

 Open access • Journal Article • DOI:10.1021/ACS.JCTC.5B00597

Merging metadynamics into hyperdynamics: accelerated molecular simulations reaching time scales from microseconds to seconds. — [Source link](#)

Kristof M. Bal, Erik C. Neyts

Institutions: University of Antwerp

Published on: 15 Sep 2015 - Journal of Chemical Theory and Computation (American Chemical Society)

Topics: Metadynamics

Related papers:

- [Hyperdynamics: Accelerated Molecular Dynamics of Infrequent Events](#)
- [Temperature-accelerated dynamics for simulation of infrequent events](#)
- [Direct observation of realistic-temperature fuel combustion mechanisms in atomistic simulations.](#)
- [Fast parallel algorithms for short-range molecular dynamics](#)
- [Escaping free-energy minima](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/merging-metadynamics-into-hyperdynamics-accelerated-1g10c6howf>

This item is the archived peer-reviewed author-version of:

Merging metadynamics into hyperdynamics : accelerated molecular simulations
Reaching time scales from microseconds to seconds

Reference:

Bal Kristof, Neyts Erik.- Merging metadynamics into hyperdynamics : accelerated molecular simulations Reaching time scales from microseconds to seconds

Journal of chemical theory and computation - ISSN 1549-9618 - 11:10(2015), p. 4545-4554

Full text (Publishers DOI): <http://dx.doi.org/doi:10.1021/acs.jctc.5b00597>

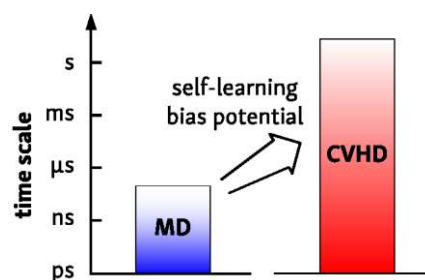
To cite this reference: <http://hdl.handle.net/10067/1281830151162165141>

Merging Metadynamics into Hyperdynamics: Accelerated Molecular Simulations Reaching Timescales from Microseconds to Seconds

Kristof M. Bal* and Erik C. Neyts

Department of Chemistry, University of Antwerp, Universiteitsplein 1, 2610 Wilrijk, Antwerp, Belgium

The hyperdynamics method is a powerful tool to simulate slow processes at the atomic level. However, the construction of an optimal hyperdynamics potential is a far from trivial task. We here propose a generally applicable implementation of the hyperdynamics algorithm, borrowing two concepts from



metadynamics. First, the use of a collective variable (CV) to represent the accelerated dynamics gives the method a very large flexibility and simplicity. Second, a metadynamics procedure can be used to construct a suitable history-dependent bias potential on-the-fly, effectively turning the algorithm into a self-learning accelerated molecular dynamics method. This collective variable-driven hyperdynamics (CVHD) method has a modular design: both the local system properties on which the bias is based, as well as the characteristics of the biasing method itself, can be chosen to match the needs of the considered system. As a result, system-specific details are abstracted from the biasing algorithm itself, making it extremely versatile and transparent. The method is tested on three model systems: diffusion on the Cu(001) surface and nickel-catalyzed methane decomposition as examples of “reactive” processes with a bond length-based CV, and the folding of a long polymer-like chain, using a set of dihedral angles as CV. Boost factors up to 10^9 , corresponding to a time scale of seconds, could be obtained while still accurately reproducing correct dynamics.

* Corresponding author, e-mail: kristof.bal@uantwerpen.be

1. INTRODUCTION

In order to gain insight into the fundamental dynamic processes of matter, the molecular dynamics (MD) method has been shown to be an indispensable tool. With MD simulations it is possible to study the dynamical evolution of an arbitrary system with atomic detail. However, MD simulations invariably suffer from severe time scale limitations. Indeed, whereas MD cannot be used to simulate time and lengths beyond the nanoscale, many relevant processes occur infrequently, beyond the μs time scale.

Several solutions to tackle the MD time scale problem exist. The family of accelerated MD methods developed by Voter and coworkers operate within the basic framework of MD, and use elevated temperatures,¹ bias potentials^{2,3} or parallelization^{4,5} to shorten the time between infrequent events. On the other end of the spectrum there are the various kinetic Monte Carlo (kMC) methods, where the full system evolution is represented as a coarse-grained sequence of infrequent events; the kMC event catalog must either be predefined,⁶ or is constructed on-the-fly during the simulation.⁷⁻¹² Finally, force-bias Monte Carlo simulations¹³⁻¹⁵ can be used to accelerate relaxation processes and push out-of-equilibrium systems towards global minima faster than MD,^{13,16} although they generally do not reproduce exact dynamic paths.¹⁶

A technique closely related to accelerated molecular dynamics is metadynamics.¹⁷ The metadynamics method was originally designed to explore reaction pathways and calculate free energy landscapes by using a history-dependent bias potential. It has been widely used in many scientific fields.^{18,19} The power of metadynamics lies in its simplicity: the bias potential is constructed as a function of a small number of *collective variables* (CVs) used to distinguish between all the relevant states of the system. Provided a CV can be developed for a certain process, it can be studied using metadynamics. Conceptually, it is very similar to hyperdynamics, in the sense that both methods rely on adding a bias potential to the global potential energy surface of the system. The bias potential ΔV is used to fill energy minima and hence shorten the waiting time between minima-to-minima transitions. In hyperdynamics, it is ensured that ΔV becomes zero in the transition state region, in which case correct relative dynamics is preserved.² Tiwary and Parrinello later recognized that if this condition is also enforced in metadynamics simulations, the method can also be used to accurately calculate the rates of slow processes,²⁰ even in complex systems.²¹

The CV-based flexibility of metadynamics, combined with the recent notion that it is able to correctly describe system dynamics, makes it a promising candidate for a more

generally applicable bias potential-based accelerated molecular dynamics method. Various hyperdynamics-based approaches have been developed²²⁻³⁰ precisely for this purpose, but these are usually specifically tailored to the (sub)class of systems they were developed for: the bias potential directly depends on a property of the studied system. In contrast, the biasing in a more general CV-based method such as metadynamics only depends on the used CVs, and applying it to a new system only requires changing the CVs, rather than adapting the overall method. Moreover, a particular point of interest in the field of hyperdynamics is the development of self-learning variants of the algorithm,³¹⁻³³ in which the applied bias potential is constructed on-the-fly during the simulation. This could also be achieved using the metadynamics algorithm, which is self-learning by construction.

However, there is a difference in scope between metadynamics rate calculations, and accelerated molecular dynamics in a general sense. Metadynamics rate calculations are used to characterize the dynamical pathways connecting a limited number of states, say, A and B , which are already known (e.g., the α and β states of alanine dipeptide²⁰ or the bound and unbound states of a protein-ligand complex²¹). In this case, the choice of CVs is dictated by the requirement that these can be used to distinguish between states A and B , while additional CVs might be required to properly differentiate between other states discovered along the $A \leftrightarrow B$ path. Such an approach, however, is not suited for a more “explorative” accelerated molecular dynamics study which aims to find the “natural” unconstrained state-to-state dynamics over long time scales, with the only state known in advance being the initial state A : an example of such a process would be the prediction of the product composition of complex chemical process, or the outcome of a growth procedure. In such an application, it is very difficult to select a proper set of CVs capable of describing this full, *a priori* unknown pathway. Indeed, metadynamics rate calculations can be used to rigorously calculate the rates of processes in a limited part of the phase space, whereas it would be interesting to have a similar method that is capable of uncovering unknown processes throughout (in principle) the full phase space.

In this paper, we describe an accelerated molecular dynamics method that combines the CV-based flexibility of metadynamics with the ability of hyperdynamics to track the natural long time scale evolution of a complex system. In Sec. 2, we describe a generic class of CVs that can be used to accelerate arbitrary processes in systems with many relevant degrees of freedom. Such a CV can then be used in a traditional hyperdynamics implementation with a predefined bias. Moreover, it is also possible to use these CVs in

a metadynamics protocol to construct the bias potential on-the-fly. To demonstrate the flexibility of our approach, the *collective variable-driven hyperdynamics* (CVHD) method, we apply it to three very different systems in Sec. 3. As a first example, diffusion on the prototypical Cu(001) model system is used as an example of a “reactive” process involving bond breaking. The second process is the folding of a long polymer-like chain is a system in which the activated process involves the rotation around bonds, changing the various dihedral angles in the chain structure. The third process is the full methane decomposition pathway on the Ni(111) surface, as an example of heterogeneous catalysis.

