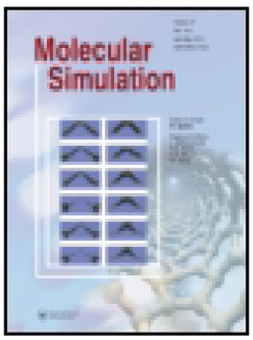

Molecular Simulation

ISSN: 0892-7022 (Print) 1029-0435 (Online) Journal homepage: https://www.tandfonline.com/loi/gmos20

Multiply accelerated ReaxFF molecular dynamics: coupling parallel replica dynamics with collective variable hyper dynamics

Karthik Ganeshan, Md. Jamil Hossain & Adri C. T. van Duin

To cite this article: Karthik Ganeshan, Md. Jamil Hossain & Adri C. T. van Duin (2019): Multiply accelerated ReaxFF molecular dynamics: coupling parallel replica dynamics with collective variable hyper dynamics, Molecular Simulation, DOI: 10.1080/08927022.2019.1646911

To link to this article: https://doi.org/10.1080/08927022.2019.1646911

Published online: 01 Aug 2019.

MOLECULAR SIMULATION
https://doi.org/10.1080/08927022.2019.1646911

Taylor & Francis
Taylor & Francis Group

Multiply accelerated ReaxFF molecular dynamics: coupling parallel replica dynamics with collective variable hyper dynamics

Karthik Ganeshan, Md. Jamil Hossain and Adri C. T. van Duin

Department of Mechanical Engineering, The Pennsylvania State University, University Park, PA, USA

To tackle the time scales required to study complex chemical reactions, methods performing accelerated molecular dynamics are necessary even with the recent advancement in high-performance computing. A number of different acceleration techniques are available. Here we explore potential synergies between two popular acceleration methods – Parallel Replica Dynamics (PRD) and Collective Variable Hyperdynamics (CVHD), by analysing the speedup obtained for the pyrolysis of n-dodecane. We observe that PRD + CVHD provides additional speedup to CVHD by reaching the required time scales for the reaction at an earlier wall-clock time. Although some speedup is obtained with the additional replicas, we found that the effectiveness of the inclusion of PRD is depreciated for systems where there is a dramatic increase in reaction rates induced by CVHD. Similar observations were made in the simulation of ethylene-carbonate/Li system, which is inherently more reactive than pyrolysis, indicate that the speedup obtained via the combination of the two acceleration methods can be generalised to most practical chemical systems.

ARTICLE HISTORY
Received 19 April 2019
Accepted 15 July 2019

KEYWORDS
Parallel replica dynamics; collective variable hyperdynamics; pyrolysis; battery electrolytes; ReaxFF

1. Introduction

While atomistic-scale simulation tools have proven to be greatly useful in understanding and developing new materials, a choice is often required to be made on the simulation method to be used based on the required chemical accuracy and computational speed [1]. *Ab initio* methods such as Density Functional Theory (DFT) and the semi-empirical Tight-Binding DFT (DFTB) rely on obtaining the interaction potentials between the atoms using approximations to the equations of motion derived from Quantum Mechanics (QM). The evaluation of the partial or full electronic structure of each atom allows these methods to be the closest in obtaining chemical accuracy; but the requirement to re-evaluate the potentials at every timestep for dynamics makes the associated dynamical method, termed *Ab initio* Molecular Dynamics (AIMD), often too computationally expensive for practical applications. Classical non-reactive molecular dynamics is a method capable of simulating several thousand atoms efficiently. This is because the dynamics is obtained using Newton's laws of motion and the interatomic potentials are evaluated using efficient empirically optimised functions. However, the inability to form or break bonds prevent the method from being useful in simulating a reactive environment. A bridge between the two classes of methods is provided by reactive MD tools, such as ReaxFF [2], COMB [3], AIREBO [4], and MEAM [5]. Though these methods retain several similarities to nonreactive MD in using the classical equations of motion, the interaction potentials between the atoms are designed to allow bond breakage and bond formation by including parameters such as the bond order. These interactions are often trained against QM-data, or experiments when available, so that they can mimic the results of QM-based methods at a fraction of the cost while being only about 2 orders of magnitude more expensive than non-reactive MD for dynamics. Using the correct training set for the given system, these methods can reach very close to the chemical accuracy that can be obtained through DFT. Therefore, reactive MD can be used for efficiently capturing both dynamic and chemical events for large systems at an acceptable efficiency. However, there is one limitation that is shared by all the above-mentioned methods that prevent long simulations – the largest time step to retain stability of the scheme. This issue arises from the requirement to resolve the vibrational modes of bonds between atoms. Typical time steps in ReaxFF are less than a femtosecond to ensure convergence of energy and prevent unphysical de/bonding interactions. Although higher order time stepping methods, such as Runge–Kutta (RK4), would generally help increase the time step while retaining stability, the short timescales in the vibrational modes prevent it from providing a practical advantage [6].

Figure 1 shows the acceleration methods currently implemented in ReaxFF, that help resolve the issue of extended residence time in a potential basin. GPU-enabled algorithms [7,8] attempt to improve the per-timestep simulation time by utilising the massively parallel architecture of GPUs, providing a 6× to 16× speedup compared to CPU-based algorithms. Although this does not tackle the time step issue, it enables ReaxFF simulations to perform at close-to non-reactive MD speeds. Parallel Replica Dynamics (PRD) [9–12] is another method that takes advantage of the parallelism. Instead of running a single simulation over multiple processors, the system is replicated over each processor

