Accelerated Path Integral Methods for Atomistic Simulations at Ultra-low Temperatures

Felix Uhl,[1, a)] Dominik Marx,[1] and Michele Ceriotti[2]

[1)] *Lehrstuhl für Theoretische Chemie, Ruhr–Universität Bochum, 44780 Bochum, Germany*

[2)] *Laboratory of Computational Science and Modelling, Institute of Materials, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland*

(Dated: October 24, 2016)

Path integral methods provide a rigorous and systematically convergent framework to include the quantum mechanical nature of atomic nuclei in the evaluation of the equilibrium properties of molecules, liquids or solids at finite temperature. Such nuclear quantum effects often are already significant for light nuclei at room temperature, but become crucial at cryogenic temperatures such as those provided by superfluid helium as a solvent. Unfortunately, the cost of converged path integral simulations increases significantly upon lowering the temperature so that the computational burden of simulating matter at the typical superfluid helium temperatures becomes prohibitive. Here we investigate how accelerated path integral techniques based on colored noise generalized Langevin equations, in particular the so called PIGLET variant, perform in this extreme quantum regime using as an example the quasi rigid methane molecule and its highly fluxional protonated cousin, CH_5^+. We show that the PIGLET technique gives a speedup of two orders of magnitude in the evaluation of structural observables and quantum kinetic energy at ultralow temperatures. Moreover, we computed the spatial spread of the quantum nuclei in CH_4 to illustrate the limits of using such colored noise thermostats close to the many body quantum ground state.

I. INTRODUCTION: MOTIVATION AND BACKGROUND

Path integral molecular dynamics[1] is nowadays a well established simulation methodology in order to consider nuclear quantum effects (NQEs) at finite temperatures even in chemically complex systems consisting of flexible molecules such as hydrogen bonded clusters in the gas phase, aqueous solutions, molecular solids or molecular adsorbates on surfaces. This quantum statistical method is particularly appealing at temperatures where NQEs are already sizable, but do not yet override the gross picture as obtained in the framework of classical statistical mechanics, which is known as the quasi classical regime. A corollary of this statement is that such numerically discretized path integral methods become increasingly cumbersome in the deep quantum regime, i.e. upon approaching the quantum ground state. The reason is that, according to Trotter's theorem applied to path integration[2], the number P of discretization points, or (imaginary) time slices or (Trotter) beads, that turns the continuous functional integral into its usual so called primitive discretized version[1], must be increased appropriately the lower the temperature becomes[3,4]. Apart from using genuine ground state quantum simulation methods that exclude any thermal excitations, and exploit the fact of working at 0 K at their advantage[5], any finite temperature approach such as path integral based techniques must face the curse of discretization at ultralow temperatures.

But why would one need to perform quantum simulations of molecular systems at ultralow albeit not at zero temperature? A major motivation comes from matrix isolation spectroscopy experiments that are carried out using helium nanodroplets[6–8]. These superfluid quantum clusters provide experimentalists with an environment that interacts only rather weakly with the molecular solute species, which they easily trap as dopants in typical pick–up experiments, yet it cools the solutes rapidly down to temperatures on the order of 1 K. At such ultralow temperatures, the molecular solutes find themselves mostly in the quantum ground state, which dramatically simplifies spectra due to the absence of significant excitations which otherwise greatly congest data recorded at higher temperatures.

Several approaches have been established during the years in order to tackle the low temperature issue in the realm of numerical path integration. Pioneering bosonic path integral simulations[9] of the λ–transition in superfluid helium at about 2 K were enabled by using a numerical pair product action[9–11] instead of the usual single particle high temperature propagator. This is particularly useful if the species and their interactions are not too complex, helium atoms with pair potential interactions being the hallmark where $P \approx 40$ is already sufficient at $T \approx 1$ K. Alternatively, higher order Trotter approximants have been devised, reinvented, and tested since several decades[12–21] which converge faster than the usual primitive approximant at a given P. More recently, yet another class of approaches has been introduced which uses colored noise thermostatting in order to add approximately some quantum fluctuations while still using the primitive approximant to take care on whatever remains to converge to the exact result as P increases[22,23].

[a)]Electronic mail: felix.uhl@rub.de

In what follows, we will explore to what extend the so called "path integral generalized Langevin equation thermostat" (PIGLET) method[23] can be used in order to approach in an efficient manner temperatures on the order of 1 K; we refer the interested reader to Refs. 23–25 for the fundamentals of using colored noise (GLE) thermostatting in order to accelerate path integral convergence in the framework of molecular dynamics simulations. It is noted in passing that PIGLET has already been extensively applied to molecule–based problems in conjunction with interaction potentials of increasing complexity such as Lennard-Jones potentials, point-charge models, polarizable force fields, and even density functional based *ab initio* path integral[26–28] simulations as exposed for instance in Refs. 24,29. Reference path integral simulations will be carried out using the standard primitive approximant using up to $P \approx 16'000$ time slices within this brute force approach. Moreover, we will also highlight some of the advantages of the PIGLET approach with respect to well-established fourth-order path integral schemes, using in particular the Takahashi–Imada[14] or Lee–Broughton[15] action as a representative example.

The molecular models to be used for this benchmarking are CH_4 and CH_5^+ in vacuum. While methane is a quasi rigid molecule that undergoes small amplitude oscillations with respect to a well defined equilibrium structure, its protonated counterpart is considered to be an utmost fluxional molecule[30–32]. Protonated methane is unique in the sense that it possesses an unusually shallow potential energy surface which supports large amplitude motion, leading to so called full hydrogen scrambling even at low temperatures[30]. Thus, while the former is expected to be a less challenging test for methods that make use of the harmonic oscillator model in one way or another such as PIGLET does, the latter species can be expected to be a tough call. In order to be computationally efficient, we use the recently developed and validated POSflex force field[33] for CH_5^+ and construct an analogous potential for CH_4. This allowed us to systematically study the convergence properties of the different approaches in terms of energy components, structural parameters and nuclear delocalization at temperatures that are relevant to different types of experiments, namely from 100 K in ion traps to 20 K in some gas phase experiments down to 1.67 K in bulk superfluid helium.

