atoms (a crucial property for efficient sampling[143]) by simply sorting the variables in ascending order, a free exploration of possible topologies is performed starting from an initial geometry and biasing all the variables with a technique like metadynamics. If, on one side, the high-dimensional character of the space prevents a straightforward reconstruction of the free energy landscape, on the other side the strong correlation between s_i variables (partly due to the sorting operation) allows efficient exploration even within expensive ab initio MD. Examples are silicon nanoclusters, isomerization/association/dissociation reactions of organic molecules[75], graphene flakes transforming into cages[144] (see Fig. 5), boron impurities in carbonates[145]. Starting from the observed transition pathways, the corresponding free energy profiles can be reliably and systematically reconstructed by means of path collective variables[96,144] (see next section).
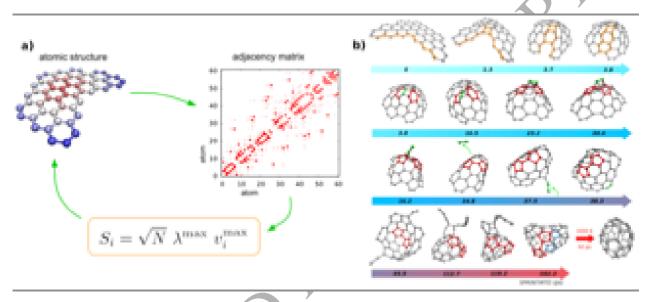

FIG. 5. a) Definition of social permutation invariant (SPRINT) coordinates based the principal eigenvalue and eigenvector of the adjacency matrix of chemical bonds.[75] For illustration, a graphene flake is colored according to the coordinates (from the most network-central atoms, in red, to the most peripheral ones, in blue). b) Complex transformation mechanisms of a graphene flake are simulated with ab initio MD applying metadynamics on the space of SPRINT coordinates. Eventually, a closed cage is obtained (see Ref.[144] for details).

A related approach, again starting from the adjacency matrix of the network but including all matrix elements, is exploited in the definition of Ref.[146], that achieved the challenging task of resolving a number of different crystalline and amorphous structures of water[147]. Interestingly, the resulting two-dimensional diagram of topological distances reproduces important features of the experimental phase diagram (for effective dimensional reduction of high dimensional spaces, see also Ref.[101]). In combination with path collective variables (discussed in the next section), such topological distances form the basis of a flexible enhanced sampling strategy adapted to phase transitions: based on a realistic interatomic potential for water (TIP4P/2005[148]) and on NPT metadynamics and umbrella sampling simulations, the new approach allowed to navigate throughout the phase diagram of water, simulating transitions among liquid, amorphous and crystal structures, including the very challenging case of ice crystallization from the liquid without cooling below the melting temperature.[99] In this field, the main collective variables employed so far were the edges of the simulation cell[149,150] or atom-centered symmetry indicators[151–154]. New approaches based on the topology of the whole interatomic network, like the one of Ref.[99], could open the way to simulations of phase transitions and nucleation in a range of different systems without fine tuning.

B. Path collective variables

A very interesting compromise between the necessity to include many relevant degrees of freedom and the convenience of sampling a low-dimensional space are path collective variables.[96] Mostly employed in biophysics, they

turn out to be exceptionally useful in simulations of chemical reactions and phase transitions. The basic idea is to start from a sequence of n reference configurations $\{x_i\}_{i=1,\dots,n}$ (where x indicates the $3N$ Cartesian coordinates) representing the initial state A, some intermediate geometries, and the final state B. Then, two collective coordinates are defined, s representing the progress along the putative pathway, and z representing the distance from it:

$$s = \frac{\sum_{i=1}^n i \; e^{-\lambda D(x(t),x_i)}}{\sum_{j=1}^n e^{-\lambda D(x(t),x_j)}} \tag{8}$$

$$z = -\frac{1}{\lambda} \log \left[\sum_{j=1}^n e^{-\lambda D(x(t),x_j)} \right] \tag{9}$$