CONTACT Adri C. T. van Duin acv13@psu.edu

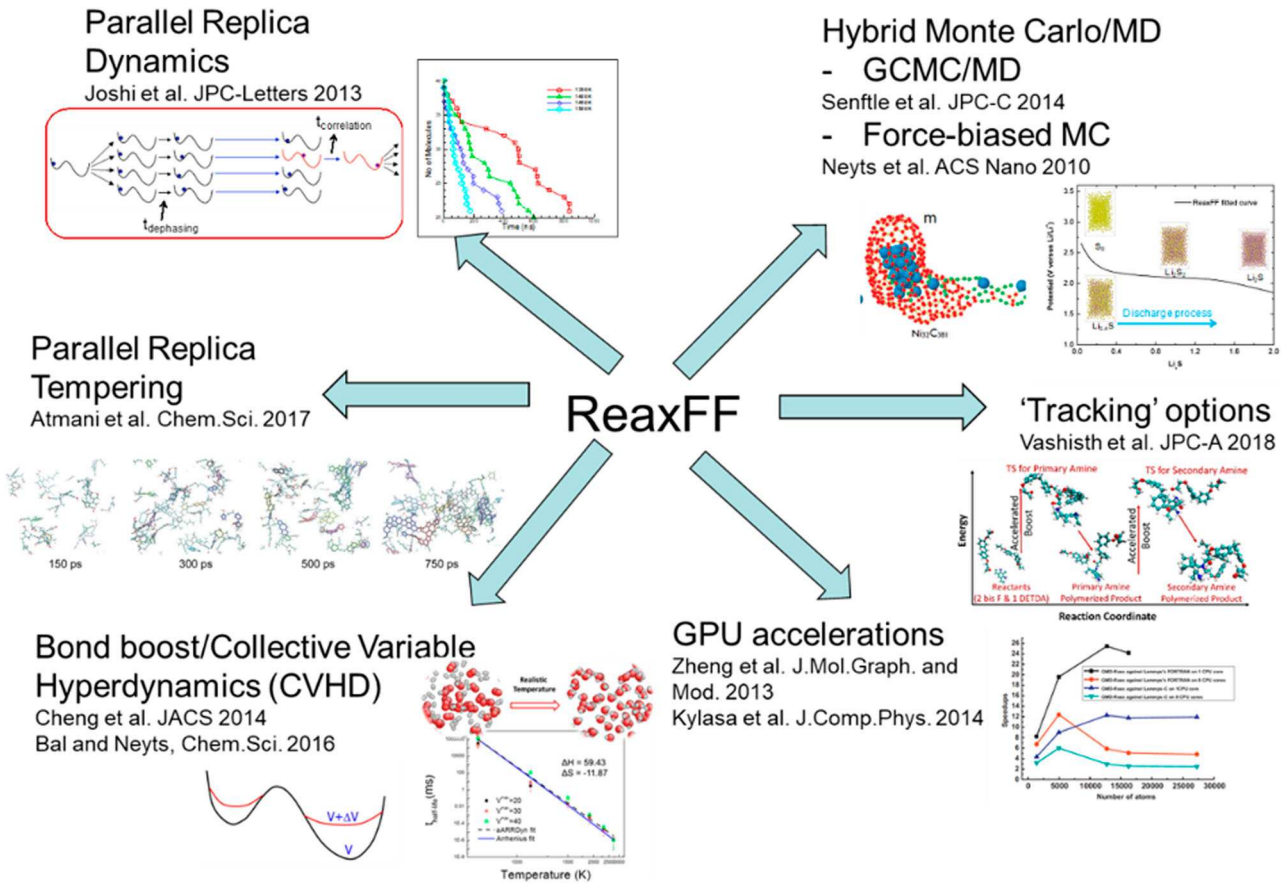

Figure 1. (Colour online) Various acceleration methods available in the ReaxFF reactive molecular dynamics environment.

such that each trajectory samples the potential basin until one observes a transition, providing a speedup directly proportional to the number of replicas used. Other methods rely on modifying the dynamics of the system to obtain the acceleration. In Temperature Accelerated Dynamics (TAD) the simulation is performed at an elevated temperature to allow quick escape of the system from the potential basin, and the correct transitions corresponding to the lower temperature are picked [13]. As at higher temperatures, entropically favoured reactions dominate over the reaction pathways corresponding to lower temperatures, caution must be exercised in ensuring only the feasible paths are picked. Methods such as Hyperdynamics [14], Metadynamics [15] and Collective-Variable Hyperdynamics (CVHD) [16], modify the free energy landscape by providing a bias potential, which fills the energy minima and decreases the energy barrier for the transitions. Force Biased Monte Carlo [17] and Grand Canonical Monte Carlo (GCMC/MD) [18] use Monte Carlo techniques to enable the system to reach a global minima faster. However, these methods may not capture the correct dynamics associated with transition. Other methods include Bond Boost [19], where a potential bias, and Tracking [20], where an additional force, is added to the bonds of labelled atoms to accelerate the dynamics.

In this work, we explore the combination of PRD and CVHD due to the mutual exclusivity in the acceleration method and generalisability to any reactive system. A descriptive introduction to the two methods and their combination is provided in the subsequent section, followed by the performance comparison against standard MD for n-dodecane pyrolysis.

2. Methods used

2.1. Parallel replica dynamics (PRD)

PRD takes advantage of the assumed ergodicity of the molecular dynamics process and the availability of several computer cores. For an ergodic process, the time average of the process in the simulation is the same as the ensemble average. Therefore, averaging over a large time domain is replaced by time averaging over multiple short simulations with different initial conditions. When the system is replicated amongst M replicas, it enables the overall simulation to sample M different escape routes out of the potential basin, c.f. Figure 2. As a consequence, the trajectory that leads the system out of the potential basin can be identified in a shorter wall-time. The simulation procedure for PRD are as follows [9,21]:

(1) Broadcast the configuration and minimise it over M replicas
(2) Initialise each replica with a unique random velocity distribution corresponding to the temperature
(3) Ensure each trajectory corresponds to a quasi-stationary state (dephase)
(4) MD run and check for any transitions

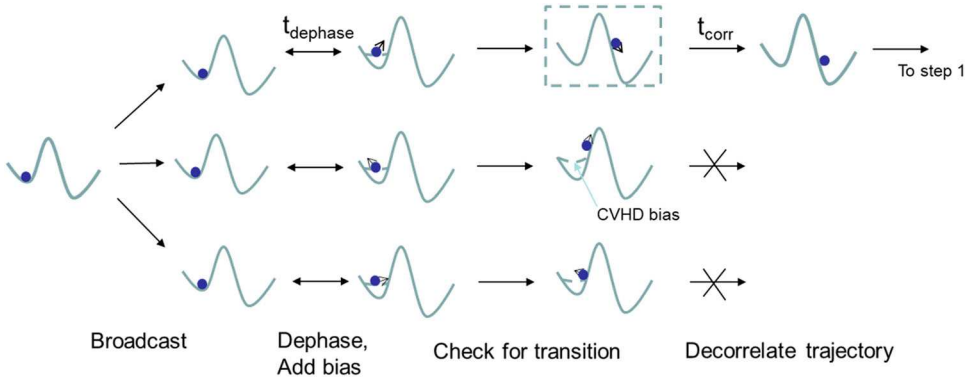

Figure 2 (Colour online) Schematic of the implementation of PRD coupled with CVHD under a representative free energy surface (FES). The arrows indicate the direction the system is moving towards, over the FES.