2. THEORY

2.1. Global structure of the CVHD method

In hyperdynamics, simulations are not carried out on the true potential energy surface $V(\mathbf{R})$, but on a modified potential $V^*(\mathbf{R})$, which is obtained by adding a bias potential $\Delta V(\mathbf{R})$:^{2,3}

$$V^*(\mathbf{R}) = V(\mathbf{R}) + \Delta V(\mathbf{R}). \quad (1)$$

A key simplification of the CVHD method is that the function $\Delta V(\mathbf{R})$ is reduced to a function $\Delta V(\eta)$ of only one parameter, the CV η . The method is modular by design, as depicted in Figure 1. The central element of the method is η , which can be thought of as a global reaction coordinate and can take continuous values between 0 and 1. To complete the method, two additional elements must be defined. These are (a) the set of local degrees of freedom that are appropriate properties to gauge the state of the system, and (b) the bias potential $\Delta V(\eta)$ that is only a function of η .

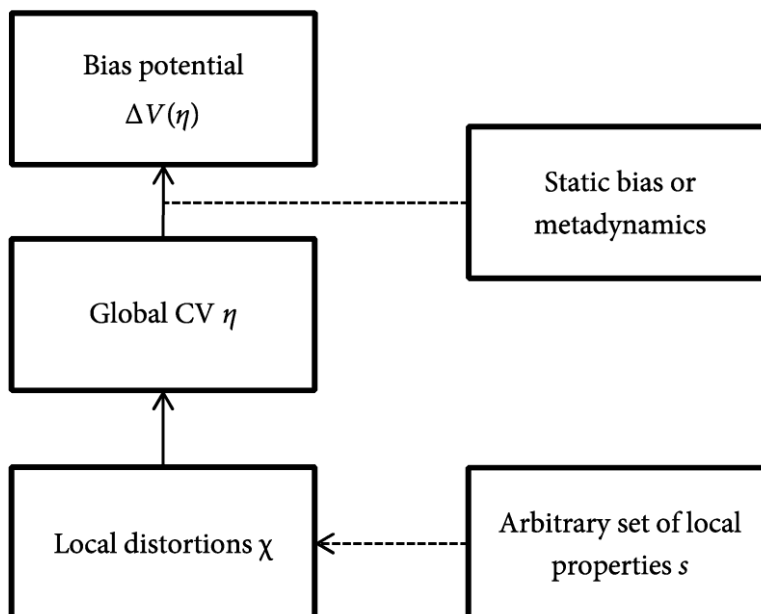


Figure 1: Schematic depiction of the structure of the CVHD algorithms. See text for details.

2.2. The CV

The key point in applying a method such as metadynamics is the choice of an appropriate CV to describe the relevant dynamics. For example, the coordination number (CN) CV³⁴ was developed for reactive systems:

$$\text{CN} = \sum_i^{\text{pairs}} \frac{1 - (r_i/d)^n}{1 - (r_i/d)^m}, \quad (2)$$

in which $n < m$ are integers, r_i the distance between an atom pair i , and d a parameter that controls the decay of the pair's contribution when r_i increases. Such type of CV is well-suited for the study of mechanisms and free energy profiles of individual reactions using metadynamics. Applying this CV, however, is only possible if the reactive centers or bonds are already known, and it is difficult to use it to discover new pathways. First, a system-wide coordination number CV will not be able to distinguish between various states. For example, consider the diffusion of a single adatom on a surface: every possible stable state will be represented by the exact same global CN value. Second, using the coordination CV would make it extremely difficult to guarantee that no bias is deposited in transition state regions. Every contact contributes equally to the total CV which, when there are many atom pairs, will lead to random fluctuations in the CN value that are larger than the change caused by a bond effectively breaking.

Very similar problems occur in the context of the so-called “accelerated molecular dynamics” (aMD) implementation of hyperdynamics,²⁴⁻²⁶ in which the bias potential depends on the instantaneous value of the system’s potential energy. Although conceptually very attractive because of its simplicity, the method has severe limitations. In practice, a bias potential is added as long as the energy is below a cutoff value. If one uses a large cutoff, needed if energy fluctuations in the system are large, many transition states will lie below the cutoff, within the boosted phase space region, thereby violating the main requirement of hyperdynamics. Using a low energy cutoff can circumvent this problem, but will not be very efficient because it will result in a small average boost. Again, if a process of interest only makes a rather small contribution to the global CV that is used to construct a bias potential, it is very difficult to either guarantee the correctness of the biased dynamics, or its efficiency.

Therefore, for this type of problems a CV must be developed in which the influence of a single degree of freedom on the total CV depends on its contribution to actual transitions, rather than treating all contributions equally. This way, large systems in which many different kinds of events occur can be studied, without having to “tag” reactive degrees of freedom in advance. Furthermore, it is impossible to distinguish all (possibly unknown) states the system will encounter using a single generic CV. Therefore, we will drop this requirement and will use a CV of which the value can only be used to describe the system’s position within the current state.

The CV we use here is inspired conceptually by the Bond Boost implementation of hyperdynamics,^{27,35} which is a quite powerful approach to accelerate reactive systems, and functionally by the CV introduced by Tiwary and van de Walle.³⁶ We generalize these approaches to a CV based on a set of *local system properties* or generalized degrees of freedom (s_1, \dots, s_N) , which groups all the relevant degrees of freedom (or CVs) of the full slow system dynamics. For each local property s_i , a *local distortion* $\chi_i = \chi(s_i)$ can now be defined, which is a function that can return values between 0 and 1. χ must be designed in such a way that if a property s_i is directly involved in a transition, and takes a value s_i^\ddagger at the corresponding dividing surface, $\chi(s_i^\ddagger) = 1$. If s_i is far enough from s_i^\ddagger and closer to its equilibrium value, $\chi_i < 1$. Whenever any property s_i is involved in a transition somewhere in the system (and hence $\chi_i = 1$), the system as a whole is about to cross a dividing surface, and the global CV describing the full system must reflect this. For this purpose, we calculate the *global distortion* χ_t as:³⁶

$$\chi_t = \left(\sum_{i=1}^N \chi_i^p \right)^{1/p}, \quad (3)$$

in which $p > 1$. While the algorithm in ref. 36 did not require continuous and vanishing derivatives at both “edges” of the CV, this is required in a hyperdynamics implementation like CVHD. Therefore, the actual CV η is here calculated from χ_t according to

$$\eta = \begin{cases} \frac{1}{2} (1 - \cos(\pi \chi_t^2)) & \text{if } \chi_t \leq 1 \\ 1 & \text{if } \chi_t > 1 \end{cases}. \quad (4)$$

In short, the exponent p is used in the calculation of χ_t to ensure that large distortions make a larger contribution to η than small ones. As a result, most of the bias energy flows into properties that are about to contribute directly to a transition (modulated by the magnitude of p), similar to what the Bond Boost method does with the breaking of bonds. Hence, the CV can be used to selectively boost large changes in small parts of the included phase space, insensitive to smaller fluctuations in the other parts. No biasing forces are applied to the system when $\chi_t > 1$: if $\chi(s)$ is properly designed, this means that no bias is added in a transition state region. Also, the specific form of χ_t does not only allow to describe transitions involving single properties ($\chi_t = 1$ because a single $\chi_i = 1$) but also the simultaneous distortion of multiple local properties (e.g., N properties all having only $\chi_i = (1/N)^{1/p}$ such that $\chi_t = 1$).³⁶ The precise value of p is not a critical parameter determining the accuracy of the CVHD method. Rather, it only controls to what extent the bias energy is distributed across the system: large values of p will suppress the influence of smaller distortions on η and will thus lead to a more localized bias potential. To optimize the method, p can be varied in function of how “concerted” events are expected to be, but its effect was found to be rather small for values between 4 and 12.

The final part of our algorithm deals with transitions. In order to be able to describe multiple consecutive events in a single biased MD run, a criterion is added to “reset” the procedure: if η remains equal to 1 during a waiting time t_w , the system is assumed to have undergone a transition, and thermalized in its new state. Then, a new property list is created, and the accelerated MD procedure is resumed.

A practical example of a property s is the stretch of a bond, which can be used to study reactive events involving bond breaking.^{27,36} Here, it is assumed that for every bond pair i with length r_i , there are distances r_i^{\min} and r_i^{\max} , which mark the begin and end point of possible reactive events. If $r_i < r_i^{\min}$, the bond is not likely to dissociate soon, and not biased, whereas if $r_i = r_i^{\max}$, the bond is about to dissociate and the system is close to a dividing surface. When the simulation starts, a list of bonds is created from all atomic pairs that are shorter apart than a cutoff r_i^{cut} . Then, local distortions can be calculated as

$$\chi_i = \begin{cases} 0 & \text{if } r_i < r_i^{\min} \\ \frac{r_i - r_i^{\min}}{r_i^{\max} - r_i^{\min}} & \text{if } r_i^{\min} < r_i < r_i^{\max} \\ 1 & \text{if } r_i > r_i^{\max} \end{cases} . \quad (5)$$

Given the relative simplicity of this implementation of η , which only has a few parameters, we expect that it is rather widely applicable. As it only requires the slow to-be-boosted process to involve bond breaking, the CV can be readily applied to a wide variety of processes. It is also not necessarily restricted to accelerated MD simulations. It could also, for example, be applied to more “traditional” metadynamics investigations such as the exploration of possible reaction pathways.