$$\tag{10}$$

where D is a suitable distance metric. In practice, it is convenient to employ equidistant reference structures $(D(x_1, x_2) \approx D(x_2, x_3) \approx \dots \approx D(x_{n-1}, x_n))$ and a parameter λ such that $\lambda D(x_i, x_{i+1})$ is the order of unity. The two main ingredients are the reference structures and the metric D. The former correspond to a putative pathway, an initial guess that can be obtained from a preliminary exploration by other means, or by some sort of interpolation between A and B. Crucially, a bad guess at this point is less likely to trap the simulations in a wrong mechanism, compared to schemes like the nudged elastic band or transition path sampling, thanks to the coordinate z that, when sampled for larger and larger values, is able to track pathways sizably different from the guess. The approach (sharing similarities with the string method[155,156]) can be applied to very different systems provided suitable D are identified.

FIG. 6. a) Schematic illustration of path coordinates tracking transformation mechanisms between reactants and products (here formamide and carbon monoxide plus ammonia) described as patterns of coordination numbers. Only information about the initial and final bond topology is included in the coordinates, without any guess about the mechanism. Free energy landscapes for b) gas phase and c) solution reactions as a function of path collective variables. In solution, pathways to formic acid are automatically discovered despite of the lack of any information about it in the coordinates (see Ref.[98] for details).
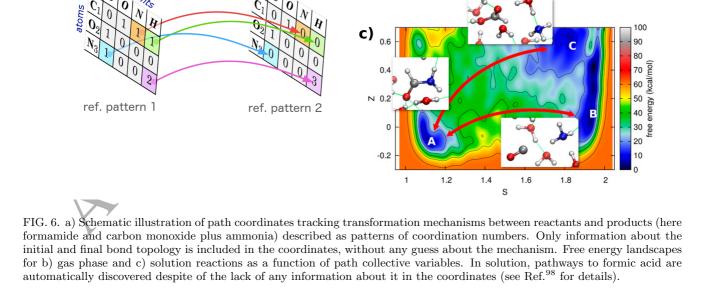
The method has been applied with success to study a number of biophysical processes, with D typically corresponding to root mean square deviations of Cartesian coordinates of selected atoms of the biomolecule. More recently, the collective variables were demonstrated to be extremely effective in the study of chemical reactions within gas phase ab

initio MD simulations combined with umbrella sampling or metadynamics[97,144,157]. Traditionally, it is very difficult to find good collective variables for concerted reactions, since simple heuristic combinations of distances, angles, etc. hardly capture the detailed pattern of simultaneous breaking/formation of several bonds. Path collective variables offer an easy and robust solution, as demonstrated by fully reversible reactive pathways, with saddle points in the free energy landscape confirmed as transition states by committor analysis[97].

On the other hand, extension to reactions in solution, however desirable, is not obvious. In this case, a suitable metric D must account for the active participation of solvent molecules, not only as a bath modifying the relative stability of solute molecules but also, e.g., exchanging protons or hydroxide anions with the solute in acid/base and hydrolysis/condensation reactions. Due to the diffusion of solvent molecules, the metric should also be invariant under permutation of identical solvent atoms and molecules. A possible solution is to employ patterns of coordination numbers between a selected set of solute atoms and all atoms available in the solute or solvent[98] (see Figure 6). In this way, high-dimensional information about complex patterns of coordination are accounted for through D (which is simply the distance between pattern tables), whereas dimensional reduction to just two variables (s and z) allows an efficient sampling and reconstruction of the free energy landscape. Application to paradigmatic prebiotic reactions (interconverting formamide, hydrogen cyanide, formic acid, etc.)[98,158] as well as to disproportionation of methanol[159] allowed a straightforward reconstruction of the reaction networks and associated free energy profiles, within a unified approach to gas phase and solutions. A remarkable result is that including only the initial A and final B state in the construction of s and z, without any guess about the pathway, is sufficient to automatically discover reaction mechanisms, including unexpected intermediates and products. Since only knowledge of reactants and (tentative) products is required, this approach greatly simplifies applications, otherwise restricted to heuristic variables and a slow trial-and-error procedure, opening the way to the systematic study of a range of reactions at realistic conditions.