(5) Decorrelate the trajectory with respect to the transition state
(6) Purge all replicas and go to step 1

A trajectory is considered to contain a 'successful event' when it has crossed the chosen potential barrier. In a reactive environment, any dominant reaction leading to a change in bonds can be considered as a successful event [10]. In some cases, the first crossing of the potential barrier may lead to a series of secondary reactions. These secondary events must not be considered as an independent successful event due to their underlying correlation to the primary event.

The checks for correlation (step 5) and dephasing (step 3) are included to ensure no error is introduced by PRD and that the mutual independence of the replicas is maintained. To broadcast the configuration with the successful event to another set of replicas and initialise with a different set of random velocities, it is required that the configuration is not correlated to its earlier transition state and any correlated secondary event must be completed. For all MD simulations, two trajectories with a small difference in initial conditions become uncorrelated after a certain period of simulated time due to the Lyapunov instability [6]. Let τ_{corr} correspond to the time for which the system continues to be correlated to the state that crossed the potential barrier and $\Delta t_{MD,\text{corr}}$ be the MD time after the potential barrier has been crossed. Ensuring that $\Delta t_{MD,\text{corr}} \geq \tau_{\text{corr}}$ guarantees that the reinitialisation of the configuration with a different set of velocities will not be correlated to the earlier transition state. Although τ_{corr} is system dependent, it is generally ~1 ps or a few Einstein vibrational periods [9]. Another criterion is that we sample only the quasi-stationary distribution during PRD. Therefore, the trajectories that leave the potential basin at time $t \leq \tau_{\text{dephase}}$ are purged and then re-initialised with different random velocities. The time of the purged simulations is not added to the effective PRD time. The mathematical formulation of PRD, including the restrictions imposed by the dephasing and correlation criteria, can be found in references [11,12].

The total simulated time is a sum of the individual time simulated by each replica, i.e.

$$t_{\text{total}} = \sum_{i}^{M} t_i, \tag{1}$$

where t_i is the time simulated by the ith replica. This is used instead of $t_{\text{eff}} = Mt_1$ *to* account for performance differences between the processors. The acceleration over standard MD is given by [21]:

$$\eta_{PRD} = M((1 + f_Q)(1 + 2Mk_g\tau_{\text{corr}}))^{-1}. \tag{2}$$

Here, f_Q is the cost associated with checking for a successful event, per time step, and k_g is the transition rate for the observed reaction. The inverse relation to k_g indicates that the most acceleration is obtained only for rare events, i.e. $k_g \rightarrow 0$. Although it is not a requirement of the method, the largest acceleration from PRD can be obtained when each replica is run on a single processor. Therefore, standard PRD quickly loses its effectiveness to methods that vary the potential energy surface, such as hyperdynamics, when using methods like AIMD or reactive MD simulations on very large systems ($N_{\text{atoms}} \gg 10^3$).

2.2. Collective variable hyperdynamics (CVHD)

Collective variable hyperdynamics is a combination of conventional metadynamics and hyperdynamics. Metadynamics [15] is a technique that is capable of calculating free-energy landscapes by using a history-dependent bias potential. The bias potential is dependent on a small number of collective variables (CVs) that describe the system. The bias potential is applied to the potential energy surface, which fills the energy minima and consequently lowers the reaction energy barrier and thus, shortens the waiting time between reactions. On the other hand, hyperdynamics is an accelerated MD technique which is similar to metadynamics in the sense that it is also based on the application of a bias potential but in hyperdynamics this bias potential is removed on-the-fly during the transition state. The free energy information is lost by the unbiasing but the hypertime can be recovered from the simulation. The hypertime is the time required to get to a similar state in an unbiased simulation. The collective variable-driven hyperdynamics (CVHD) method [16] is a self-learning variant of the hyperdynamics algorithm that combines the hyperdynamics with CV-based feature of metadynamics in which an appropriate history-dependent bias potential can slowly be grown on the fly during the simulation to track the long time-scale evolution of the system.

The choice of an appropriate collective variable (CV) is essential for a CVHD simulation. The CV is based on a set of local system properties or degrees of freedom ($s_1, \ldots, s_N$), which collects all the relevant degrees of freedom of the full slow system dynamics. Following the description from Bal and Neyts [16], for each local property s_i, a local distortion ($\chi_i = \chi(s_i)$) is defined and can have a value between 0 and 1, the value of 1 being at the transition state. The global distortion is defined as

$$\chi_t = \left(\sum_{i=1}^{N} \chi_i^p \right)^{\frac{1}{p}} \tag{3}$$

in which $p > 1$. The CV η is defined as

$$\eta = \frac{1}{2}(1 - \cos(\pi \chi_t^2)), \quad \text{if } \chi_t \leq 1 \tag{4a}$$

$$\eta = 1, \quad \text{if } \chi_t > 1 \tag{4b}$$

For reaction events involving bond breaking, the property s is considered as the stretch of a bond. For every bond pair i with length r_i, there are distances $r_i^{\min}$ and $r_i^{\max}$, which mark the start and end point of possible reactive events. If $r_i < r_i^{\min}$, the bond is not likely to dissociate soon, and is not biased. Whereas, if $r_i = r_i^{\max}$, the bond is about to dissociate. For $r_i^{\min} < r_i < r_i^{\max}$ the local distortion is defined as

$$\chi_i = \frac{r_i - r_i^{\min}}{r_i^{\max} - r_i^{\min}} \tag{5}$$

The bias potential ΔV which is deposited at intervals τ_G is defined as a function of the CV η as a Gaussian shaped hill with width δ and height w_k as follows

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

The bias potential being dependent on the CV makes it history dependent. Gaussian shaped hills are added until a transition state is reached. Similar to the choice of the CV, choosing a suitable bias is hard and error-prone.