Finally, as explained above, the basic functional form of Eq. (5) could also be applied to properties other than bond lengths to study a wide variety of other processes. Indeed, as long as slow events lead to a significant distortion of a small subset in a large collection of local parameters, the CVHD algorithm will be efficient. In the given example, we use bond distances to construct the global distortion function χ_i , but the set of system properties s_i could also be something else, as we will show in Sec. 3.3. A simple modification of Eq. (5) could be the use of an atomic strain, rather than bonds, which is a suitable property to study dislocation nucleation.³⁰ The formalism as outlined in this section is sufficiently general and flexible to be used a starting point for the development of such new CVs, without having to modify the full boosting algorithm.

2.3. The bias potential

If η is now used as a CV in a metadynamics simulation, the bias potential $\Delta V(\eta)$ is slowly “grown” at intervals of τ_G by the metadynamics algorithm in the form of Gaussian functions with width δ and height w :

$$\Delta V(\eta) = \sum_{k < n_G} w_k \exp\left(-\frac{(\eta - \eta(k\tau_G))^2}{2\delta^2}\right). \quad (6)$$

This means that the bias potential is history-dependent, as the Gaussian bias functions depend on the CV value at the moment they were deposited. The nature of the metadynamics method ensures that this bias potential matches the underlying free energy surface of the studied processes and should thus guarantee a safe and efficient bias. An important difference between the CVHD method and metadynamics, however, is that the bias deposited in a state A is deleted once a new state is reached: even when state A is visited again later, the bias potential must be generated again. This is because the CV η does not have to distinguish all possible states in the system, like in conventional metadynamics¹⁷ and metadynamics rate calculations,²⁰ but is only required to identify the system’s position in the current state. Note that this also explains the use of a criterion based on t_w to determine whether a new state has been reached, rather than having this explicitly reflected by the CV value. The CVHD method thus gains flexibility in the size of the phase space that can be explored, but at the same time loses some efficiency.

The use of CVs is not limited to metadynamics, and could also be the basis of a traditional “static” hyperdynamics potential. η can satisfy the important constraint that no bias can be added in the transition state region, provided a proper form of χ is selected (or r_i^{\max} in Eq. (5)). This means that a simple linear function of η can be an appropriate bias potential:

$$\Delta V(\eta) = \Delta V^{\max} (1 - \eta) , \quad (7)$$

in which ΔV^{\max} is the maximal bias strength. Just as is the case for traditional hyperdynamics implementations, ΔV^{\max} must be chosen appropriately: large enough to obtain a substantial boost, but not larger than the barriers of interest. Note that this static approach requires some *a priori* knowledge of the possible events in the system, in contrast to the dynamic application (see below).

The expression for the bias potential itself should preferably not be specifically tailored for a particular system or process. Rather, in the spirit of other CV-based methods, all the complexity of the investigated process should be abstracted by the proper choice of a CV. As a result, the simple linear function of Eq. (7) is sufficient. As such, the CVHD family of methods is highly flexible. Both the underlying local property (as implemented by its associated distortion function) as the bias function itself can be changed to fit

the needs of the studied process. The only characteristic that is shared by all (sub)methods, as Figure 1 shows, is the CV η .

2.4. Comparison of static and dynamic CV-based bias

An important advantage of a metadynamics-based accelerated MD protocol (*dynamically biased* CVHD, dCVHD) is that it does not require the definition of an analytical bias potential function. Rather, a suitable bias potential is constructed on-the-fly. This means that as long the system dynamics can be represented by a combination of local contributions such as in Eq. (3), the algorithm can adapt to all kinds of processes with unknown activation barriers. As such, the algorithm can be interpreted as a self-learning implementation of hyperdynamics. However, if the studied system is reasonably well-characterized, a predefined *static CV-based bias* (*statically biased* CVHD, sCVHD) could be more suitable. Such a bias can be more efficient as it will immediately start at its full strength and does not need to be built up during the simulation. Therefore, it can be expected that the performance of the dynamic bias method will not be as good as for the static bias in the case of relatively fast successive events. Furthermore, a static bias eliminates the additional simulation parameters introduced by metadynamics: the Gaussian hill width, height and deposition rate. Both methods do, however, share the same fundamental structure (see Figure 1).

It must be stressed that although the dCVHD method uses a metadynamics protocol to construct the bias potential, it is strictly speaking not a metadynamics method. Indeed, the deposited bias is not stored during the whole simulation, as the CV η does not span the whole phase space, but only the current state. Once a new state is reached, all deposited bias is deleted and bias deposition is initiated again. Furthermore, bias is never deposited in transition state regions. Therefore, the CVHD method cannot be used to calculate free energy profiles, which require extensive complete sampling of a limited part of the phase space and for which “traditional” metadynamics is required.

In both the static and the dynamic approach, the speed-up relative to MD (the boost factor) is an ensemble average over the biased potential energy surface involving the bias potential $\Delta V(\eta)$:

$$\text{Boost} = \left\langle e^{\beta \Delta V(\eta)} \right\rangle, \quad (8)$$

in which $\beta = 1/k_B T$. The effectively simulated physical time, or the *hypertime*, can then be calculated by multiplying this boost factor with the MD time.² It must be noted that

accurate hypertime calculations require good sampling of the regions with large boost, which limits the imposed strength of a static bias potential.²⁷ In the case of a dynamic bias, however, the deposition of a large bias in a certain region of the CV space implies this region is frequently visited, because the number of hills deposited in a specific region depends on the time the system spends there. This means that the dynamic bias method can apply strong bias potentials while still maintaining a high accuracy of the calculated hypertime.

2.5. Critical aspects

When applying the CVHD method, some important considerations must be kept in mind. We here summarize the main critical aspects and limitations of the CVHD method in its current form.

Every valid accelerated MD method employing a bias potential has to ensure that this *bias vanishes in the transition state regions*. In our algorithms, the key to achieving this is the choice of χ . Our static bias approach will vanish at the transition state if for every property s_i , $\chi(s_i^\ddagger) = 1$. In the case of Eq. (5), this means that every r_i^{\max} of a bond should be smaller than the corresponding r_i at a transition state r_i^\ddagger . Parametrizing χ_i functions to achieve this behavior is not a trivial task, as it requires some knowledge of the transition states the system will encounter: a first estimation of “safe” parameter choices can be obtained from some (presumed) relevant transition states that are already known, and verified against other transitions that are discovered during an initial simulation. When using a dynamic bias, it is possible to impose the same constraint on r_i^{\max} as in the static case, which will cause all bias force to vanish when $r_i > r_i^{\max}$. The metadynamics algorithm, in contrast, will by default keep depositing bias at any time. This will become problematic when the system starts spending a large part of its time at the boundaries of the well ($\eta \sim 1$) once the bias is at its full strength in the minima, a situation that becomes more common with increasing system size.²⁷ If this happens, a large bias could be deposited at the boundary of the CV, which is unphysical and will negatively impact the accuracy of the hypertime calculation [Eq. (8)] and the dynamics in general. A simple solution to avoid this bias pile-up is to restrict the metadynamics algorithm from depositing any bias at large η values (for example, at $\eta > 0.9$). Additional control of the magnitude of deposited bias can be achieved through the well-tempered implementation of metadynamics.³⁷

A second phenomenon that can cause excess dynamic bias deposition is the presence of so-called *hidden CVs*. In such a situation, the used CV does not include all relevant “slow” degrees of freedom, which might block certain dividing surfaces, leading the system to cross other dividing surfaces. Especially when low-lying pathways are not included in the CV, this will lead to an erroneous overbiasing, a problem that is discussed in more detail in ref. 18. When using the functional form of the CVs used in the CVHD method, one has to ensure that all relevant local properties are included in the CV. This is both an inherent strength and a limitation of the CVHD method. On one hand, if it is possible to describe the full “slow” dynamics by a set of simple local properties, such as for example chemical reactions involving bond breaking being fully described by bond length local properties, no hidden CVs will be present. On the other hand, the requirement that the system dynamics can be decomposed into contributions by a small number of highly localized interactions renders the CVHD method impractical to study systems in which this is not the case: especially complex biological processes involving various types of non-bonded interactions are very difficult to study this way. Fortunately, the latter use case can be covered by metadynamics rate calculations, employing a well-tailored set of CVs,²¹ whereas the CVHD method is more suitable for, e.g., the prediction of reaction product compositions or the properties of grown materials.