II. THEORETICAL BACKGROUND

A. Path Integral Molecular Dynamics

The Hamiltonian for N nuclei in the Born–Oppenheimer approximation is given by

$$\mathcal{H} = \sum_{i=1}^{N} \frac{\hbar^2}{2m_i} \nabla_i^2 + V(\mathbf{q}), \qquad (1)$$

where $\mathbf{q} = \{\mathbf{q}_i\}$ denotes the set of all nuclear positions. The statistical mechanics of a system described by such Hamiltonian in the canonical ensemble at inverse temperature $\beta = 1/k_B T$ is described by the thermal density matrix, from which the canonical quantum mechanical partition function can be obtained as

$$\mathcal{Z} = \mathrm{Tr}\left(e^{-\beta \mathcal{H}}\right) = \int d\mathbf{q}\, \langle \mathbf{q} | e^{-\beta \mathcal{H}} | \mathbf{q} \rangle. \qquad (2)$$

To be able to compute this trace for general Born–Oppenheimer potentials, it is useful to factorize the Boltzmann operator using the Trotter theorem — effectively introducing P complete sets of position eigenstates, each of which leads to the evaluation of the off diagonal components of a Boltzmann operator at a higher temperature β/P. The factorized partition function can be evaluated by introducing a mixed basis consisting of momentum eigenfunctions for the kinetic part and position eigenfunctions for the potential term, and using a high temperature expansion for $e^{-\beta \mathcal{H}}$

$$\mathcal{Z} = \lim_{P \to \infty} \prod_{i=1}^{N} \left(\frac{m_i P}{2\pi \beta \hbar^2}\right)^{3P/2} \int \prod_{i=1}^{N} d\mathbf{q}_i^{(1)} .. d\mathbf{q}_i^{(P)}$$
$$\exp\left[-\frac{\beta}{P} \sum_{j=1}^{P} \left(\sum_{i=1}^{N} \frac{m_i \omega_P^2}{2} \left(\mathbf{q}_i^{(j+1)} - \mathbf{q}_i^{(j)}\right)^2 \right.\right.$$
$$\left.\left. + V\left(\{\mathbf{q}_i\}^{(j)}\right) \right) \right], \qquad (3)$$

where $\omega_P = P/\beta\hbar$, and cyclic boundary conditions $\mathbf{q}_i^{(j+P)} \equiv \mathbf{q}_i^{(P)}$ are implied. Truncating the expansion at a finite value of P introduces formally an error of the order of $(\beta/P)^2$ within this primitive approximant to the high temperature propagator. While this discretized partition function can be fully evaluated in terms of the positions of the P copies of the physical system using stochastic Monte Carlo methods, for the purpose of sampling one can introduce a set of fictitious classical momenta $\mathbf{p}$ conjugate to $\mathbf{q}$ with associated fictitious masses $m_i'^{(j)}$, leading to the usual expression for the path integral molecular dynamics Hamiltonian,

$$\mathcal{H}_P(\{\mathbf{q}\}, \{\mathbf{p}\}) = \sum_{j=1}^{P} \left[\sum_{i=1}^{N} \frac{\mathbf{p}_i^{(j)2}}{2m_i'^{(j)}} + \right.$$
$$\left. \sum_{i=1}^{N} \frac{m_i^{(j)} P^2}{2(\beta\hbar)^2} \left(\mathbf{q}_i^{(j+1)} - \mathbf{q}_i^{(j)}\right)^2 + V\left(\{\mathbf{q}_i\}^{(j)}\right) \right] \qquad (4)$$

which describes the movement of a classical ring polymer with a potential $V(\mathbf{q})$ acting on each of the P beads.

B. Integrating and Thermostatting PIMD

The path integral Hamiltonian (4) has gained a reputation for being associated with difficult sampling problems, due to the presence of a stiff harmonic term together with the physical potential. In fact, integration

and sampling can be made more manageable precisely by exploiting the large fraction of the free ring polymer normal modes that is associated with quasi harmonic vibrations that are much faster than those associated with V. Several techniques have been proposed to this end, involving scaling of the dynamical masses of different normal modes, or the so called staging transformation of ring polymer coordinates[34–36]. Here, for reasons that will become apparent further below, we will use physical masses for the integration of the PIMD, setting $m' = m$. Nevertheless, it is possible to use a time step comparable to the one that should be used in a classical molecular dynamics trajectory, by applying a multiple time stepping procedure[37], in which the free particle contribution to the full Hamiltonian (4),

$$\mathcal{H}_P^0(\{\mathbf{q}\}, \{\mathbf{p}\}) = \sum_{i=1}^{N} \sum_{k=1}^{P} \left(\frac{[\tilde{\mathbf{p}}_i^{(k)}]^2}{2m_i'^{(k)}} + \frac{1}{2} m_i^{(k)} \omega^{(k)2} [\tilde{\mathbf{q}}_i^{(k)}]^2 \right), \quad (5)$$

is evolved exactly and the forces arising due to the physical potential are applied in the outer loop. As shown, the propagation of $\mathcal{H}_P^0$ is best carried out by working in the normal mode representation of the ring polymer, that is

$$\tilde{\mathbf{p}}_i^{(k)} = \sum_{j=1}^{P} \mathbf{p}_i^{(j)} C_{jk} \quad \text{and} \quad \tilde{\mathbf{q}}_i^{(k)} = \sum_{j=1}^{P} \mathbf{q}_i^{(j)} C_{jk}, \quad (6)$$

where C_{jk} is an orthogonal transformation matrix whose elements are listed, for instance, in Ref. 22, and $\tilde{\mathbf{p}}^{(k)}$ is the normal mode associated with the ring polymer vibration of frequency $\omega^{(k)} = 2\omega_P \sin((k-1)\pi/P)$. In practice, the forward and backward transformations can be performed efficiently using fast Fourier transforms, which becomes crucial in cases involving a very large number of replicas.