Using a bond length based CV, Bal and Neyts [22] applied the CVHD method in the pyrolysis and oxidation of n-dodecane at low temperatures and reached time scales which are unattainable by conventional molecular dynamics simulations.

2.3. CVHD coupled with PRD

The standard MD in each PRD replica can be replaced with an individually accelerated MD, such as CVHD or hyperdynamics, to obtain additional speedup per timestep. Such a combination of PRD and Hyperdynamics has been evaluated by Voter and Germann for epitaxial growth of Ag layer [23]. The combination multiplies the effective increase in time of PRD to that of hyperdynamics, as described in the following equation [24], borrowing the definition of the bias potential from Equation (6).

$$t_{\text{eff}} = \sum_{i}^{M} t_i = \sum_{i}^{M} \sum_{j}^{N} \Delta t_{MD,i} \exp[\Delta V(r_i(t_j))/k_b T]. \tag{7}$$

It must be noted that the inclusion of hyperdynamics increases the rate of reaction observed in PRD, i.e. k_g has to be replaced by $\eta_{HD} k_g$ where η_{HD} is the boost provided by hyperdynamics. This increased rate of reaction diminishes the acceleration provided by PRD because of its inverse relation to k_g, as implied by Equation (2). Therefore, the bias potential must be appropriately matched to maintain reasonable benefits in using the coupling to PRD.

Here, we explore the combination of PRD with CVHD to eliminate the requirement to have a predefined boost potential. While both the methods reset after a transition has occurred, it is not required that the same transition is targeted. Therefore, a transition with a higher barrier can be targeted by PRD while CVHD is permitted to reset after transitions that occur at lower barriers. To provide an example, a large barrier transition could be a chemical reaction and a small barrier transition could be one associated with an internal rotational barrier for a partially conjugated hydrocarbon (e.g. butadiene). With this separation of event checks, we can utilise a lower CVHD boost potential to ensure that the smaller barriers are not bypassed while maintaining reasonable effectiveness of PRD. Figure 2 shows a schematic of a PRD + CVHD simulation under a representative free energy surface.

At this juncture, it is important to recognise the key difference between the two acceleration techniques. By imposing a physically consistent bias potential formulation, CVHD directly alters the interatomic interactions of the system when compared to standard MD. Whereas, PRD must be perceived more as a sampling technique rather than a hyperdynamics-like acceleration tool. Each replica is an independent MD simulation with a small change in initial condition and is imposed with no additional constraints. The dephasing step ensures that the initial condition is not a transition state. Hence, the physics within each replica would be consistent with standard MD and parameters such as initiation time of a reaction, reaction rates, stabilisation of reaction products, etc. remain unaltered. The acceleration in PRD, however, is the effective decrease in the waiting time required to observe the successful transition. The probability of waiting time for the primary transition is given by [23]

$$p(t_{\text{total}}) = k_g \exp(-k_g t_{\text{total}}), \tag{8}$$

where t_{total} is evaluated using Equation (1) and k_g is the effective rate constant of the primary reaction. With the increase in the number of replicas, the probability of waiting time $p(t_{\text{total}})$ decreases exponentially, providing a sense of acceleration. Therefore, the acceleration in PRD works by sampling the phase space and targeting the trajectory that completes a transition first. Similarly, in CVHD coupled with PRD (PRD + CVHD), each replica must be treated as an independent CVHD simulation with k_g being replaced with the CVHD-accelerated rate constant $k_{g,\text{eff}}$, and t_{total} replaced with t_{eff} from Equation (7). While the physics simulated in PRD + CVHD would be similar to a CVHD-only simulation, there is an acceleration obtained in adding PRD to CVHD because the sampling of the phase space is performed after each rare event. As we know the reactions predicted by CVHD is consistent with standard MD [22], we can conclude that the reaction

products obtained between standard MD, PRD, CVHD and PRD + CVHD would be consistent.

While both the methods contribute towards the acceleration of the primary reaction, the behaviour is altered when dealing with secondary reactions and becomes highly dependent on how the PRD event checks are performed. In CVHD, the bias potential is reset after the primary reaction is complete, and the collective variables are slowly deposited again without stopping the MD simulation. This sudden drop in acceleration often translates to a disparity between the CPU time taken and the simulated MD time when comparing primary and secondary reactions. For example, in the pyrolysis simulation presented in the subsequent section, the primary reaction takes only 8 times the CPU time taken by the secondary reactions although the physical time (or hypertime) required for the primary reaction is longer by 3 orders of magnitude. In PRD, the acceleration is provided only until the targeted reaction is complete. This limitation arises from the requirement that the trajectory must be made uncorrelated to its history before broadcasting it to another set of replicas, while the replicas that did not observe a successful reaction are purged. Therefore, if the secondary reactions are not targeted in the event check, the PRD component of the simulation drops to a single replica that provides no further acceleration to the secondary reactions for a period of τ_{corr} (note that the time scales of the secondary reactions must be lesser than τ_{corr}). Consequently, in the method of PRD + CVHD, the acceleration experienced by the primary and secondary reactions is governed by the individual behaviour of PRD and CVHD, although both the reactions experience an acceleration over standard MD. In most cases, the secondary reactions are several in number and/or are unknown while taking place at time scales several orders of magnitude lower than the primary reactions. Thus, in the subsequent implementation and examples of its use, the PRD event check is targeted only at the primary reaction that has a high barrier.