An important disadvantage of CV-based methods (and many conventional hyperdynamics implementations as well) is their poor *scaling with the system size*: in large systems, events will occur more frequently, leading to additional overhead. Representing all to-be-accelerated dynamics by a single CV does not allow for a parallel treatment of events, leading to a deterioration of the boost in such large systems. Possibly, our method could be improved by incorporating aspects of a recently proposed “local” variant of the Bond Boost algorithm.³⁸ Similarly, existing solutions to deal with the small-barrier problem,³⁹ to improve the calculation of the hypertime,⁴⁰ or to construct appropriate CVs on-the-fly⁴¹ could also be beneficial for the methods presented here.

2.6. Related methods

To achieve additional boost, a joint application of our method with parallel replica^{4,5} or force-bias Monte Carlo¹³⁻¹⁶ methods could be possible. Parallel replica is not restricted to transition state theory and is therefore very generic. It is, unlike CVHD, not limited to transitions that can be localized in small parts of the system. This way, parallel replica can be used to accelerate transitions between large superstates of the system, rather than the microstate dynamics on which CVHD acts. Force-bias Monte Carlo methods, on the

other hand, do not require the definition of “events” at all: this makes them very suited to study processes in which “fast”, unbiased degrees of freedom can still provide a bottleneck in the case of very strongly biased “slow” dynamics. Of course, as is the case with conventional metadynamics simulations, bias deposition in dynamically biased simulations should always be sufficiently slow, in order to allow the system to equilibrate along the “fast” degrees of freedom not included in the CV.

A few other self-learning or adaptive hyperdynamics methods also exist. The recently proposed HD-MD method has a bias dependent on the system’s instantaneous potential energy (like the aMD method) and uses short MD runs within each local energy basin to parameterize a bias potential function on-the-fly.³¹ These parameters are chosen based on the desired boost, which can therefore be made equal in every new basin (this is not necessarily the case in CVHD). The fact that the bias potential is only a direct function of the system’s potential energy makes the method conceptually very simple. However, because the method is based on aMD and has no explicit way to detect transition states, the bias potential may be nonzero at dividing surfaces, in contrast to CVHD. Also within the context of the aMD method, a simpler algorithm, dubbed the Adaptive Hyperdynamics Method (AHD) method, was proposed.³² In the AHD method, the threshold energy (below which biasing is allowed) is adjusted in time intervals τ by comparing the current potential energy of the system to the minimal energy in the previous interval, and changing the threshold accordingly. Although very simple, the algorithm causes the biasing force to be discontinuous due to the stepwise modification of ΔV , which is not the case when metadynamics is used to deposit bias. Finally, within the framework of the Bond Boost method, a self-learning hyperdynamics method was derived by Perez and Voter.³³ This algorithm calculates an on-the-fly estimate of a pairwise potential of mean force (PMF), which is then used to iteratively improve the bias strength. The method was found to be very efficient in finding exactly the optimal boost for a given process, but lacks the generality and simplicity inherent to using generic CVs and a metadynamics-like self-learning bias.

3. APPLICATION

3.1. General methodology

All simulations were carried out using the LAMMPS package⁴² and a modified version of the Colvars module.⁴³ All modifications and their descriptions are provided in the Supporting Information. To control the system temperature, ensure its homogeneity, and allow for swift decorrelation, a Langevin-type thermostat⁴⁴ with relaxation time of 1 ps

was employed. The equations of motion were integrated with a MD time step of 1 fs, except when using the ReaxFF potential, which required a 0.1 fs time step. Boost and hypertime were calculated by evaluating Eq. (8) at every step; we did not compensate for the overhead induced by the bias calculation as it was found to be insignificant compared to the evaluation of the interatomic potential.

3.2. Bond-based CV: diffusion on Cu(001)

As an illustration of the bond-based CV in Eq. (5), we apply both the static and the dynamic CVHD methods to diffusion on the Cu(001) surface. Specifically, we apply the methods to the diffusion of adatoms, dimers and vacancies, which can all diffuse through simple hopping mechanisms. Copper adatom and dimer diffusion, however, can also occur by a two-atom exchange.⁴⁵ Thus, we can assess the performance of the bias methods for a set of competitive mechanisms with different characteristics in terms of the number of atoms and bonds involved, and minimal r_i^\ddagger . The same processes have also been previously studied with the Bond Boost method,²⁷ which allows us to directly compare the performance of our generic CV-based methods to a dedicated hyperdynamics implementation.

The studied system consisted of a six-layer slab, each layer containing 50 atoms. The Cu-Cu interactions were described using a standard EAM potential.^{46,47} The two bottom layers were kept fixed and, depending on the studied mechanism, an adatom or dimer was placed on top of the slab, or a vacancy was created by removing an atom from the top layer.

Climbing Image Nudged Elastic Band (cNEB)⁴⁸ calculations were used to obtain information on the minimal r_i^\ddagger values of breaking bonds associated with every mechanism. Of all the mechanisms considered, these are the smallest for adatom hopping, where the two partially broken bonds have a length of 3.3 Å at the transition state. Therefore, a global r_i^{\max} value of 3.3 Å was used for each bond in all simulations. The r_i^{\min} value for every bond was chosen to be the average bond length as obtained from an equilibration run of t_w (after an initial waiting time t_w before detecting a transition), and r_i^{cut} was a global constant of 3.0 Å. We furthermore set $p = 8$ and $t_w = 5$ ps. In the dynamic bias simulations, Gaussian hills with a height of 0.005 eV and width of 0.025 were added with an interval of 1 ps. Well-tempered metadynamics³⁷ with a bias temperature $\Delta T = 2000$ K is used to deposit the dynamic bias. In the hyperdynamics simulations, we set $\Delta V^{\max} = 0.3$ eV, as a compromise between boost and accurate hypertime sampling.

Reaction rates Γ for all event types were quantified by counting the observed number of diffusion events n_i and dividing this by the calculated hypertime, $\Gamma_i = n_i/t_{\text{hyper}}$. At temperatures between 150 and 600 K, rates were obtained and averaged over multiple runs of 2×10^8 steps (1×10^8 for vacancies). Then, the calculated rates were fitted to the Arrhenius relation $\Gamma(T) = \Gamma_0 e^{-\beta E_A}$. The fitted prefactors Γ_0 and activation energies E_A can be readily compared to those reported by Miron and Fichthorn in their Bond Boost study of the same system.²⁷

It can be seen from Table 1 that both the static and dynamic biasing methods yield the same Γ_0 and E_A , both in excellent agreement with the Bond Boost result and the static barriers calculated by the cNEB method. We have also explicitly compared CVHD-obtained rates with MD results and found excellent agreement, as discussed in the Supporting Information. The performance of both CV-based methods, as expressed by the achieved boost is, as Figure 2 shows, quite comparable, yielding a boost up to 3×10^4 for adatom diffusion at a temperature of 250 K and up to 10^9 for vacancy diffusion at 150 K. The methods show the same basic behavior typical for hyperdynamics methods, with a boost that declines with increasing temperature, due to the inverse temperature dependence of β . Furthermore, it can be seen that the relative efficiency of the static bias method compared to its dynamic counterpart improves with increasing temperature. This is because waiting times between events are shorter at higher temperatures which, as a result, puts the dynamic method at disadvantage. This difference becomes irrelevant at lower temperatures where waiting times, even with full-strength static bias, are much longer than the time needed by the metadynamics protocol to construct the dynamic bias.

Table 1: Prefactors (Γ_0) and activation energies (E_A) for elementary diffusion processes on the Cu(001) surface, as obtained from dynamic boost (DB) and static boost (SB) simulations (see text for details). For comparison, the Bond Boost (BB) estimates²⁷ and static barriers computed with the cNEB method are also included. DB and SB error bars reflect the 90% confidence interval.