IX. CHALLENGES TO BE FACED

The techniques presented in this review alleviate to a large extent the time scale limitations of brute force MD simulations, and yielded qualitative and quantitative insight into many different transformation processes. At the same time, there is an active struggle to make methodological improvements, since important challenges remain to be tackled. Condensed matter features a vast range of systems and problems, spanning materials science, chemistry, and biophysics: despite intriguing similarities, crucial specific features differ between systems and must be taken into account in the application of enhanced sampling methods. This situation adversely affects portability, and usually the study of a new specific system entails a large investment of expert human work to fine tune the algorithms (it can be argued that nowadays the human cost is often the main limitation of these techniques, rather than computer resources). An emblematic example is the quest for reaction coordinates, often performed heuristically by trial and error: starting from a preliminary idea of the problem, metastable states and transformation pathways are explored, and if the outcome is not satisfactory, or if unexpected states and mechanisms are discovered, the simulations are redesigned and repeated. Formulations of reaction coordinates able to cover a wide spectrum of systems in a unified way are therefore a key step in the direction of easier, more reliable, and more reproducible simulations. Similarly, software tools that allow to directly compare the performance of different enhanced sampling approaches go in the same direction[160–162]. The stakes are high: today a chemical reaction in solution including hundreds of atoms, an example of challenging problem, can be simulated at ab initio MD level within one week, and in few years it will take about one day or less. This means that the power of modern computer facilities, once combined with reliable sampling algorithms yielding reproducible results, has the potential to boost the systematic prediction of experimental results in the near future.

The development of enhanced sampling would also greatly benefit from the introduction of agreed benchmark systems of higher complexity than those customarily used (e.g., alanine dipeptide or small Lennard-Jones clusters), better approximating the rich behavior observed in realistic systems of interest. A significant example is the set of extensive MD trajectories of proteins generated with the Anton machine[163] that helped reviving the race to efficient sampling methods for protein folding. If several groups attempted to critically assess and compare the performance of different algorithms[130,164–168], a non-trivial task, still there is ample room for improvements. Finally, an important challenge for the future is the discovery of systematic approaches to tackle the kinetics of complex systems in an affordable and accurate way. This is a large field of research on its own, until now quite fragmented into different communities focused on physics, chemistry or biology. Two examples of problems of strategic importance, today unsolved, where the kinetics versus thermodynamics competition plays a central role are the prediction of viable routes for the synthesis of new materials with desired structure, and the prediction of interactomes, i.e., the network of protein-protein interactions in living organisms. Theories of the kinetics of rare events lay at the foundation of all

attempts to explore free energy landscapes.

* fabio.pietrucci@impmc.upmc.fr

[1] P. Kollman, Chem. Rev. **93**, 2395 (1993).

[2] P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler, Annu. Rev. Phys. Chem. **53**, 291 (2002).

[3] A. F. Voter, F. Montalenti, and T. C. Germann, Annu. Rev. Mater. Res. **32**, 321 (2002).

[4] C. Chipot and A. Pohorille, *Free energy calculations* (Springer-Verlag Berlin Heidelberg, 2007).

[5] A. Laio and F. L. Gervasio, Rep. Progr. Phys. **71**, 126601 (2008).

[6] C. Dellago and P. G. Bolhuis, *Transition path sampling and other advanced simulation techniques for rare events* (Springer Berlin Heidelberg, 2008).

[7] R. J. Allen, C. Valeriani, and P. R. ten Wolde, J. Phys.: Condens. Matter **21**, 463102 (2009).

[8] P. G. Bolhuis and C. Dellago, Rev. Comput. Chem. **27**, 111 (2010).

[9] V. Leone, F. Marinelli, P. Carloni, and M. Parrinello, Curr. Opin. Struct. Biol. **20**, 148 (2010).

[10] C. D. Christ, A. E. Mark, and W. F. Van Gunsteren, J. Comput. Chem. **31**, 1569 (2010).

[11] T. Lelièvre, G. Stoltz, and M. Rousset, *Free energy computations: A mathematical perspective* (World Scientific, 2010).