Evolving the dynamics in normal modes coordinates is not only beneficial for the stability of the integrator. Having access to the different dynamical time scales also allows one to tackle the ergodicity problems that might affect this strongly anisotropic quasi harmonic problem. Thermostats can be applied in the normal modes representation, and their coupling parameters be tuned to obtain the most efficient sampling of ring polymer configurations using either deterministic[36] or stochastic[22] approaches. Following the latter avenue, this leads to the so called Path Integral Langevin Equation (PILE) thermostat when the underlying thermostat is based on white noise Langevin dynamics, and simply amounts to applying the transformation

$$\tilde{\mathbf{p}}_i^{(k)} \leftarrow e^{-(\Delta t/2)\gamma^{(k)}} \tilde{\mathbf{p}}_i^{(k)} + \sqrt{\frac{m_i'^{(k)}}{\beta P} \left(1 - e^{-(\Delta t)\gamma^{(k)}}\right)} \boldsymbol{\xi}_i^{(k)} \quad (7)$$

before and after the velocity–Verlet step of the integrator. In this equation, ξ denotes uncorrelated Gaussian random numbers, Δt is the integration time step, and $\gamma^{(k)}$

is the friction parameter that determines the coupling strength. It can be shown that the most efficient configurational sampling can be achieved by setting $\gamma^{(k)} = \omega^{(k)}$. A similar tuning of the coupling strength can also be performed using (massive) Nosé–Hoover chain thermostats[38] in the path integral context[39]. As we will demonstrate further down, the two approaches perform similarly at the level of accuracy whereas PILE is computationally less complex.

C. PIGLET Thermostatting

The choice of an efficient integrator and thermostat are important to thoroughly sample the path integral partition function. However, the convergence of the expectation values of estimators to their quantum mechanical limit cannot be accelerated by conventional thermostats that enforce a canonical distribution of positions and momenta.

In recent years, however, it has been shown that stochastic thermostats based on *colored* noise can be modified to induce a frequency dependent thermalization with the aim to mimic approximately quantum mechanical fluctuations – at the expense of breaking detailed balance[40,41]. This basic idea can be combined with PIMD, using the colored noise thermostat to accelerate the convergence with the number of beads[24]. The stationary distribution of a generalized Langevin equation (GLE) can be predicted analytically for a one dimensional harmonic oscillator of frequency ω and mass m. Subsequently, the parameters of the noise can be tuned to obtain an expectation value for $\langle q^2 \rangle = k_B \tilde{T}(\omega)/m\omega^2$, where $\tilde{T}(\omega)$ can be interpreted as corresponding to a frequency dependent temperature.

Then, when one applies such GLE to a path integral simulation of the same oscillator, the fluctuations will be modified from the case of canonical sampling, and yield

$$\langle q^2 \rangle_P = \frac{1}{P} \sum_{k=1}^{P} \left\langle \tilde{q}^{(k)2} \right\rangle = \sum_{k=1}^{P} \frac{k_B \tilde{T}^{(k)}(\sqrt{\omega^2 + \omega^{(k)2}})}{m(\omega^2 + \omega^{(k)2})}, \quad (8)$$

where we considered the most general case in which a different GLE has been applied to each of the normal modes, so the frequency dependence of the temperature $\tilde{T}^{(k)}(\omega)$ is mode specific. Note that a crucial aspect here is that the normal mode transformation is an orthogonal one, and that introducing a harmonic potential with frequency ω on top of the path integral molecular dynamics Hamiltonian of a free particle (5) amounts to a shift of the ring polymer frequencies *but leaves the eigenvectors unchanged*. This property means that it is possible to apply the GLE in the free ring polymer normal modes basis, without the need of knowing the vibrational frequencies and displacement patterns associated with the physical potential.

By adjusting the GLE parameters associated with different normal modes one can enforce the quantum me-

chanical expectation value for $\langle q^2 \rangle$ even for a small number of beads, as well as the bead–bead correlations that might be needed to accelerate convergence of the estimators of other properties. In the case of the PIGLET method[23], the simplifying assumption is made that $\tilde{T}^{(1)}(\omega) = T$ — that is, the centroid thermostat fulfills the classical fluctuation–dissipation theorem — and that $\tilde{T}^{(k)}(\omega) = \tilde{T}(\omega)$ — that is, all other ring polymer modes are subject to the same GLE which is tuned to enforce $\langle q^2 \rangle = \hbar/2m\omega \coth \hbar\omega\beta/2$. This is sufficient to guarantee that any structural (purely q–dependent) observable will be exact in the harmonic limit, and so will be the important centroid virial estimator of the particle kinetic energy — that also enters estimators for the pressure. While fluctuations can be made exact for any P with a harmonic potential, convergence for an arbitrary interaction potential V is guaranteed within PIGLET by the fact that in the $P \to \infty$ limit the GLE thermostat reduces to canonical sampling. Asymptotically, convergence will be achieved in a similar way as for conventional PIMD based on the primitive approximant. As we will see, other quantities that depend upon different bead–bead correlations converge much more slowly upon application of a PIGLET thermostat, just because no provision is made for them to be exact in the harmonic limit.

The procedure for obtaining the GLE parameters that give the desired $T(\omega)$ dependence, as well as efficient coupling to all vibrations, is described in detail in Ref. 42,43. The parameters that we used in this work at higher temperatures were available from the online repository Ref. 44, whereas we optimized new parameter sets for the lowest temperatures using the open–source tools available at Ref. 45.

While all of the techniques we discussed so far are available in the i–PI software[46], given the peculiar nature of the simulations we performed here — with a relatively inexpensive force field to be computed for tens of thousands of path integral replicas — we found that the communication overhead associated with the i–PI client/server model was too large to make this approach viable. We have therefore freshly implemented both PILE and PIGLET thermostats directly into the CP2k program package[47–49], which will prove useful for any future studies using CP2k.

III. COMPUTATIONAL DETAILS

In order to study systematically the behavior of molecular systems at low temperatures we decided to investigate different cryogenic regimes from moderate (100 K) to intermediate (20 K), and finally down to ultra low temperature (1.67 K) as being relevant to species solvated in superfluid helium. The temperatures were controlled using the PILE and PIGLET thermostats. For one case, massive Nosé–Hoover chain thermostatting[38] was used to benchmark the PILE thermostat (Figure 1

and 2 panel F). In the present case, the PILE thermostat is demonstrated to give the same average values as well as similar error bars for all tested properties compared to the Nosé–Hoover thermostat, but is computationally less demanding.

We selected two extreme cases as our tests so as to systematically scrutinize the behavior of PIGLET in different scenarios. Firstly, we considered protonated methane (CH_5^+) which is known to have a very flat and anharmonic hypersurface with respect to the angular coordinates[50], that results into utmost fluxional behavior even at very low temperatures[30]. Secondly, we simulated the very stiff, almost perfectly harmonic methane molecule (CH_4). For the description of both molecules we used the recently introduced POSflex model as previously parameterized for protonated methane[33]. For the quasi rigid methane molecule the angular part is defined according to Eq. (7) and Table I of Ref. 51, whereas the radial term is parameterized herein following the general procedure described in Ref. 33; the resulting force field parameters are compiled in Table I.