Process	Γ_0^{DB} (THz)	Γ_0^{SB} (THz)	Γ_0^{BB} (THz)	E_A^{DB} (eV)	E_A^{SB} (eV)	E_A^{BB} (eV)	E_A^{cNEB} (eV)
Vacancy hop	$83e^{\pm 0.1}$	$54e^{\pm 0.5}$	$54e^{\pm 0.5}$	0.44 ± 0.01	0.43 ± 0.02	0.44 ± 0.03	0.44
Adatom hop	$54e^{\pm 0.3}$	$33e^{\pm 0.1}$	$40e^{\pm 0.5}$	0.53 ± 0.01	0.51 ± 0.01	0.52 ± 0.03	0.51
Adatom ex- change	$430e^{\pm 1.0}$	$130e^{\pm 1.2}$	$270e^{\pm 0.6}$	0.76 ± 0.04	0.71 ± 0.05	0.73 ± 0.04	0.71
Dimer hop	$34e^{\pm 0.2}$	$21e^{\pm 0.2}$	$30e^{\pm 0.7}$	0.51 ± 0.01	0.49 ± 0.02	0.47 ± 0.03	0.49
Dimer ex- change	$137e^{\pm 1.3}$	$213e^{\pm 1.2}$	$190e^{\pm 0.8}$	0.74 ± 0.06	0.76 ± 0.05	0.71 ± 0.06	0.70

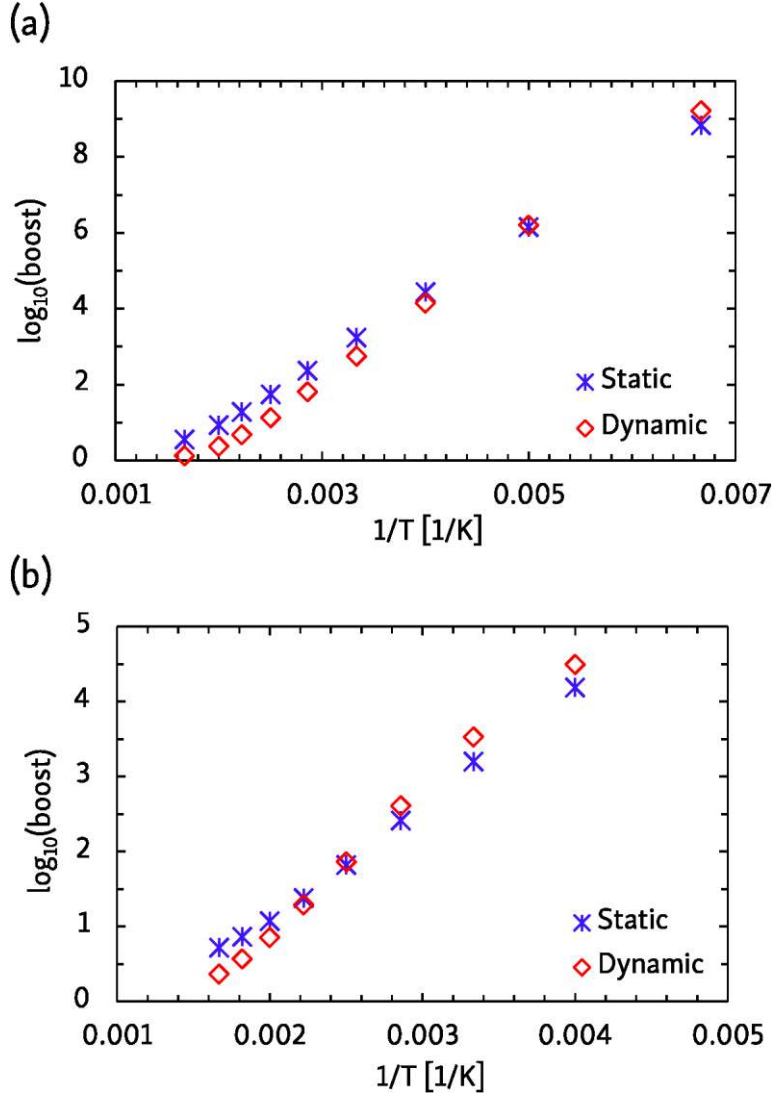


Figure 2: Calculated boost factor as function of the temperature, of both the static as the dynamic CVHD method, for (a) vacancy and (b) adatom diffusion on Cu(001). (As explained in the text, dimer diffusion behaves essentially the same as adatom diffusion.)

Not all processes occur at similar rates at every temperature. As Figure 2(a) shows, vacancy diffusion was studied at temperatures as low as 150 K, where we were able to observe about 60 events over a total MD time of 0.5 μ s: this is because the associated hypertime reached about 500 s. All other processes, however, are much slower at this temperature. For example, according to the kinetic parameters in Table 1, adatom hopping will be about 2000 times slower, with an average waiting time in the order of 10^4 s, explaining why we were able to observe the latter process with an appreciable frequency only from 250 K and higher, as depicted in Figure 2(b). The higher barriers and waiting times of the adatom diffusion processes, compared to vacancy diffusion, also increase the relative efficiency of the dynamic bias method, as expected. Similarly,

exchange processes could only be observed starting from 400 K, and show the same boost characteristics as their respective hopping counterparts. In general, the bond-based CV as used in the CVHD algorithm performs about as well as the Bond Boost method, albeit being part of a much more generic framework. Furthermore, because all relevant processes were already well-characterized, use of the dynamic biasing method has no added advantage since an optimized static bias could be applied at lower cost.

3.3. Dihedral-based CV: folding of a helix

To demonstrate the flexibility of the CVHD framework, we study a process governed by a different local property or CV: the folding of extended chain to a helix. This process is very different from the Cu(001) diffusion example in two ways. First, the activated processes underpinning the system evolution are not bond breaking, but rotation around bonds, changing the dihedral angles. Second, whereas the Cu(001) system remained in equilibrium, and every state was associated with the same handful of possible escape pathways, a folding chain may visit a much larger number of different states, all of which may have wildly different and unpredictable kinetic and thermodynamic stabilities. The dihedral angle is also a four-atom local property, setting it apart from the pairwise properties that are commonly used in hyperdynamics methods such as the Bond Boost method.

We use a simple model system, consisting of a chain of 50 connected beads, interacting with harmonic bond and angle functions. The dihedral potential energy term⁴⁹ is designed in such a way that there are three minima, of which the g^- state is preferred to the t and g^+ states. In our parameterization, the barrier for the $t \rightarrow g^-$ transition of a single dihedral is about 0.37 eV (or ~ 8 kcal/mol), and the g^- state is about 0.7 eV more stable than the two other options, thus favoring a helix over an extended chain. Further details are provided in the Supporting Information. Similar to the bond-based system property of Eq. (5), the local distortion function can be calculated from a dihedral angle ϕ_i as

$$\chi_i = \begin{cases} \left| \frac{\phi_i - \phi_i^{\text{ref}}}{\Delta\phi_i} \right| & \text{if } \phi_i \in [\phi_i^{\text{ref}} \pm \Delta\phi_i] \\ 1 & \text{if otherwise} \end{cases} \quad (9)$$

Here, ϕ_i^{ref} is a reference dihedral angle, which is determined as the dihedral angle of the closest local minimum (here, $\pm 60^\circ$ or $\pm 180^\circ$) at the moment the property list is created.

$\Delta\phi_i$ is the maximal deviation from ϕ_i^{ref} that keeps ϕ_i far enough from transitions, which we set to be 40° . The choice of $\Delta\phi_i$, like all parameterizations of a local distortion χ , requires some *a priori* knowledge of the system; in our case, we use a well-defined model potential, but in more complicated systems one should always verify that for all observed events the requirement that $\eta = 1$ at the transition state is satisfied. The other parameters do not depend on the local property used, and we set $p = 8$ and $t_w = 5$ ps. In the dynamic bias simulations, Gaussian hills with a height of 0.005 eV and width of 0.05 were added with an interval of 1 ps and a bias temperature $\Delta T = 1000$ K. In the case of the static bias, we use $\Delta V^{\text{max}} = 0.15$ eV.

We start every simulation from a fully extended chain, with all dihedrals in the t conformation (see Figure 3(a)), and run unbiased MD, static and dynamic bias simulations at 300 K, while monitoring the number of dihedral rotations. Because of the way our bias potential is designed, it is highly unlikely any other transitions will occur before there are no more $t \rightarrow g^-$ events possible. Therefore, 47 transitions were always found to be sufficient to reach a perfect helix [Figure 3(b)].

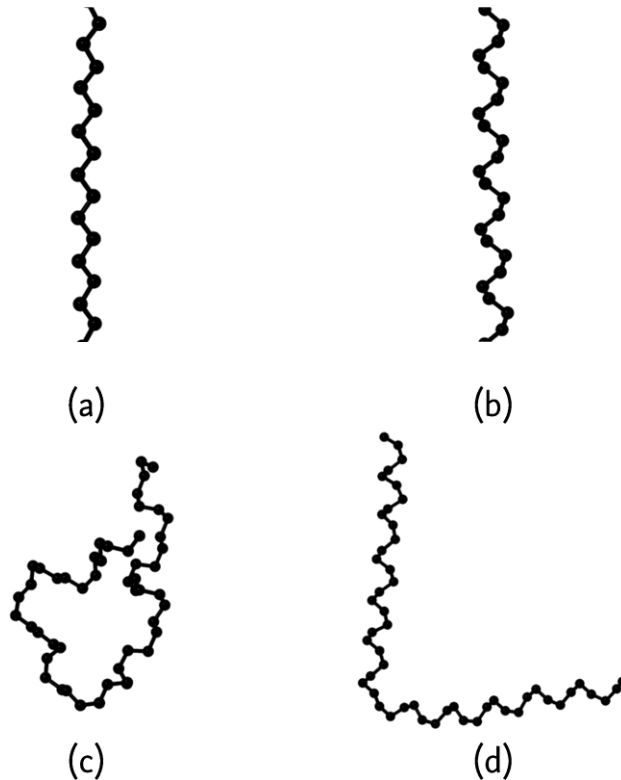


Figure 3: Some possible states of the helix model system: (a) fully extended, (b) perfect helix, (c) half-folded after 23 folding events and (d) after 46 (or, rather, $46 + 2n$) events.