3. Simulation setup

The pyrolysis of n-dodecane was considered to study the effectiveness of PRD + CVHD in comparison to standard CVHD, as the process follows first order kinetics [10] and a reference CVHD parameters can be obtained from Bal and Neyts [22]. Here, 24 n-dodecane molecules were considered in a 50 Å × 50 Å × 50 Å periodic box, at a simulation temperature of 1200 and 1500 K, with a density of 0.054 kg dm^{-3}. The parameters for CVHD from [22] were retained – a Gaussian height of 0.25 kcal mol^{-1} was used with the gaussian half width of 0.025. The frequency of deposition of 0.2 ps and a 1 ps waiting period to check for an event was used. A local bond-breaking event check was implemented, based on which the boosted potential in Equation (6) is altered. To maintain similarities with the reference [22], the ReaxFF force field parameters for C/H/O described by Chenoweth et al. [2] was used. A monitoring script was written in Bash to implement PRD by initiating and terminating replica MD runs based on event detection. The trajectory where one dodecane molecule completes a reaction is considered to contain the successful event for PRD. A sample pyrolysis reaction is shown in Figure 3, where one of the 24 n-dodecane molecules decomposes into C_4H_9 and C_8H_{17} radicals.

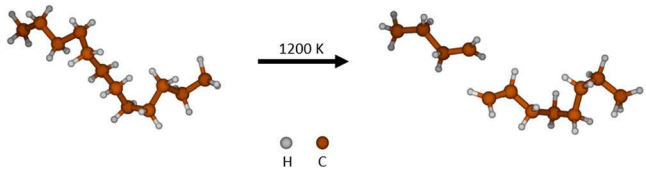

Figure 3. (Colour online) One of the 24 n-dodecane molecules decomposing into C4H9 and C8H17 radicals.

The check for the transition is made by setting a bond order cutoff of 0.3 and counting the number of dodecane molecules based on the connectivity of each atom. Based on the analysis of the velocity autocorrelation of heptane [10], a τ_{corr} of 10 ps and $\tau_{dephase}$ of 5 ps was chosen. A higher τ_{corr} was opted to allow sufficient time for the secondary reactions to complete. The PRD algorithm mentioned in section 2.1 was slightly modified – the minimisation of the system before checking for a transition was replaced by a primary check for the event followed by a confirmation check after $\Delta t_{MD} = \tau_{buffer}$. This eliminates any false positives generated by quickly reversible reactions without the requirement of pausing the MD run. Therefore, the PRD script does not interfere with the MD run until the reaction is complete, confirmed and passed the correlation time. Here, a τ_{buffer} of 2 ps is considered. As the event detection is computationally cheap, a moderately frequent event check is carried out at every 0.1 ps to prevent extended runtimes after a successful transition is made. The PRD implementation, including the event checks, are external to the CVHD/MD run; which eliminates any additional computational overhead experienced by the simulation. However, there is a caveat – one additional processor that cannot be involved in the MD simulation is required to run the PRD monitoring and event check programmes. As the process is serial and generally not computationally taxing, it can be accommodated on a local machine that is capable of initiating MD runs on another batch of processors. All MD/CVHD simulations were performed in ReaxFF within the ADF suite, distributed by SCM (see www.scm.com). ADF version later than 2017-r60492 is required to ensure each replica is initiated with a different random velocity.

4. Results and discussion

Figure 4(a,b) shows the difference in the time simulated by standard MD, PRD, CVHD and PRD + CVHD. Here, the initial stage of n-dodecane pyrolysis simulated by a 4-replica CVHD simulation is considered for illustrative purposes. Similar to the observation made by Voter and Germann [23], we see that over a relatively large number of steps, CVHD outperforms MD and standard PRD by simulating a time scale that is several orders of magnitude longer. As the acceleration provided by CVHD slowly builds up with time, a larger speedup is obtained using PRD during very early stages of the simulation, as indicated by Figure 4(b). It must be noted that in this work, we use dynamically biased CVHD instead of statically biased CVHD. In the case with statically biased CVHD we would expect a constant boost for every timestep, providing a linear increase in simulated time, similar to PRD. However, dynamic biasing of the system allows the desirable

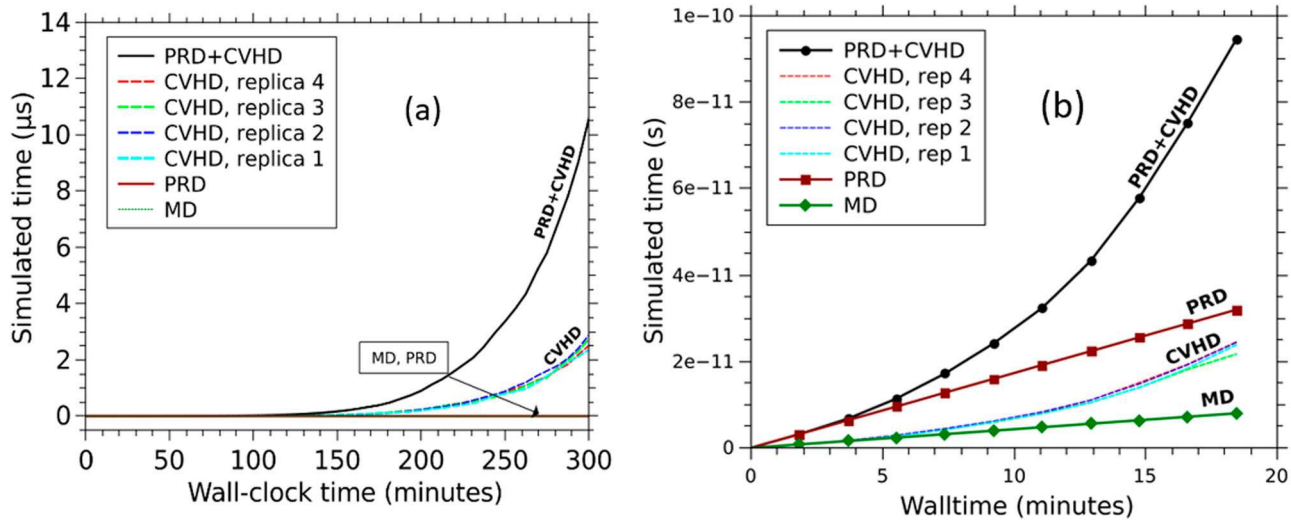

Figure 4. (Colour online) The total simulated time vs wallclock time for different methods. (a) Time simulated at a later stage in simulation, where CVHD >> PRD (b) During the initial phase of the simulation, where PRD > CVHD.

self-learning mechanism of biasing based on the requirements of the system.