Table I. Parameterization of the anharmonic radial C–H potential of the POSflex model for CH_4 as defined in Ref. 33.

Parameter	Value
r_0	1.094 Å
k_2	1.283 E_h Å^{-2}
k_3	−6.627 E_h Å^{-3}
k_4	16.949 E_h Å^{-4}

For every selected combination of temperature, Trotter bead number, and thermostat we performed production runs of 900 ps with a time step of 0.25 fs, but we selected for further analysis one every 10th path configuration. For simulations with more than 64 beads we only stored 64 equally spaced replicas, discarding the others. The average Takahashi–Imada estimates[14] for the quantum kinetic energy were obtained from the PILE simulations in a post processing scheme[52–55] (using an approximate virial estimator[55] to compare with the virial estimates of the kinetic energy as obtained from the other methods).

IV. RESULTS AND DISCUSSION

A. Kinetic and Potential Energies

This section is structured in the following way. We first present our convergence study of quantum kinetic and potential energies as a function of the bead number. Next we discuss structural properties where we only present selected cases for the sake of clarity. All following figures are organized such that the left side (panel A,C,E) contains the results for methane whereas the right side (panel B,D,F) contains the results obtained for protonated methane. Panel A and B always contain results

from simulations at $100\,\mathrm{K}$, panel C and D from simulations at $20\,\mathrm{K}$, and panel E and F from simulations at $1.67\,\mathrm{K}$.
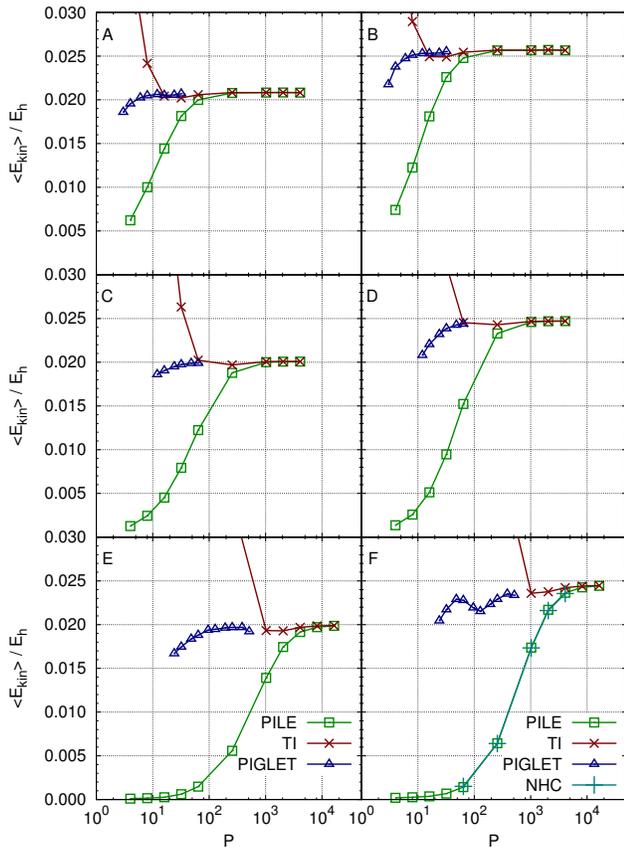

Figure 1. Convergence behavior of the average kinetic energy of CH_4 (A,C,E) and CH_5^+ (B,D,F) with the bead number P at $100\,\mathrm{K}$ (A,B), $20\,\mathrm{K}$ (C,D), and $1.67\,\mathrm{K}$ (E,F). The thermostats used in each simulation are indicated as follows: green squares (PILE), blue triangles (PIGLET), and turquoise pluses (massive Nosé–Hoover Chains) only in panel F (note that the NHC and PILE data are identical on the shown scale). The Takahashi–Imada estimates (red crosses) were obtained from the PILE simulations by *a posteriori* statistical reweighing (see text).

The convergence behavior of the average virial kinetic energy as generated by the distinct path integral simulation techniques can be analyzed based on Fig. 1. As expected, to converge the kinetic energy with a given level of accuracy, independent of the choice of thermostat, higher bead numbers are necessary for lower temperatures.

When considering the PILE simulations (green squares) at $100\,\mathrm{K}$ at least 256 beads are needed for convergence at the meV level. At a temperature of $20\,\mathrm{K}$ at least 1024 beads are needed, and for $1.67\,\mathrm{K}$ at least 8192 beads are required to reach such level of convergence. Importantly, the convergence of the kinetic energy is strictly monotone from the low bead limit all the way up to the

limiting behavior in case of PILE thermostatting (akin to the well-established massive Nosé–Hoover chain thermostats[38]). In the asymptotic regime, the convergence can be improved by using a higher order path integral scheme. Here, we used in particular a Takahashi–Imada estimator for the kinetic energy (specifically a hybrid thermodynamic/virial form[55] that does not suffer from increased variance with large P). All fourth-order Hamiltonians contain a term proportional to the square modulus of the physical force, which is problematic when one wants to use them in a PIMD context, since the molecular dynamics integrator then contains hard-to-compute second derivatives of the potential (which is particularly hampering in the realm of *ab initio* path integral[26–28] simulations). In order to circumvent this problem, we used the well known strategy of statistically reweighing second order trajectories[52,53,56]. It is noted that outside of the asymptotic regime the reweighed fourth order estimators dramatically overestimate the energy and, moreover, that convergence appears to be non monotone. This is, in fact, a manifestation of the pathological statistical behavior of the reweighing procedure when the difference between the sampled (Trotter) and the target (Takahashi–Imada) ensembles is too large. For low numbers of beads, the discrepancy is so large that the reweighed estimators are affected by a systematic error as well as by a huge statistical uncertainty. This problem is not specific to the TI method, but affects any fourth-order scheme when sampled with a reweighing approach. For the closely related Suzuki–Chin path integrals[16,17] it has been shown that the crucial parameter that determines the sampling efficiency (which should be on the order of unity or smaller to make reweighing viable) scales as $N_\mathrm{f}\beta^6/P^4$, where N_f is the number of degrees of freedom of the system[56]. The discrepancy between the sampled and target ensemble is slightly smaller for TI, which is the reason we used it here, but it shows a similar asymptotic behavior with system size and beads number.