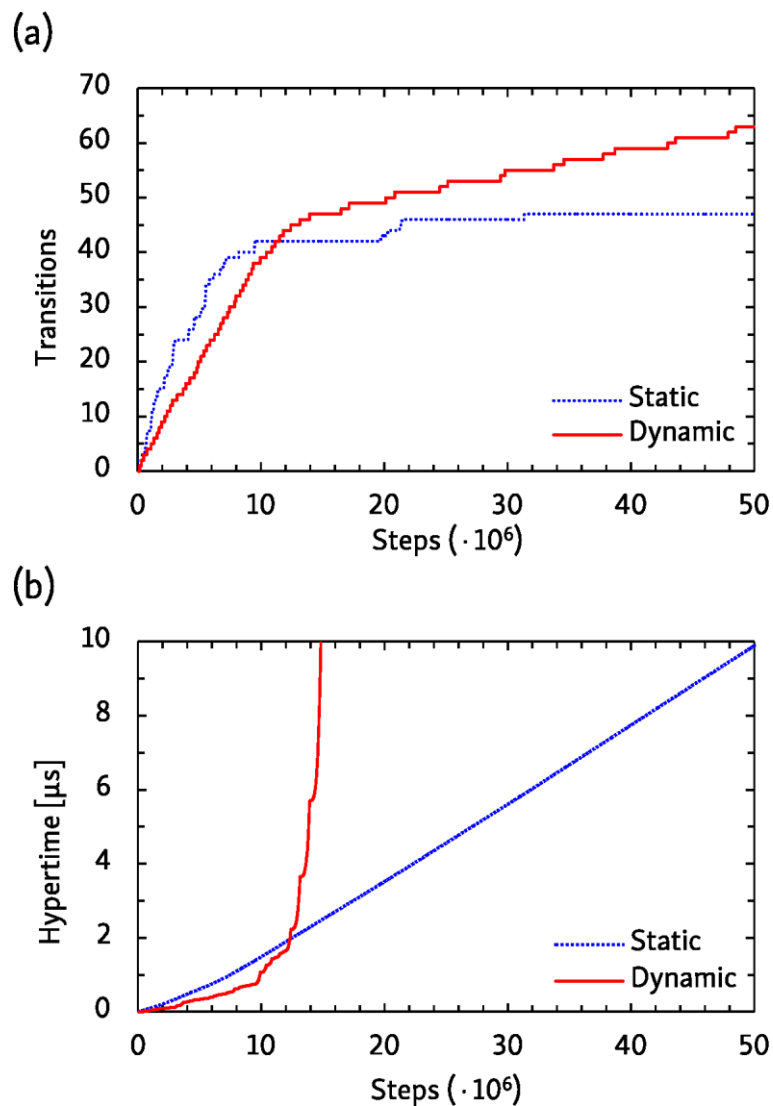


Figure 4: (a) Number of transitions as a function of simulation time in the helix folding test system, and (b) hypertime as a function of simulation time in the helix folding test system, for both static and dynamic bias (see text for details).

Examples of the system evolution in terms of the number of transitions are depicted in Figure 4(a) for both static and dynamic bias. It can be seen that initially, using a static bias is initially more efficient than a dynamic bias. As discussed before, a well-tailored static bias has the advantage of starting at its full strength, whereas a dynamic bias takes time to be constructed. However, we notice that after 42 transitions, the dynamically biased simulation surpasses the statically biased case. Indeed, because the studied process obeys first order kinetics in the number of t dihedrals, the waiting time between events increases the closer the system is to the fully folded state. The choice of $\Delta V^{\max} = 0.15$ eV was calibrated to the initial phase of the folding process, and is very well-suited for this first stage with relatively short waiting times, but becomes less effi-

cient when the number of t dihedrals is low and the waiting time is increased. The dynamic biasing scheme, on the other hand, will keep strengthening the bias potential while waiting for a transition to occur. The dCVHD method thus dynamically uses a larger bias in the case of long waiting times.

A different look at this subject is given by Figure 4(b), which compares the hypertime reached by both methods as a function of the simulation time. Whereas this quantity linearly increases for the static bias, indicating a constant boost, it shows an exponential increase in the case of the dynamic bias. Because of this property, the dynamic biasing scheme resulted in a perfect helix about two times faster than the static scheme.

The achieved boost factors of both biasing methods are, as Figure 4(b) shows, quite different. On one hand, the static biasing scheme provides a constant boost of about 200. On the other hand, the dynamic scheme adapts itself to match the boost requirements of the current state in which the system resides – arriving at an accumulated boost of about 500 when completing the folding process. This also explains the different behavior of both methods after the perfect helix is formed: any transition from the g^- state has a barrier that are about twice as high as the one associated with the $t \rightarrow g^-$ transition, which means that the statically biased simulation will not be able to escape from the helix state within a reasonable computational time. However, in the dynamically biased simulation, the bias will slowly be increased until a transition to a less favorable state [Figure 3(d)] can occur. Such a process only takes place after about 0.1 ms, much longer than the 5 μ s required to obtain the folded helix. The escape pathways from this unfavorable “kinked” helix have lower barriers than from the perfect helix. As a result, the construction of a suitably strong bias to return to the latter state requires fewer simulation steps than the reverse reaction, explaining the successive occurrence of short and long “steps” in the dynamic transitions curve in Figure 4(a).

Table 2: Average folding and half-folding times (μ s) for the helix model system, as obtained from unbiased MD and CV-based simulations with static and dynamic bias, respectively. All error bars are at the 90% confidence level.

Process	Unbiased MD	Static bias	Dynamic bias
Half folding	0.45 ± 0.08	0.38 ± 0.05	0.39 ± 0.04
Full folding	6 ± 2	7 ± 2	6 ± 1

Although the general performance characteristics of both accelerated methods in the case of the folding model system have now been established, their accuracy must still be ascertained. Table 2 therefore collects the average times for full folding (47 transitions) and obtaining a half-folded structure, such as the one depicted in Figure 3(c) (23 transitions), as obtained by biased simulations but also long unbiased MD runs. For both processes, all considered methods are in excellent agreement with each other.

3.4. A complex, realistic process: methane dissociation on Ni(111)

Finally, as an example of the kind of complex dynamics that can be accessed with CVHD simulations, we consider the catalytic dissociation of methane on the Ni(111) surface. This process is important not only in methane reforming processes, but also in chemical vapor deposition growth of carbon nanostructures. The initial dissociative adsorption of CH₄ has an activation energy in the order of 20 kcal/mol,⁵⁰ rendering direct MD simulations of this process difficult; to observe appreciable CH₄ dissociation, previous simulation attempts were required to use elevated temperatures⁵¹⁻⁵³ (up to 1500 K), instead of experimental temperatures of 800-1000 K, or only focus on plasma-activated species (CH_x radicals).^{54,55} From a technical point of view, this type of reaction is a useful additional test case for the CVHD method, being both an example of a system with a phase boundary, and of heterogeneous catalysis in general.

The methane dissociation process is modeled starting from a single CH₄ molecule above a six-layer nickel slab (64 atoms per layer), with the two bottom layers held fixed. A reflective wall is used at a z-height of 20 Å, leading to a gas phase volume of about 19.9×17.4×8 Å³. The interatomic interactions are described by the ReaxFF potential⁵⁶ as implemented in LAMMPS⁵⁷, using the Ni/C/H parameter set of Mueller et al.⁵⁸ and the QEq method⁵⁹ to calculate atomic charges.