[12] A. Pohorille, C. Jarzynski, and C. Chipot, J. Phys. Chem. B **114**, 10235 (2010).

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

[14] X. Biarnés, S. Bongarzone, A. V. Vargiu, P. Carloni, and P. Ruggerone, J. Comput. Aid. Mol. Des. **25**, 395 (2011).

[15] F. Baftizadeh, P. Cossio, F. Pietrucci, and A. Laio, Curr. Phys. Chem. **2**, 79 (2012).

[16] S. Bonella, S. Meloni, and G. Ciccotti, Eur. Phys. J. B **85**, 97 (2012).

[17] L. Sutto, S. Marsili, and F. L. Gervasio, WIREs Comput. Mol. Sci. **2**, 771 (2012).

[18] C. Abrams and G. Bussi, Entropy **16**, 163 (2013).

[19] M. A. Rohrdanz, W. Zheng, and C. Clementi, Annu. Rev. Phys. Chem. **64**, 295 (2013).

[20] G. Bussi and D. Branduardi, Rev. Comput. Chem. **28**, 1 (2015).

[21] P. Bolhuis and C. Dellago, Eur. Phys. J. Special Topics **224**, 2409 (2015).

[22] J. Comer, J. C. Gumbart, J. Henin, T. Lelievre, A. Pohorille, and C. Chipot, J. Phys. Chem. B **119**, 1129 (2015).

[23] M. Luitz, R. Bomblies, K. Ostermeir, and M. Zacharias, J. Phys.: Condens. Matter **27**, 323101 (2015).

[24] R. Elber, J. Chem. Phys. **144**, 060901 (2016), http://dx.doi.org/10.1063/1.4940794.

[25] O. Valsson, P. Tiwary, and M. Parrinello, Annu. Rev. Phys. Chem. **67** (2016).

[26] A. P. Usher, *A History of Mechanical Inventions: Revised Edition* (Courier Corporation, 2013).

[27] P. Hänggi, P. Talkner, and M. Borkovec, Rev. Mod. Phys. **62**, 251 (1990).

[28] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, Science **220**, 671 (1983).

[29] W. Andreoni, Z. Phys. D: At., Mol. Clusters **19**, 31 (1991).

[30] A. T. Brünger, P. D. Adams, and L. M. Rice, Structure **5**, 325 (1997).

[31] M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos, Rev. Mod. Phys. **64**, 1045 (1992).

[32] U. H. Hansmann, Chem. Phys. Lett. **281**, 140 (1997).

[33] Y. Sugita and Y. Okamoto, Chem. Phys. Lett. **314**, 141 (1999).

[34] D. M. Zuckerman and E. Lyman, J. Chem. Theory Comput. **2**, 1200 (2006).

[35] E. Rosta and G. Hummer, J. Chem. Phys. **131**, 165102 (2009).

[36] T.-Q. Yu, J. Lu, C. F. Abrams, and E. Vanden-Eijnden, Proc. Natl. Acad. Sci. U.S.A. **113**, 11744 (2016).

[37] D. J. Earl and M. W. Deem, Phys. Chem. Chem. Phys. **7**, 3910 (2005).

[38] P. Liu, B. Kim, R. A. Friesner, and B. Berne, Proc. Natl. Acad. Sci. U.S.A. **102**, 13749 (2005).

[39] E. Marinari and G. Parisi, Europhys. Lett. **19**, 451 (1992).

[40] C. Zhang and J. Ma, J. Chem. Phys. **132**, 244101 (2010).

[41] G. Gobbo and B. J. Leimkuhler, Phys. Rev. E **91**, 061301 (2015).

[42] N. Lenner and G. Mathias, J. Chem. Theory Comput. **12**, 486 (2016).

[43] M. R. So/rensen and A. F. Voter, J. Chem. Phys. **112**, 9599 (2000).

[44] L. Rosso, P. Minry, Z. Zhu, and M. E. Tuckerman, J. Chem. Phys. **116**, 4389 (2002).