Small systems such as those considered here are the ideal scenario to employ these higher order techniques, but the statistical problems are expected to worsen with lower T. Indeed, at our highest temperature of $100\,\mathrm{K}$ the systematic errors due to reweighing level off at about $P = 16$, and TI gives reasonable results already at $P = 64$ for both molecules. Lowering the temperature to $20\,\mathrm{K}$, the TI results reach a sweet spot for error cancellation at 64 beads before the agreement relative to the converged result slightly deteriorates at $P = 256$ and finally converges at 1024 beads. For $1.67\,\mathrm{K}$, at ultra low temperatures, the same non monotone convergence is exacerbated with a clear cross over from a regime dominated by the reweighing instabilities and thus dramatic overestimation of the kinetic energy up to a Trotter number as large as $P = 1024$, to the asymptotic convergence from there onwards.

Having in mind the aim to get even faster convergence with moderate bead numbers, without the issue

of statistical reweighing, we explored the performance of PIGLET thermostatting down to ultra low temperatures. Here the energy can be pushed close to convergence for relatively low bead numbers at the price, however, of non monotone convergence in the limit of striving for ultimate accuracy in the sense of fully converging the path integral discretization, $P \to \infty$. An extreme case of such non monotone convergence can be found for CH_5^+ at 1.67 K (panel F in Figure 1), where the PIGLET kinetic energy (blue triangles) is close to the converged value already at 48 beads, but declines for $P = 128$ where it yields a worse estimate than for 32 Trotter beads before it eventually reaches its converged value. By decomposing the kinetic energy into its atomic contributions, we observed that this non monotone behavior can be ascribed to the protons, whereas the kinetic energy of the carbon atom converges regularly. While we could not pin–point more precisely the origin of the non monotone behavior, we suspect that the problem could be related to zero point energy leakage[57], where different degrees of freedom that are thermostatted at different effective temperatures exchange energy through anharmonic couplings. Increasing the number of replicas reduces the gradient in $T(\omega)$ — as a larger fraction of quantum fluctuations is dealt with by the path integral — but at the same time increases the number of vibrational modes through which zero point energy can leak. A milder manifestation of this effect is probably involved in the small drop of the energy that is observed for the PIGLET simulations of methane at 1.67 K (panel E) and $P = 512$.

This non monotone behavior of the PIGLET thermostat with the bead number does not only arise for the kinetic, but also for the potential energy as evidenced by the data depicted in Fig. 2. It can be seen in all studied cases that the PIGLET thermostat underestimates the potential energy limit for low values of P. For larger bead number, however, the potential energy is slightly overestimated but approaches the correct limit for the largest investigated bead numbers. This region of overestimation of the potential energy seems to be system independent for our investigated cases and seems to be dictated only by the temperature, peaking at roughly 8, 48, and 384 beads for the temperatures 100 K, 20 K, and 1.67 K respectively.

While this non monotone convergence could be problematic when aiming at ultimate accuracy, it is clear that for *all* the cases we considered the PIGLET thermostatted simulations yield a *dramatic* improvement of *both* kinetic and potential energy for *any* number of beads compared to the standard PILE or NHC approaches. If one is satisfied with a residual error on the order of a few percent, PIGLET enables one to reduce the number of replicas needed for convergence by a factor between 10 and 100. For instance, in the case of CH_5^+ at 100 K the bead number can be lowered from $P = 256$ to 16 to achieve the same level of convergence (with 16–fold speedup) and for CH_4 at 1.67 K the bead number can be lowered by a factor of 42 (from $P = 4096$ to 96) with no
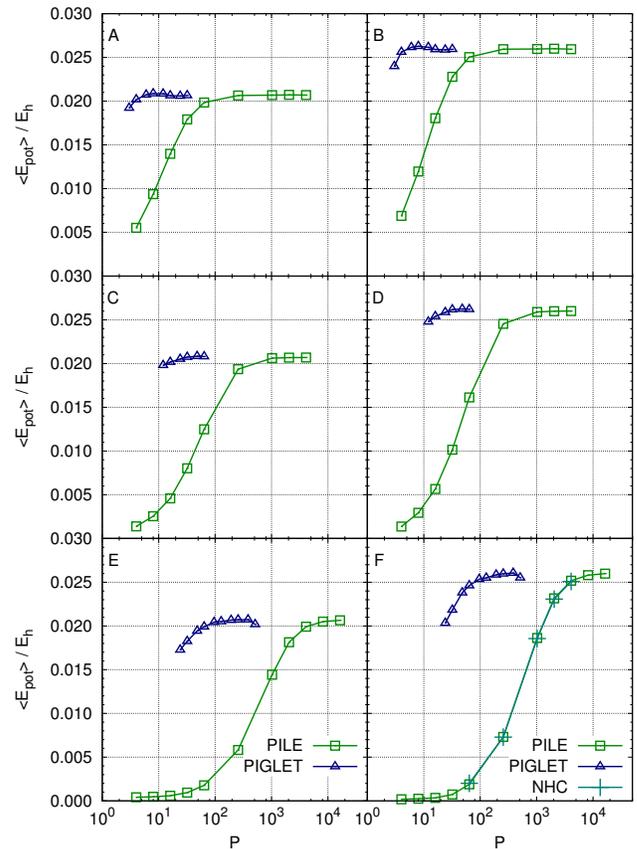

Figure 2. Convergence behavior of the average potential energy of CH_4 (A,C,E) and CH_5^+ (B,D,F) with the bead number P at 100 K (A,B), 20 K (C,D), and 1.67 K (E,F). The thermostats used in each simulation are indicated as follows: green squares (PILE), blue triangles (PIGLET), and turquoise pluses (massive Nosé–Hoover Chains) only in panel F (note that the NHC and PILE data are identical on the shown scale).

loss of accuracy. Note that thermostatting via PIGLET only requires generating about 20 random number per particle per replica per step, an operation whose computational overhead is barely noticeable even in a case such as the present in which the interatomic forces are obtained from an inexpensive empirical force field. As a consequence, the decrease in the bead number enabled by using PIGLET directly translates into a speedup.