Although the above analysis would indicate that using more replicas would proportionally increase efficiency, it is important to recall Equation (2). The inverse relation to M and k_g suggests that there is depreciating gains for large number of replicas, for long simulations that involve several transitions. To evaluate the performance, a PRD + CVHD simulation was set up with 10, 20, 40, 70 and 100 replicas, with the results averaged over 3 sample runs. Each replica simulation was performed on a single Intel Xeon E5 2.5 GHz processor. A standard CVHD simulation was performed to serve as a point of reference, with the time averaged over 10 samples. Figure 5(a,b) shows the speedup obtained by the inclusion of PRD over standard CVHD, evaluated by using the relation

$$\eta_{\text{eff}} = \frac{\text{walltime for 1 replica}}{\text{walltime for } M \text{ replicas}}. \tag{9}$$

The theoretical estimate of speedup over CVHD is obtained by following relation.

$$\eta_{\text{eff}} = \frac{\eta_{PRD+CVHD}}{\eta_{CVHD}} = M\frac{1 + 2k_{g,\text{eff}}\tau_{\text{corr}}}{1 + 2Mk_{g,\text{eff}}\tau_{\text{corr}}} \tag{10}$$

Here, $k_{g,\text{eff}}$ is the effective reaction rate observed by PRD due to the acceleration provided by CVHD. From the analysis conducted by Chowdhury et al. for n-Dodecane pyrolysis [25], the expected rate constant for the reaction at 1200 K is $\approx$ 6.1E + 04 s^{-1} and $k_g\tau_{\text{corr}}$ $\approx$ 6.1E-07. However, due to the acceleration provided by CVHD, a value of $k_{g,\text{eff}}\tau_{\text{corr}} = 0.075$ was used for the evaluation of the estimated speedup at 1200 K. We can immediately see that PRD doesn't scale linearly with the number of replicas; as a direct consequence of the inverse relation to $Mk_{g,\text{eff}}$ predicted by Equation (2). Therefore, only an incremental speedup can be expected when larger number of replicas are used. When conducting a similar experiment at 1500 K, we observe that the speedup obtained by PRD is

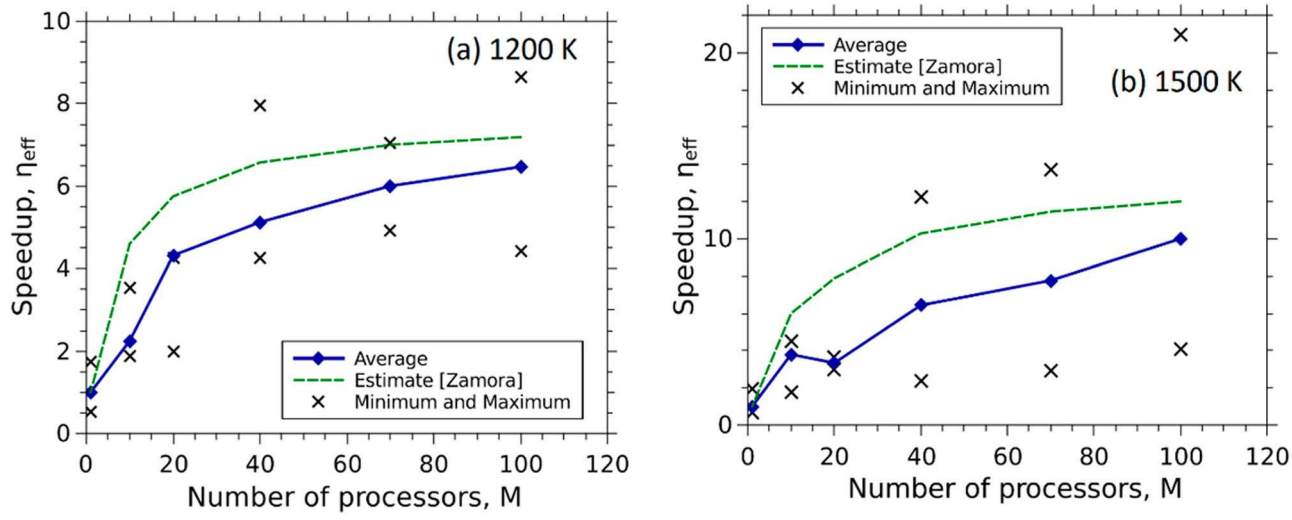

Figure 5. (Colour online) The speedup in simulation obtained by the combination of PRD and CVHD, compared against standard CVHD. (a) Simulation temperature of 1200 K (b) Simulation temperature of 1500 K.

Table 1. Comparison of effective simulated time and hypertime of the trajectory with a successful transition across different number of replicas.

Number of replicas used	Average simulated time (seconds)	Successful trajectory hypertime (seconds)
1	1.71e-5	1.71e-5
10	2.27e-5	2.59e-6
20	2.55e-6	9.55e-8
40	7.582e-6	2.49e-7
70	1.36e-5	2.03e-7
100	3.33e-5	3.98e-7

larger at 1500 K than at 1200 K, as indicated by Figure 5(b). We believe this could be attributed again to the boost provided by CVHD. As the reaction rate is inherently higher at 1500 K, the assistance required from CVHD is lowered. Therefore, the effective rate of the reaction is lower than that of the boosted 1200 K simulation. The perceived lower reaction rate at 1500 K contributes to the increased effectiveness of PRD. It is important to note that the data presented in Figure 5(a,b) correspond to the speedup obtained by using PRD in comparison to standard CVHD. Caution must be exercised when following this trend for different temperatures, especially lower temperatures such as 1000 K and 700 K. The effective increase in the rate of the reaction by CVHD is governed by the bias potential properties set a priori by the user. Hence, even at lower temperatures, the effectiveness of the combination of PRD and CVHD can be significant if the bias potential applied is moderated. Therefore, the trend of the effectiveness of the combination of the two methods must be perceived as a function of the effective rate of reaction of the system, rather than a function a particular system property. The total time required to complete the decomposition of three n-dodecane molecules is of the order of 2 μs. Hence, in non-accelerated MD, 2×10^{10} steps or 72 thousand hours would be required. Due to the randomness of the initialised velocity distribution in each system, there will be statistical noise in the results obtained for performance comparison. However, many more samples can be used to obtain a more accurate representation of the average performance.