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

[46] J. B. Abrams and M. E. Tuckerman, J. Phys. Chem. B **112**, 15742 (2008).

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

[48] A. F. Voter, J. Chem. Phys. **106**, 4665 (1997).

[49] D. Hamelberg, J. Mongan, and J. A. McCammon, J. Chem. Phys. **120**, 11919 (2004).

[50] D. R. Glowacki, E. Paci, and D. V. Shalashilin, J. Phys. Chem. B **113**, 16603 (2009).

[51] T. Huber, A. E. Torda, and W. F. van Gunsteren, J. Comput. Aid. Mol. Des. **8**, 695 (1994).

[52] H. Grubmüller, Phys. Rev. E **52**, 2893 (1995).

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

[54] S. Marsili, A. Barducci, R. Chelli, P. Procacci, and V. Schettino, J. Phys. Chem. B **110**, 14011 (2006).

[55] E. Darve, D. Rodríguez-Gómez, and A. Pohorille, J. Chem. Phys. **128**, 144120 (2008).

[56] J. Henin, G. Fiorin, C. Chipot, and M. Klein, J. Chem. Theory Comput. **6**, 35 (2010).

[57] J. Schlitter, M. Engels, and P. Krueger, J. Mol. Graph. **12**, 84 (1994).

[58] S. Izrailev, S. Stepaniants, B. Isralewitz, D. Kosztin, H. Lu, F. Molnar, W. Wriggers, and K. Schulten, "Steered molecular dynamics," in *Computational Molecular Dynamics: Challenges, Methods, Ideas: Proceedings of the 2nd International Symposium on Algorithms for Macromolecular Modelling, Berlin, May 21–24, 1997* (Springer Berlin Heidelberg, Berlin, Heidelberg, 1999) pp. 39–65.

[59] M. Marchi and P. Ballone, J. Chem. Phys. **110**, 3697 (1999).

[60] E. Paci and M. Karplus, J. Mol. Biol. **288**, 441 (1999).

[61] E. Carter, G. Ciccotti, J. T. Hynes, and R. Kapral, Chem. Phys. Lett. **156**, 472 (1989).

[62] D. Cvijovic and J. Klinowski, Science **267**, 664 (1995).

[63] X. Wu and S. Wang, J. Chem. Phys. **110**, 9401 (1999).

[64] F. Marinelli, Biophys. J. **105**, 1236 (2013).

[65] O. Valsson and M. Parrinello, Phys. Rev. Lett. **113**, 090601 (2014).

[66] A. Laio, A. Rodriguez-Fortea, F. L. Gervasio, M. Ceccarelli, and M. Parrinello, J. Phys. Chem. B **109**, 6714 (2005).

[67] G. Bussi, A. Laio, and M. Parrinello, Phys. Rev. Lett. **96**, 090601 (2006).

[68] Y. Crespo, F. Marinelli, F. Pietrucci, and A. Laio, Phys. Rev. E **81**, 055701 (2010).

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

[70] P. Tiwary and M. Parrinello, J. Phys. Chem. B **119**, 736 (2015).

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

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

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

[74] G. A. Tribello, M. Ceriotti, and M. Parrinello, Proc. Natl. Acad. Sci. U.S.A. **107**, 17509 (2010).

[75] F. Pietrucci and W. Andreoni, Phys. Rev. Lett. **107**, 085504 (2011).

[76] P. L. Geissler, C. Dellago, D. Chandler, J. Hutter, and M. Parrinello, Science **291**, 2121 (2001).

[77] J. M. Park, A. Laio, M. Iannuzzi, and M. Parrinello, J. Am. Chem. Soc. **128**, 11318 (2006).

[78] M. Grünwald and C. Dellago, Nano Lett. **9**, 2099 (2009).

[79] P. G. Bolhuis, Proc. Natl. Acad. Sci. U.S.A. **100**, 12129 (2003).

[80] T. S. Van Erp and P. G. Bolhuis, J. Comput. Phys. **205**, 157 (2005).

[81] R. J. Allen, P. B. Warren, and P. R. Ten Wolde, Phys. Rev. Lett. **94**, 018104 (2005).