The simulations are carried out at 800 K, applying the bond-based CV of Eq. (5) with $p = 6$ and $t_w = 0.1$ ps to C-H bonds, and dynamic biasing with a deposition stride of 10 fs, a hill width of 0.025, a hill height of 0.25 kcal/mol and a bias temperature of 4000 K. Compared to the previously discussed processes, metal-catalyzed methane decomposition poses two additional challenges. First, the general problem of thermostating gas phase species is that it is a poor model of energy exchange in such a system: in reality, this only occurs at discrete moments in time during collisions. Also, a Langevin thermostat distorts the diffusion path of gas-phase particles. However, thermostating the methane molecule is necessary to dissipate the excess energy introduced by the dynamic bias procedure, and to avoid unphysical heating of the molecule. As a compro-

mise, we apply a separate Langevin thermostat to the CH_4 molecule that only acts on the vibrational and rotational degrees of freedom, and leaves its translational motion untouched. The second problem is specific to the successive dehydrogenation pathway of methane, in which the separate steps have very different bond lengths at the transition state (r^\ddagger) for the dissociating bonds, ranging from 1.55 Å for $\text{CH} \rightarrow \text{C} + \text{H}$ to 1.80 Å for $\text{CH}_3 \rightarrow \text{CH}_2 + \text{H}$ on the Ni(111) surface.⁶⁰ Safe r^{max} values for the former were found to perform very poorly when attempting to boost the latter process. Therefore, we used a global r^{max} value of 1.8 Å and $r^{\text{min}} = 0.9$ Å. Although the safety of such a setting is not completely guaranteed, we found that little to no bias was effectively deposited in the transition state regions of the “unsafe” cases.

Table 3: Average reaction time for all elementary reaction steps of the full methane dehydrogenation process $\text{CH}_4(\text{g}) \rightarrow \text{C}(\text{ad}) + 4\text{H}(\text{ad})$, on Ni(111) at 800 K, as obtained from ReaxFF dCVHD simulations. Reaction times are given as a 90% confidence interval.

Process	Reaction time
$\text{CH}_4(\text{g}) \rightarrow \text{CH}_3(\text{ad}) + \text{H}(\text{ad})$	4–9 μs
$\text{CH}_3(\text{ad}) \rightarrow \text{CH}_2(\text{ad}) + \text{H}(\text{ad})$	0.09–0.22 μs
$\text{CH}_2(\text{ad}) \rightarrow \text{CH}(\text{ad}) + \text{H}(\text{ad})$	37–91 ps
$\text{CH}(\text{ad}) \rightarrow \text{C}(\text{ad}) + \text{H}(\text{ad})$	0.3–0.8 ms

We carried out 15 independent simulations of 10^6 steps, corresponding to a MD time of 100 ps each, and were always able to observe the full methane decomposition process. For every elementary step in the reaction, we calculated the average reaction time, summarized in Table 3. These results demonstrate the usefulness and power of a dynamic biasing method. Indeed, methane decomposition at 800 K is a process that consists of rather fast steps such as the dissociation of a C-H bond of adsorbed CH_2 (which takes about 50 ps), to the very slow decomposition of adsorbed CH, which requires more than 0.1 ms. Studying this reaction sequence with a static bias would therefore not be achievable; the vast time scale spread of the various elementary processes is illustrated in Figure 5. Boosts of 2×10^6 are achieved.

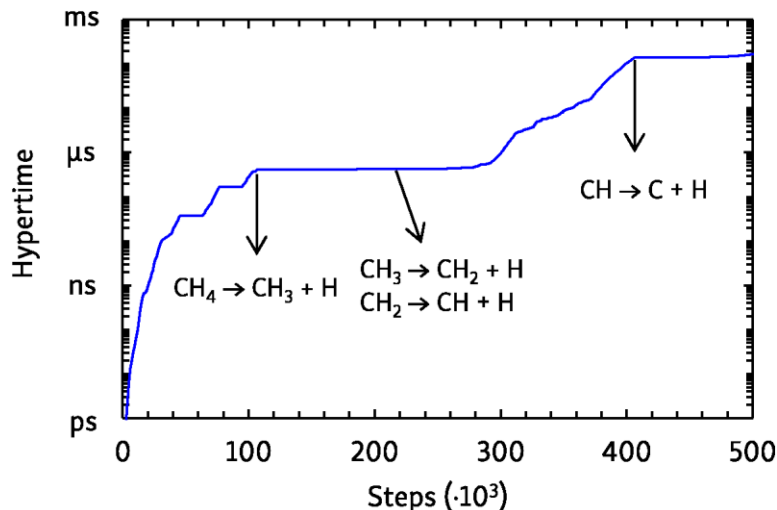


Figure 5: Hypertime evolution in a dCVHD simulation of methane decomposition on Ni(111), at 800 K. The observed elementary steps are shown at the time step at which they occurred.

The very long reaction time of the various reactions means that we cannot verify the accuracy of the values in Table 3 by direct comparison to MD simulations: even at a temperature of 1000 K we did not observe any reaction within 100 ps. It is, however, possible to compare the relative rates of the elementary steps to estimations based on differences between their respective activation energies.⁶⁰ The barrier for C-H dissociation in adsorbed CH_3 is 10 kcal/mol higher than for adsorbed CH_2 , meaning that the latter is about 1000 times faster than the former, in agreement with our findings. Similarly, the dissociation of adsorbed CH has a barrier that is 14 kcal/mol higher than the dissociation step involving adsorbed CH_3 , leading to a rate that differs by an order of magnitude of 10^4 , again in agreement with the results in Table 3. Finally, according to kinetic theory, the initial CH_4 pressure is about 40 bar, with a flux to the surface of 0.3 ps^{-1} . Considering a dissociation barrier of 19 kcal/mol,⁵⁸ we can make a crude estimation of the average reaction time to be $0.5 \text{ } \mu\text{s}$, which is also in line with our observations.

4. CONCLUSION

We have developed a theoretical framework, the collective variable-driven hyperdynamics (CVHD) method, which is an implementation of the hyperdynamics method that includes some of the strengths of metadynamics. The CVHD method is intended to be used as an accelerated molecular dynamics method, in which the waiting time between infrequent events is shortened by adding a bias potential to the energy minima in the system, without requiring *a priori* knowledge of the pathways and states that will be

encountered. From metadynamics, the method borrows the concept of using collective variables (CVs) to describe the system's dynamics and to express the bias potential as a simple function of the CV, essentially a generalized reaction coordinate, giving rise to the *statically biased* CVHD (sCVHD) method. However, the metadynamics algorithm can also be used to dynamically build up a suitable bias potential for every new potential energy basin the system encounters. This *dynamically biased* CVHD (dCVHD) method is effectively a self-learning hyperdynamics implementation and does not require an *a priori* knowledge of the activation barriers the system can encounter during its long time scale evolution.

A key point of the CVHD methods is its modular design. All relevant dynamics is represented by a single global CV η , which measures the distortion associated with an arbitrary set of local degrees of freedom. Both the biasing method, which depends on the CV, as the local properties on which the CV depends can be chosen independently to be optimal for the system studied. In this work, we have demonstrated the applicability of the bond length and dihedral angle local properties, in the study of solid state diffusion and heterogeneous catalysis, and chain folding, respectively.

If the studied process is already well-characterized and all relevant activation barriers are known, using a static bias (the sCVHD method) is the optimal choice: a well-optimized static bias can be constructed, and the on-the-fly construction of a dynamic bias will only cause additional overhead. On the other hand, in systems undergoing a more complex evolution, using a dynamical metadynamics-based bias may be the more optimal choice, as it is generally not possible to construct a single static bias that is both safe and efficient for every process encountered. This ability of the dCVHD method to adapt its bias to the specific requirements of the system at any time is an important advantage of the method. Irrespective of their relative efficiency, however, both biasing methods give rise to a correct sequence of state-to-state transitions.

Although the CVHD method is inherently flexible in the kind of local properties it can use to calculate its global CV, its performance does not seem to suffer from this genericity. For example, in the case of low-temperature diffusion on the Cu(001) with the bond length local property, accelerations as large 10^9 can be obtained, corresponding to physical times up to several seconds. In general, the CVHD method is about as efficient as the Bond Boost implementation of hyperdynamics, but has the added advantage of being more general. The local distortion functions developed so far already span a large range of processes and systems, and additional ones can be incorporated to further ex-

tend the scope of the method. We therefore believe that the CVHD method will be a valuable tool in the study of slow or activated processes in a wide range of scientific fields including growth, conformational sampling, and catalysis.

ASSOCIATED CONTENT

Supporting Information

Description and codes of the CVHD implementation, explicit comparison of diffusion rates from CVHD and MD for Cu/Cu(001), and description of the potential energy functions of the dihedral test system. This material is available free of charge via the Internet at <http://pubs.acs.org>.

AUTHOR INFORMATION

Corresponding author

Kristof Bal
Universiteitsplein 1, 2610 Antwerp, Belgium
+32-3-265.23.69
kristof.bal@uantwerpen.be

Notes

The authors declare no competing financial interest.

ACKNOWLEDGEMENTS

K.M.B. is funded as PhD fellow (aspirant) of the FWO-Flanders (Fund for Scientific Research-Flanders), Grant 11V8915N. The computational resources and services used in this work were provided by the VSC (Flemish Supercomputer Center) and the HPC infrastructure of the University of Antwerp (CalcUA), funded by the Hercules Foundation and the Flemish Government – department EWI.