Table 1 shows the cumulative net simulated time across the 3 stages of the PRD simulation. We see that the effective simulated time across multiple number of replicas mostly remains within the same order of magnitude as that of a single CVHD simulation. Comparing the values to the hypertime observed by the successful trajectory in PRD, we can see that the successful trajectory experiences a lower hypertime corresponding to the speedup provided by PRD. Therefore, it can

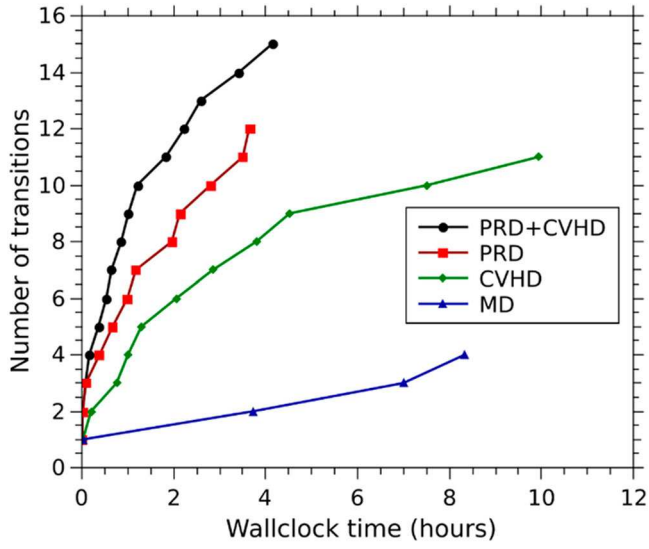

Figure 7. (Colour online) Number of reactive transitions observed in the system vs. wall-clock time.

be inferred that the acceleration provided by CVHD is lowered with coupled with PRD; and the rate of reaction captured is independent of the number of replicas, provided a reasonable boost potential is used in CVHD.

In addition to the pyrolysis of n-dodecane, eReaxFF MD [26] was used to perform a comparative study between MD, PRD and PRD + CVHD for the ring opening reaction of ethylene-carbonate (EC)/Li system. A system comprising of 201 EC molecules with 15 Li^+/explicit-electron pairs in a periodic cube of dimension 40 Å at density of 0.46 g cm^{-3} was simulated at 300 K using the NVT ensemble with Berendsen thermostat under a damping frequency of 100 fs. 10 replicas were used in both PRD and PRD + CVHD simulations and CVHD parameters – gaussian height of 0.5 kcal mol^{-1} with a half-width of 0.025 at a deposition frequency of 0.5 ps were considered for CVHD and PRD + CHD. For the comparison of accelerations, a complete ring opening reaction of the EC molecule was considered as a successful transition. Figure 6 illustrates the simulated ring opening reaction of the EC/Li system. Similar behaviour, as that of n-dodecane pyrolysis, is seen when comparing the performance between the different acceleration tools. Although the EC/Li system is more reactive (key barrier around 10 kcal mol^{-1}) in comparison to dodecane pyrolysis (key barrier around 60 kcal mol^{-1} in the presence of radicals; around 80 kcal mol^{-1} for homolytic scission), the results shown in Figure 7 provide evidence to generalise that the

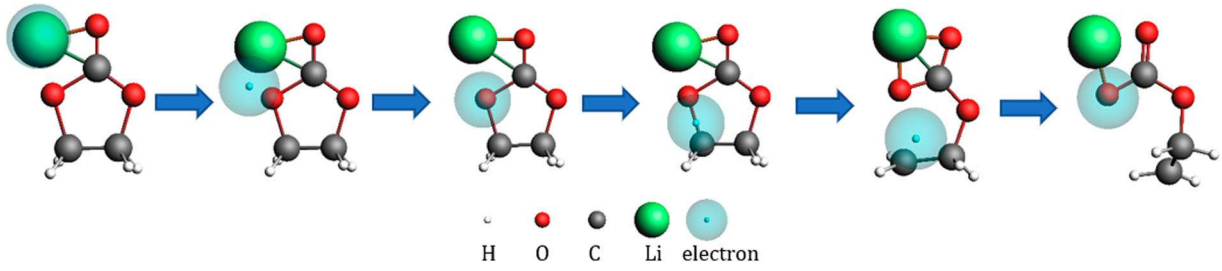

Figure 6. (Colour online) Ring opening reaction of ethylene carbonate induced by Li.

acceleration provided by the combination of PRD and CVHD is greater than the methods considered independently.

5. Conclusion

We have analysed the performance of the combination of PRD and CVHD for the pyrolysis of n-dodecane. Although the combination of the two provides a method capable of simulating milliseconds to seconds time period in a reasonable wall-clock time, there is a decrease in the speedup obtained in PRD as the rate of reaction increases, especially because of the boost provided by CVHD. The effective simulated time across different combinations of PRD added to CVHD lies within the same order of magnitude. This indicates that the observed reaction mechanism is independent of the number of replicas chosen and using larger number of replicas only save wall-clock time required. Similar speedup obtained for the system of ethylene carbonate/Li indicates that the combined method can be advantageous even for a moderately more reactive system.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

KG and ACTvD acknowledge funding from the Fluid Interface Reactions, Structures and Transport (FIRST) and the Multi-Scale Fluid-Solid Interactions in Architected and Natural Materials (MUSE) Centers, which are Energy Frontier Research Centers funded by the U.S. Department of Energy (DOE), Office of Science, Basic Energy Sciences and from NSF [CDS&E 1807740]. MJH and ACTvD acknowledge funding from the project sponsored by the Army Research Laboratory [Cooperative Agreement Number W911NF-12-2-0023].