[82] A. Dickson and A. R. Dinner, Annu. Rev. Phys. Chem. **61**, 441 (2010).

[83] A. K. Faradjian and R. Elber, J. Chem. Phys. **120**, 10880 (2004).

[84] S. Kumar, J. M. Rosenberg, D. Bouzida, R. H. Swendsen, and P. A. Kollman, J. Comput. Chem. **13**, 1011 (1992).

[85] B. Roux, Computer Phys. Commun. **91**, 275 (1995).

[86] F. Zhu and G. Hummer, J. Comput. Chem. **33**, 453 (2012).

[87] E. Rosta and G. Hummer, J. Chem. Theory Comput. **11**, 276 (2014).

[88] H. Wu, F. Paul, C. Wehmeyer, and F. Noé, Proc. Natl. Acad. Sci. U.S.A. , 201525092 (2016).

[89] E. Thiede, B. Van Koten, J. Weare, and A. R. Dinner, arXiv:1603.04505 (2016).

[90] M. Sprik and G. Ciccotti, J. Chem. Phys. **109**, 7737 (1998).

[91] D. Branduardi, F. L. Gervasio, A. Cavalli, M. Recanatini, and M. Parrinello, J. Am. Chem. Soc. **127**, 9147 (2005).

[92] L. Zheng, M. Chen, and W. Yang, Proc. Natl. Acad. Sci. U.S.A. **105**, 20227 (2008).

[93] J. F. Dama, *Contributions towards the theory of the bottom-up coarse-graining of complex molecules*, Ph.D. thesis, The University of Chicago (2016).

[94] T. Lelievre, M. Rousset, and G. Stoltz, Nonlinearity **21**, 1155 (2008).

[95] P. Tiwary, J. F. Dama, and M. Parrinello, J. Chem. Phys. **143**, 234112 (2015).

[96] D. Branduardi, F. L. Gervasio, and M. Parrinello, J. Chem. Phys. **126**, 054103 (2007).

[97] G. A. Gallet, F. Pietrucci, and W. Andreoni, J. Chem. Theory Comput. **8**, 4029 (2012).

[98] F. Pietrucci and A. M. Saitta, Proc. Natl. Acad. Sci. U.S.A. **112**, 15030 (2015).

[99] S. Pipolo, M. Salanne, G. Ferlat, S. Klotz, A. M. Saitta, and F. Pietrucci, arXiv:1703.00753 (2017).

[100] M. Ceriotti, G. A. Tribello, and M. Parrinello, Proc. Natl. Acad. Sci. U.S.A. **108**, 13023 (2011).

[101] G. A. Tribello, M. Ceriotti, and M. Parrinello, Proc. Natl. Acad. Sci. U.S.A. **109**, 5196 (2012).

[102] P. L. Geissler, C. Dellago, and D. Chandler, J. Phys. Chem. B **103**, 3706 (1999).

[103] A. Berezhkovskii and A. Szabo, J. Chem. Phys. **122**, 014503 (2005).

[104] R. B. Best and G. Hummer, Proc. Natl. Acad. Sci. U.S.A. **102**, 6732 (2005).

[105] B. Peters, Annu. Rev. Phys. Chem. **67**, 669 (2016).

[106] K. Fukui, J. Phys. Chem. **74**, 4161 (1970).

[107] K. Ishida, K. Morokuma, and A. Komornicki, J. Chem. Phys. **66**, 2153 (1977).

[108] A. J. Ballard and C. Dellago, J. Phys. Chem. B **116**, 13490 (2012).

[109] R. G. Mullen, J.-E. Shea, and B. Peters, J. Chem. Theory Comput. **10**, 659 (2014).

[110] Y. Yonetani, J. Chem. Phys. **143**, 044506 (2015).

[111] M. Dametto, D. Antoniou, and S. Schwartz, Mol. Phys. **110**, 531 (2012).

[112] D. Munoz-Santiburcio, M. Kosa, A. Hernandez-Laguna, C. I. Sainz-Diaz, and M. Parrinello, J. Phys. Chem. C **116**, 12203 (2012).