REFERENCES

- [1] Sørensen, M. R.; Voter, A. F. *J. Chem. Phys.* **2000**, *112*, 9599–9606.
- [2] Voter, A. F. *J. Chem. Phys.* **1997**, *106*, 4665–4677.
- [3] Voter, A. F. *Phys. Rev. Lett.* **1997**, *78*, 3908–3911.
- [4] Voter, A. F. *Phys. Rev. B* **1998**, *57*, R13985–R13988.

- [5] Perez, D.; Uberuaga, B. P.; Shim, Y.; Amar, J. G.; Voter, A. F. *Annu. Rep. Comput. Chem.* **2009**, *5*, 79–98.
- [6] Voter, A. F. *Phys. Rev. B* **1986**, *34*, 6819–6829.
- [7] Henkelman, G.; Jónsson, H. *J. Chem. Phys.* **2001**, *115*, 9657–9666.
- [8] Xu, L.; Henkelman, G. *J. Chem. Phys.* **2008**, *129*, 114104.
- [9] Xu, H.; Osetsky, Y. N.; Stoller, R. E. *Phys. Rev. B* **2011**, *84*, 132103.
- [10] El-Mellouhi, F.; Mousseau, N.; Lewis, L. J. *Phys. Rev. B* **2008**, *78*, 153202.
- [11] Béland, L. K.; Osetsky, Y. N.; Stoller, R. E.; Xu, H. *Comput. Mater. Sci.* **2015**, *100*, 124–134.
- [12] Peter, E. K.; Shea, J.-E. *Phys. Chem. Chem. Phys.* **2014**, *16*, 6430–6440.
- [13] Timonova, M.; Groenewegen, J.; Thijsse, B. J. *Phys. Rev. B* **2010**, *81*, 144107.
- [14] Neyts, E. C.; Thijsse, B. J.; Mees, M. J.; Bal, K. M.; Pourtois, G. *J. Chem. Theory Comput.* **2012**, *8*, 1865–1869.
- [15] Mees, M. J.; Pourtois, G.; Neyts, E. C.; Thijsse, B. J.; Stesmans, A. *Phys. Rev. B* **2012**, *85*, 134301.
- [16] Bal, K. M.; Neyts, E. C. *J. Chem. Phys.* **2014**, *141*, 204104.
- [17] Laio, A.; Parrinello, M. *Proc. Natl. Acad. Sci. USA* **2002**, *99*, 12562–12566.
- [18] Barducci, A.; Bonomi, M.; Parrinello, M. *WIREs Comput. Mol. Sci.* **2011**, *1*, 826–843.
- [19] Zheng, S.; Pfaendtner, J. *Mol. Simul.* **2015**, *41*, 55–72.
- [20] Tiwary, P.; Parrinello, M. *Phys. Rev. Lett.* **2013**, *111*, 230602.
- [21] Tiwary, P.; Limongelli, V.; Salvalaglio, M.; Parrinello, M. *Proc. Natl. Acad. Sci. USA* **2015**, *112*, E386–E391.
- [22] Pal, S.; Fichthorn, K. A. *Chem. Eng. J.* **1999**, *74*, 77–83.
- [23] Wang, J.-C.; Pal, S.; Fichthorn, K. A. *Phys. Rev. B* **2001**, *63*, 085403.
- [24] Hamelberg, D.; Mongan, J.; McCammon, J. A. *J. Chem. Phys.* **2004**, *120*, 11919–11929.
- [25] Markwick, P. R. L.; McCammon, J. A. *Phys. Chem. Chem. Phys.* **2011**, *13*, 20053–20065.

- [26] Pierce, L. C.; Salomon-Ferrer, R.; de Oliveira, C. A. F.; McCammon, J. A.; Walker, R. C. *J. Chem. Theory Comput.* **2012**, *8*, 2997–3002.
- [27] Miron, R. A.; Fichthorn, K. A. *J. Chem. Phys.* **2003**, *119*, 6210–6216.
- [28] Becker, K. E.; Mignogna, M. H.; Fichthorn, K. A. *Phys. Rev. Lett.* **2009**, *102*, 046101.
- [29] Lin, Y.; Fichthorn, K. A. *Phys. Rev. B* **2012**, *86*, 165303.
- [30] Hara, S.; Li, J. *Phys. Rev. B* **2010**, *82*, 184114.
- [31] Paz, S. A.; Leiva, E. P. M. *J. Chem. Theory Comput.* **2015**, *11*, 1725–1734.
- [32] Hirai, H. *J. Chem. Phys.* **2014**, *141*, 234109.
- [33] Perez, D.; Uberuaga, B. P.; Shim, Y.; Amar, J. G.; Voter, A. F. *Annu. Rep. Comput. Chem.* **2009**, *5*, 79–98.
- [34] Iannuzzi, M.; Laio, A.; Parrinello, M. *Phys. Rev. Lett.* **2003**, *90*, 238302.
- [35] Fichthorn, K. A.; Miron, R. A.; Wang, Y.; Tiwary, Y. *J. Phys.: Condens. Matter* **2009**, *21*, 084212.
- [36] Tiwary, P.; van de Walle, A. *Phys. Rev. B* **2013**, *87*, 094304.
- [37] Barducci, A.; Bussi, G.; Parrinello, M. *Phys. Rev. Lett.* **2008**, *100*, 020603.
- [38] Kim, S. Y.; Perez, D.; Voter, A. F. *J. Chem. Phys.* **2013**, *139*, 144110.
- [39] Miron, R. A.; Fichthorn, K. A. *Phys. Rev. Lett.* **2004**, *93*, 128301.
- [40] Kim, W. K.; Falk, M. L. *J. Chem. Phys.* **2014**, *140*, 044107.
- [41] Tribello, G. A.; Ceriotti, M.; Parrinello, M. *Proc. Natl. Acad. Sci. USA* **2010**, *107*, 17509–17514.
- [42] Plimpton, S. *J. Comput. Phys.* **1995**, *117*, 1–19.
- [43] Fiorin, G.; Klein, M. L.; Hémin, J. *Mol. Phys.* **2013**, *111*, 3345–3362.
- [44] Bussi, G.; Parrinello, M. *Phys. Rev. E* **2007**, *75*, 056707.
- [45] Boisvert, G.; Lewis, L. J. *Phys. Rev. B* **1997**, *56*, 7643–7655.
- [46] Daw, M. S.; Baskes, M. I. *Phys. Rev. B* **1984**, *29*, 6443–6453.
- [47] Foiles, S. M.; Baskes, M. I.; Daw, M. S. *Phys. Rev. B* **1986**, *33*, 7983–7991.
- [48] Henkelman, G.; Uberuaga, B. P.; Jónsson, H. *J. Chem. Phys.* **2000**, *113*, 9901–9904.

- [49] Guo, Z.; Thirumalai, D. *J. Mol. Biol.* **1996**, *263*, 323–343.
- [50] Bengaard, H. S.; Nørskov, J. K.; Sehested, J.; Clausen, B. S.; Nielsen, L. P.; Molenbroek, A.; Rostrup-Nielsen, J. R. *J. Catal.* **2002**, *209*, 365–384.
- [51] Mueller, J. E.; van Duin, A. C. T.; Goddard, W. A., III *J. Phys. Chem. C* **2010**, *114*, 5675–5685.
- [52] Liu, B.; Lusk, M. T.; Ely, J. F. *Surf. Sci.* **2012**, *606*, 615–623.
- [53] Shibuta, Y.; Arifin, R.; Shimamura, K.; Oguri, T.; Shimojo, F.; Yamaguchi, S. *Chem. Phys. Lett.* **2013**, *565*, 92–97.
- [54] Somers, W.; Bogaerts, A.; van Duin, A. C. T.; Neyts, E. C. *J. Phys. Chem. C* **2012**, *116*, 20958–20965.
- [55] Somers, W.; Bogaerts, A.; van Duin, A. C. T.; Huygh, S.; Bal, K. M.; Neyts, E. C. *Catal. Today* **2013**, *211*, 131–136.
- [56] van Duin, A. C. T.; Dasgupta, S.; Lorant, F.; Goddard, W. A., III *J. Phys. Chem. A* **2001**, *105*, 9396–9409.
- [57] Aktulga, H. M.; Fogarty, J. C.; Pandit, S. A.; Grama, A. Y. *Parallel Comput.* **2012**, *38*, 245–259.
- [58] Mueller, J. E.; van Duin, A. C. T.; Goddard, W. A., III *J. Phys. Chem. C* **2010**, *114*, 4939–4949.
- [59] Rappé, A. K.; Goddard, W. A., III *J. Phys. Chem.* **1991**, *95*, 3358–3363.
- [60] Mueller, J. E.; van Duin, A. C. T.; Goddard, W. A., III *J. Phys. Chem. C* **2009**, *113*, 20290–20306.