B. Nuclear Quantum Effects

While energies provide an easily accessible and well defined indicator of convergence in terms of simple numbers, structural observables in terms of distribution functions are of more immediate interest in view of their direct connection to both NQEs on structure and also to experimental findings. In Fig. 3 and Fig. 4 we compiled representative distribution functions for the CH and HH dis-
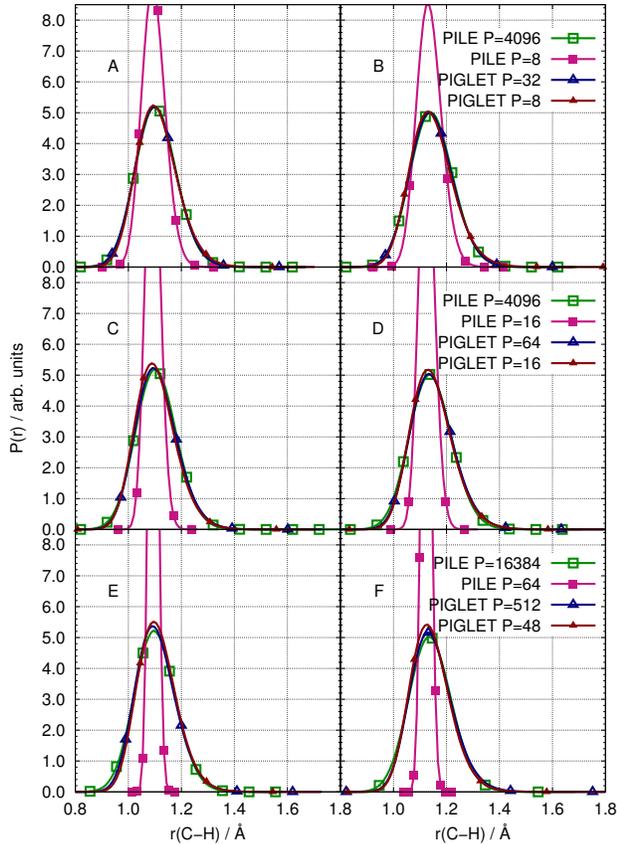
Figure 3. C–H distance distribution functions of CH_4 (A,C,E) and CH_5^+ (B,D,F). The panels correspond to $100\,K$ (A,B), $20\,K$ (C,D), and $1.67\,K$ (E,F) where representative combinations of thermostat type and bead number are selected as indicated. For each system and temperature we report the PILE and PIGLET data obtained with the largest bead numbers (the former serving as our benchmarks) and compare them with results obtained with significantly lower Trotter numbers as indicated.

tance, respectively. For the sake of clarity, we neither report the results for all possible internal coordinates nor for all simulation parameters such as bead number and thermostatting scheme. Instead, we chose to only present some representative distributions at a few specific bead numbers. For PILE as well as for PIGLET we report as a reference the highest number of beads we considered. In order to highlight the convergence properties of PIGLET compared to PILE, we also show the results for the lowest bead numbers that gives near-converged results with PIGLET. The key observation of highest practical relevance that can be extracted from Figs. 3 and 4 is that the PIGLET data obtained with the *lowest* discretization number P almost perfectly fit the *converged* PILE curves and features deviations which are hardly visible on the intrinsic variations of these distribution functions. Note that at $1.67\,K$ we found well converged PIGLET results with $P = 48$ already, but we did not run the correspond-

ing bead number for PILE so that we compared to the closest match for PILE being $P = 64$.

The results are very consistent across different temperatures and between the two molecules we considered. The highest number of beads we considered for PIGLET (which is about 10 to 50 times smaller than P used for the converged PILE benchmarks) agrees with the reference simulations within the accuracy that is provided by the graphs as depicted. Moreover, when using PIGLET it is possible to achieve nearly quantitative accuracy while reducing even further the number of replicas. For instance, at $100\,K$, already 8 replicas yield an almost perfect match with the fully converged distributions and even $P = 16$ gives only minor deviations at $20\,K$.

The most striking case is that of the ultra low temperature, $1.67\,K$, where 48 replicas suffice to capture the essential features of NQEs on the molecular structure of both quasi-rigid methane and fluxional protonated methane as described by the probability distribution functions in Figs. 3 and 4. In all these cases, conventional PILE thermostatting yields much too narrow distributions which is even associated with a shift of the maximum and thus of the most probable quantum structure in the most extreme case of CH_5^+ at cryogenic temperature.

As already alluded to, the HH distance distribution functions in Fig. 4 show similar trends. At all temperatures PIGLET can reproduce almost perfectly the fully converged PILE results with a considerable reduction in the number of beads and thus compute time. Moreover, they even semi quantitatively capture the shape of these probability distributions also at much smaller Trotter numbers – conditions for which PILE fails to reproduce even qualitatively the correct quantum mechanical behavior due to NQEs. A case that is worth discussing in more detail is that of CH_5^+ at $1.67\,K$. At this ultra low temperature, the positive impact of the colored noise in the low bead number regime is particularly striking. While PIGLET captures well the fluxional nature of the H atoms already with 48 beads, the distribution obtained with PILE at $P = 64$ breaks down in a multitude of sharper peaks, corresponding to proton–proton distances according to the equilibrium structure. In particular, distances around $1.2\,Å$ that occur during the scrambling motion via the C_{2v} transition state, and therefore are the signature of quantum fluxionality, are significantly underpopulated. We conclude that PILE simulations using such a bead number are unable to even qualitatively describe the fluxionality of CH_5^+ in the limit of ultra low temperatures such as those provided by superfluid helium.