[113] A. Ma and A. R. Dinner, J. Phys. Chem. B **109**, 6769 (2005).

[114] W. E, W. Ren, and E. Vanden-Eijnden, Chem. Phys. Lett. **413**, 242 (2005).

[115] B. Peters and B. L. Trout, J. Chem. Phys. **125**, 054108 (2006).

116 B. Peters, G. T. Beckham, and B. L. Trout, J. Chem. Phys. **127**, 034109 (2007).
117 W. Lechner, J. Rogal, J. Juraszek, B. Ensing, and P. G. Bolhuis, J. Chem. Phys. **133**, 174110 (2010).
118 B. Peters, P. G. Bolhuis, R. G. Mullen, and J.-E. Shea, J. Chem. Phys. **138**, 054106 (2013).
119 L. V. Nedialkova, M. A. Amat, I. G. Kevrekidis, and G. Hummer, J. Chem. Phys. **141**, 114102 (2014).
120 P. Tiwary and B. J. Berne, Proc. Natl. Acad. Sci. U.S.A. **113**, 2839 (2016).
121 T. S. van Erp, M. Moqadam, E. Riccardi, and A. Lervik, J. Chem. Theory Comput. **12**, 5398 (2016).
122 C. J. Dsilva, R. Talmon, C. W. Gear, R. R. Coifman, and I. G. Kevrekidis, SIAM J. Appl. Dyn. Syst. **15**, 1327 (2016).
123 N. Schaudinnus, B. Bastian, R. Hegger, and G. Stock, Phys. Rev. Lett. **115**, 050602 (2015).
124 M. Chen, T.-Q. Yu, and M. E. Tuckerman, Proc. Natl. Acad. Sci. U.S.A. **112**, 3235 (2015).
125 D. J. Wales, M. A. Miller, and T. R. Walsh, Nature **394**, 758 (1998).
126 F. Pietrucci, A. V. Vargiu, and A. Kranjc, Sci. Rep. **5** (2015).
127 P. Cossio, F. Marinelli, A. Laio, and F. Pietrucci, J. Phys. Chem. B **114**, 3259 (2010).
128 F. Marinelli, F. Pietrucci, A. Laio, and S. Piana, PLoS Comput. Biol. **5**, e1000452 (2009).
129 N. Todorova, F. Marinelli, S. Piana, and I. Yarovsky, J. Phys. Chem. B **113**, 3556 (2009).
130 F. Pietrucci, L. Mollica, and M. Blackledge, J. Phys. Chem. Lett. **4**, 1943 (2013).
131 D. Granata, C. Camilloni, M. Vendruscolo, and A. Laio, Proc. Natl. Acad. Sci. U.S.A. **110**, 6817 (2013).
132 Y. Wang, X. Chu, S. Longhi, P. Roche, W. Han, E. Wang, and J. Wang, Proc. Natl. Acad. Sci. U.S.A. **110**, E3743 (2013).
133 S. Singh, C.-C. Chiu, A. S. Reddy, and J. J. de Pablo, J. Chem. Phys. **138**, 155101 (2013).
134 D. Granata, F. Baftizadeh, J. Habchi, C. Galvagnion, A. De Simone, C. Camilloni, A. Laio, and M. Vendruscolo, Sci. Rep. **5** (2015).
135 F. Pietrucci, F. Marinelli, P. Carloni, and A. Laio, J. Am. Chem. Soc. **131**, 11811 (2009).
136 A. Kranjc, S. Bongarzone, G. Rossetti, X. Biarns, A. Cavalli, M. L. Bolognesi, M. Roberti, G. Legname, and P. Carloni, J. Chem. Theory Comput. **5**, 2565 (2009).
137 F. Baftizadeh, X. Biarnes, F. Pietrucci, F. Affinito, and A. Laio, J. Am. Chem. Soc. **134**, 3886 (2012).
138 F. Baftizadeh, F. Pietrucci, X. Biarnés, and A. Laio, Phys. Rev. Lett. **110**, 168103 (2013).
139 F. Pietrucci and A. Laio, J. Chem. Theory Comput. **5**, 2197 (2009).