References

[1] Ohno K, Esfarjani K, Kawazoe Y. Computational materials science: from ab initio to Monte Carlo methods. Heidelberg: Springer-Verlag; 2018.
[2] Chenoweth K, Van Duin ACT, Goddard WA. ReaxFF reactive force field for molecular dynamics simulations of hydrocarbon oxidation. J Phys Chem A. 2008;112:1040–1053.
[3] Shan TR, Devine BD, Kemper TW, et al. Charge-optimized many-body potential for the hafnium/hafnium oxide system. Phys Rev B Condens Matter Mater Phys. 2010;81:1–12.
[4] Stuart SJ, Tutein AB, Harrison JA. A reactive potential for hydrocarbons with intermolecular interactions. J Chem Phys. 2000;112:6472–6486.
[5] Baskes MI. Application of the embedded-atom method to covalent materials: a semiempirical potential for silicon. Phys Rev Lett. 1987;59:2666–2669.
[6] Frenkel D, Smit B, Tobochnik J, et al. Understanding molecular simulation. Comput Phys [Internet]. 1997;11:351. Available from: http://scitation.aip.org/content/aip/journal/cip/11/4/10.1063/1.4822570
[7] Kylasa SB, Aktulga HM, Grama AY. PuReMD-GPU: a reactive molecular dynamics simulation package for GPUs. J Comput Phys. 2014;272:343–359. DOI:10.1016/j.jcp.2014.04.035.
[8] Zheng M, Li X, Guo L. Algorithms of GPU-enabled reactive force field (ReaxFF) molecular dynamics. J Mol Graph Model. 2013;41:1–11. DOI:10.1016/j.jmgm.2013.02.001.
[9] Voter AF. Parallel replica method for dynamics of infrequent events. Phys Rev B Condens Matter Mater Phys. 1998;57:R13985–R13988.
[10] Joshi KL, Raman S, Van Duin ACT. Connectivity-based parallel replica dynamics for chemically reactive systems: from femtoseconds to microseconds. J Phys Chem Lett. 2013;4:3792–3797.
[11] Le Bris C, Lelièvre T, Luskin M, et al. A mathematical formalization of the parallel replica dynamics. Monte Carlo Methods Appl [Internet]. 2012;18:1–21. Available from: https://www.degruyter.com/view/j/mcma.2012.18.issue-2/mcma-2012-0003/mcma-2012-0003.xml
[12] Lelièvre T. Accelerated dynamics: mathematical foundations and algorithmic improvements. Eur Phys J Spec Top. 2015;224:2429–2444.
[13] Voter AF, Sørensen MR. Accelerating atomistic simulations of defect dynamics: hyperdynamics, parallel replica dynamics, and temperature-accelerated dynamics. MRS Proc [Internet]. 1998;538:427. Available from: http://journals.cambridge.org/abstract_S194642740028980X
[14] Voter AF. Hyperdynamics: accelerated molecular dynamics of infrequent events. Phys Rev Lett [Internet]. 1997;78:3908–3911. Available from: https://link.aps.org/doi/10.1103/PhysRevLett.78.3908
[15] Laio A, Parrinello M. Escaping free-energy minima. Proc Natl Acad Sci. 2002;99:12562–12566.
[16] Bal KM, Neyts EC. Merging metadynamics into hyperdynamics: accelerated molecular simulations reaching time scales from microseconds to seconds. J Chem Theory Comput. 2015;11:4545–4554.
[17] Neyts EC, Shibuta Y, Van Duin ACT, et al. Catalyzed growth of carbon nanotube with definable chirality by hybrid molecular dynamics-force biased Monte Carlo simulations. ACS Nano. 2010;4:6665–6672.
[18] Senftle TP, Van Duin ACT, Janik MJ. Determining in situ phases of a nanoparticle catalyst via grand canonical Monte Carlo simulations with the ReaxFF potential. Catal Commun. 2014;52:72–77. DOI:10.1016/j.catcom.2013.12.001.
[19] Cheng T, Jaramillo-Botero A, Goddard WA, et al. Adaptive accelerated ReaxFF reactive dynamics with validation from simulating hydrogen combustion. J Am Chem Soc. 2014;136:9434–9442.
[20] Vashisth A, Ashraf C, Zhang W, et al. Accelerated ReaxFF simulations for describing the reactive cross-linking of polymers. J Phys Chem A. 2018;122:6633–6642.
[21] Zamora RJ, Perez D, Martinez E, et al. Accelerated molecular dynamics methods in a massively parallel world. In: Handbook of materials modeling [Internet]. Cham: Springer International Publishing; 2018. p. 1–28. Available from: http://link.springer.com/10.1007/978-3-319-42913-7_25-1
[22] Bal KM, Neyts EC. Direct observation of realistic-temperature fuel combustion mechanisms in atomistic simulations. Chem Sci. 2016;7:5280–5286.
[23] Voter AF, Germann TC. Accelerating the dynamics of infrequent events: combining hyperdynamics and parallel replica dynamics to treat epitaxial layer growth. MRS Proc [Internet]. 1998;528:221. Available from: http://journals.cambridge.org/abstract_S1946427400258387.
[24] Voter AF, Montalenti F, Germann TC. Extending the time scale in atomistic simulation of materials. Annu Rev Mater Res. 2002;32:321–346.
[25] Ashraf C, Shabnam S, Jain A, et al. Pyrolysis of binary fuel mixtures at supercritical conditions: a ReaxFF molecular dynamics study. Fuel. 2019;235:194–207. DOI:10.1016/j.fuel.2018.07.077.
[26] Islam MM, Kolesov G, Verstraelen T, et al. EReaxFF: a pseudoclassical treatment of explicit electrons within reactive force field simulations. J Chem Theory Comput. 2016;12:3463–3472.