There is, however, a price to pay for the much faster convergence of structural observables and energies when using PIGLET acceleration as we will clearly disclose in the following by analyzing the quantum delocalization properties of the protons upon approaching ultra low temperatures. Colored noise techniques like PIGLET introduce artificial, frequency dependent fluctuations that
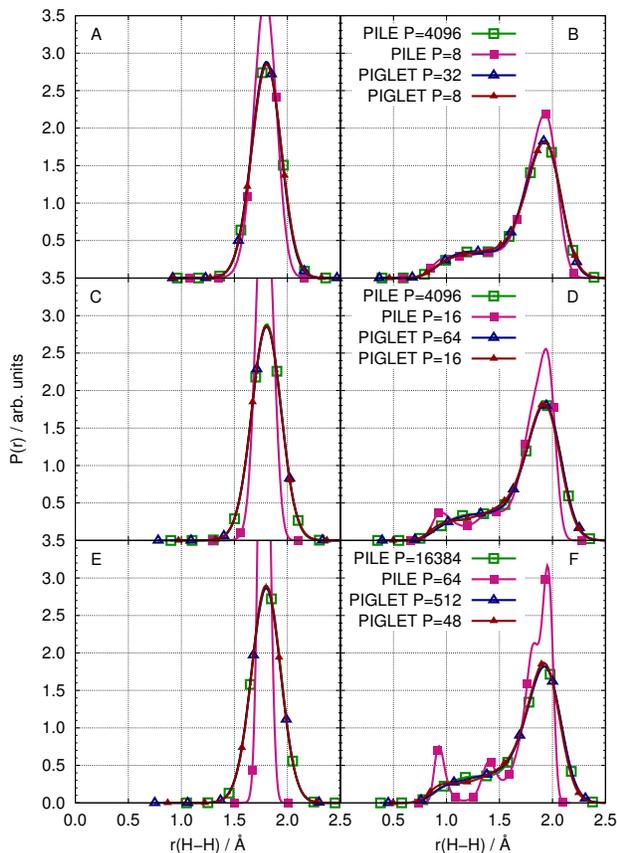
Figure 4. H–H distance distribution functions of CH_4 (A,C,E) and CH_5^+ (B,D,F). The panels correspond to $100\,K$ (A,B), $20\,K$ (C,D), and $1.67\,K$ (E,F) where representative combinations of thermostat type and bead number are selected as indicated. For each system and temperature we report the PILE and PIGLET data obtained with the largest bead numbers (the former serving as our benchmarks) and compare them with results obtained with significantly lower Trotter numbers as indicated.

are carefully designed to reproduce exactly those for a fully converged calculation with a purely harmonic potential. In the most basic setup[24], only the marginal distribution of individual beads is tuned to reproduce convergent results — which is already enough to accelerate convergence of any observable that can be written as a position dependent operator. Other quantum mechanical observables (such as those that are momentum dependent) also depend on cross correlations between different replicas. Each of such correlation has to be enforced separately, as it is done for instance in the PIGLET scheme to guarantee efficient estimation of the centroid virial kinetic energy estimator[23].

A quantity that is often used to assess the quantum mechanical spread of the nuclei is the the average radius of gyration of the ring polymer. For non spherically symmetric potentials, as in the present case, this idea can be generalized as worked out in Ref. 58, to obtain an

'effective radius' that quantifies the spatial spread of an arbitrary distribution. In Fig. 5 we analyze, as an example, the average effective radius for the protons in CH_4 for our three different temperatures as function of the beads number. As qualitatively expected, the effective
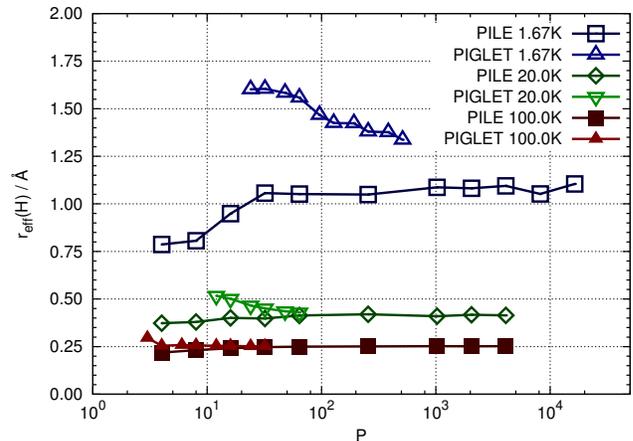

Figure 5. Convergence behavior of the effective radius, averaged over the four protons in CH_4, as a function of the bead number P at 100, 20, and $1.67\,K$ using PILE and PIGLET thermostats as indicated.

protonic radius increases when the temperature is lowered. At $100\,K$ the PILE effective radius converges to approximately $0.25\,\text{Å}$, whereas at $20\,K$ the converged number increases to about $0.4\,\text{Å}$. By lowering the temperature even further to $1.67\,K$ the converged spread seems to be in between roughly 1.0 and $1.1\,\text{Å}$. As it can be seen from the figure, the effective radius estimates from the PIGLET runs *systematically overestimate* the benchmark data even when using the *best discretization* of the path integral, and moreover exhibit slower convergence than what we saw for all other properties discussed so far. This is to be expected, as the effective radius estimator depends on beads cross correlations that have not been included in the target distribution for designing the colored noise parameters. The relative weight of different kinds of correlations depends on the temperature, which is apparent in the change in the speed of convergence at different temperatures. At $100\,K$, PIGLET seems to have an edge on PILE, but already at $20\,K$ the opposite is true. At the lowest temperature we considered, the convergence of PIGLET has become dramatically slow and deviates by more than 20% from the converged PILE values even with $P = 512$ beads. Despite this fundamental caveat, PIGLET nevertheless is able to describe the qualitative behavior of the quantum spread as a function of temperature down to superfluid helium temperatures. We close this critical discussion by mentioning that it would be possible — in principle — to extend the PIGLET framework to include also the correlations that contribute to the estimator for the effective radius, a line of research that we will explore in future work.

V. CONCLUSIONS AND OUTLOOK

In the present study we have investigated the applicability of colored noise acceleration of path integral molecular dynamics simulations to the study of ultra cold systems, such as for instance molecules embedded in superfluid helium droplets, pushing the method to its limits by considering the overly anharmonic, fluxional molecule CH_5^+ at $T = 1.67K$. The results are extremely encouraging since the number of replicas that is required to converge the kinetic or potential energy energy within a few percent relative to the $P \to \infty$ limit is significantly reduced, by a factor between 10 and 100. Structural distribution functions show similarly improved convergence: even at the lowest temperature and for the utmost anharmonic protonated methane case, PIGLET simulations with fewer than 100 beads give quantum probability distributions that are almost indistinguishable from the fully converged results which require more than 10'000 replicas. Contrary to all high-order factorization schemes of the propagator, the present approach can be readily applied to rather complex interaction potentials of relevance to molecule–based systems, including pointcharge models and polarizable force fields, without having to worry about the need of computing second- or even higher-order derivatives of the potential, or having to deal with statistically inefficient reweighting schemes. These virtues are particularly rewarding in the realm of *ab initio* path integral[26–28] simulations where the interactions are computed on-the-fly[28] from the fully self-consistent electronic structure as the path integral is sampled.