140 D. J. Bicout and A. Szabo, J. Chem. Phys. **109**, 2325 (1998).
141 G. Hummer, New J. Phys. **7**, 34 (2005).
142 K. Bryan and T. Leise, Siam Rev. **48**, 569 (2006).
143 D. J. Wales and P. Salamon, Proc. Natl. Acad. Sci. U.S.A. **111**, 617 (2014).
144 F. Pietrucci and W. Andreoni, J. Chem. Theory Comput. **10**, 913 (2014).
145 E. Balan, F. Pietrucci, C. Gervais, M. Blanchard, J. Schott, and J. Gaillardet, Geochim. Cosmoch. Acta **193**, 119 (2016).
146 G. A. Gallet and F. Pietrucci, J. Chem. Phys. **139**, 074101 (2013).
147 F. Pietrucci and R. Martoňák, J. Chem. Phys. **142**, 104704 (2015).
148 J. L. Abascal and C. Vega, J. Chem. Phys. **123**, 234505 (2005).
149 R. Martoňák, A. Laio, and M. Parrinello, Phys. Rev. Lett. **90**, 075503 (2003).
150 T.-Q. Yu and M. E. Tuckerman, Phys. Rev. Lett. **107**, 015701 (2011).
151 D. Quigley and P. M. Rodger, J. Chem. Phys. **128**, 154518 (2008).
152 W. Lechner and C. Dellago, J. Chem. Phys. **129**, 114707 (2008).
153 A. Samanta, M. E. Tuckerman, T.-Q. Yu, and E. Weinan, Science **346**, 729 (2014).
154 F. Giberti, M. Salvalaglio, and M. Parrinello, IUCr J. **2**, 256 (2015).
155 W. E, W. Ren, and E. Vanden-Eijnden, J. Phys. Chem. B **109**, 6688 (2005).
156 L. Maragliano, A. Fischer, E. Vanden-Eijnden, and G. Ciccotti, J. Chem. Phys. **125**, 024106 (2006).
157 D. Branduardi, M. De Vivo, N. Rega, V. Barone, and A. Cavalli, J. Chem. Theory Comput. **7**, 539 (2011).
158 M. Ferus, F. Pietrucci, A. Saitta, A. Knizek, P. Kubelik, O. Ivanek, V. Shestivska, and S. Civis, Proc. Natl. Acad. Sci. U.S.A. **114**, 4306 (2017).
159 G. Cassone, F. Pietrucci, F. Saija, F. Guyot, and A. M. Saitta, Chem. Sci. **8**, 2329 (2017).
160 M. Bonomi, D. Branduardi, G. Bussi, C. Camilloni, D. Provasi, P. Raiteri, D. Donadio, F. Marinelli, F. Pietrucci, R. A. Broglia, *et al.*, Computer Phys. Commun. **180**, 1961 (2009).
161 G. A. Tribello, M. Bonomi, D. Branduardi, C. Camilloni, and G. Bussi, Computer Phys. Commun. **185**, 604 (2014).
162 G. Fiorin, M. L. Klein, and J. Henin, Mol. Phys. **111**, 3345 (2013).
163 K. Lindorff-Larsen, S. Piana, R. O. Dror, and D. E. Shaw, Science **334**, 517 (2011).
164 X. Huang, G. R. Bowman, and V. S. Pande, J. Chem. Phys. **128**, 205106 (2008).
165 H. Huang, E. Ozkirimli, and C. B. Post, J. Chem. Theory Comput. **5**, 1304 (2009).
166 L. Sutto, M. DAbramo, and F. L. Gervasio, J. Chem. Theory Comput. **6**, 3640 (2010).
167 M. A. Cuendet and M. E. Tuckerman, J. Chem. Theory Comput. **10**, 2975 (2014).
168 A. C. Pan, T. M. Weinreich, S. Piana, and D. E. Shaw, J. Chem. Theory Comput. **12**, 1360 (2016).