When pushing an approximate method to such extreme conditions, however, it is most interesting to investigate when and why it fails, and how it could be improved to overcome its shortcomings. Here we find, for instance, that the convergence of PIGLET is not monotone as a function of the discretization number, particularly for the kinetic energy of CH_5^+ at 1.67 K. The colored noise thermostat has to enforce a wildly non uniform effective temperature on physical vibrations of different frequency, a task that is made even harder, at low temperatures, by the presence of a quasi continuum set of ring polymer vibrations. A better understanding of the behavior of GLE thermostats in quasi harmonic systems, and of how the thermostat coupling strength should be tuned to contrast zero point energy leakage between different modes might help designing GLEs which can cope more stably with unusually anharmonic problems. Another aspect in which GLE frameworks need to be improved is the evaluation of estimators that depend on off diagonal bead correlations. As we reveal in detail, the effective radius of the ring polymer gets significantly and systematically overpredicted by PIGLET simulations and, in addition, convergence at the lowest temperatures is slower than by conventional path integral molecular dynamics.

Even with these caveats, our results underscore the great potential of the combination of non equilibrium colored noise thermostats and path integral molecular dynamics for the study of matter at extreme levels of quantum mechanical behavior. Thanks also to the native implementation of these techniques in the CP2k simulation package, this work will pave the way to the study of low temperature physics and cryochemistry, such as helium droplet spectroscopy of molecules, together with an *ab initio* treatment of interatomic forces that is enabled by the recently developed hybrid *ab initio*/bosonic PIMD/PIMC simulation technique[59,60].

REFERENCES

[1] M. E. Tuckerman, *Statistical Mechanics: Theory and Molecular Simulation* (Oxford University Press, Oxford, 2010).

[2] H. Kleinert, *Path Integrals in Quantum Mechanics, Statistics, Polymer Physics, and Financial Markets* (World Scientific, Singapore, 2004).

[3] R. M. Fye, Phys. Rev. B **33**, 6271 (1986).

[4] R. M. Fye and R. T. Scalettar, Phys. Rev. B **36**, 3833 (1987).

[5] J. Grotendorst, D. Marx, and A. Muramatsu, eds., *Quantum Simulations of Complex Many–Body Systems: From Theory to Algorithms* (John von Neumann Institute for Computing, FZ Jülich, 2002).

[6] S. Goyal, D. L. Schutt, and G. Scoles, Phys. Rev. Lett. **69**, 933 (1992).

[7] J. P. Toennies and A. F. Vilesov, Angew. Chem. Int. Ed. Eng. **43**, 2622 (2004).

[8] K. Szalewicz, Int. Rev. Phys. Chem. **27**, 273 (2008).

[9] D. M. Ceperley and E. L. Pollock, Phys. Rev. Lett. **56**, 351 (1986).

[10] J. A. Barker, J. Chem. Phys. **70**, 2914 (1979).

[11] D. M. Ceperley, Rev. Mod. Phys. **67**, 279 (1995).

[12] M. Suzuki, Commun. math. Phys. **51**, 183 (1976).

[13] H. D. Raedt and B. D. Raedt, Phys. Rev. A **28**, 3575 (1983).

[14] M. Takahashi and M. Imada, J. Phys. Soc. Jpn. **53**, 3765 (1984).

[15] X.-P. Li and J. Q. Broughton, J. Chem. Phys. **86**, 5094 (1987).

[16] M. Suzuki, Phys. Lett. A **201**, 425 (1995).

[17] S. A. Chin, Phys. Lett. A **226**, 344 (1997).

[18] F. R. Krajewski and M. H. Müser, Phys. Rev. B **65**, 174304 (2002).

[19] R. E. Zillich, J. M. Mayrhofer, and S. A. Chin, J. Chem. Phys. **132**, 044103 (2010).

[20] A. Perez and M. E. Tuckerman, J. Chem. Phys. **135**, 064104 (2011).

[21] R. E. Zillich, J. Comput. Phys. **301**, 111 (2015).

[22] M. Ceriotti, M. Parrinello, T. E. Markland, and D. E. Manolopoulos, J. Chem. Phys. **133**, 124104 (2010).

[23] M. Ceriotti and D. E. Manolopoulos, Phys. Rev. Lett. **109**, 100604 (2012).

[24] M. Ceriotti, D. E. Manolopoulos, and M. Parrinello, J. Chem. Phys. **134**, 84104 (2011).

[25] M. Ceriotti, in *Computational Trends in Solvation and Transport in Liquids*, edited by G. Sutmann, J. Grotendorst, G. Gompper, and D. Marx (Institute for Advanced Simulation and Jülich Supercomputing Centre, Forschungszentrum Jülich, 2015) pp. 1–17.

[26] D. Marx and M. Parrinello, Z. Phys. B (Rapid Note) **95**, 143 (1994).

[27] D. Marx, Lect. Notes Phys. **704**, 507 (2006).

[28] D. Marx and J. Hutter, *Ab Initio Molecular Dynamics: Basic Theory and Advanced Methods* (Cambridge University Press, Cambridge, 2009).

[29] M. Ceriotti, J. Cuny, M. Parrinello, and D. E. Manolopoulos, Proc. Natl. Acad. Sci. USA **110**, 15591 (2013).

[30] D. Marx and M. Parrinello, Nature **375**, 216 (1995).

[31] O. Asvany, P. Kumar P, B. Redlich, I. Hegemann, S. Schlemmer, and D. Marx, Science **309**, 1219 (2005).

[32] S. D. Ivanov, O. Asvany, A. Witt, E. Hugo, G. Mathias, B. Redlich, D. Marx, and S. Schlemmer, Nat. Chem. **2**, 298 (2010).

[33] F. Uhl, L. Walewski, H. Forbert, and D. Marx, J. Chem. Phys. **141**, 104110 (2014).

[34] M. Sprik, M. L. Klein, and D. Chandler, Phys. Rev. B **31**, 4234 (1985).

[35] E. L. Pollock and D. M. Ceperley, Phys. Rev. B **30**, 2555 (1984).

[36] M. E. Tuckerman, B. J. Berne, G. J. Martyna, and M. L. Klein, J. Chem. Phys. **99**, 2796 (1993).

[37] M. Tuckerman, B. J. Berne, and G. J. Martyna, J. Chem. Phys. **97**, 1990 (1992).

[38] G. J. Martyna, M. L. Klein, and M. Tuckerman, J. Chem. Phys. **97**, 2635 (1992).

[39] M. E. Tuckerman, D. Marx, M. L. Klein, and M. Parrinello, J. Chem. Phys. **104**, 5579 (1996).

[40] M. Ceriotti, G. Bussi, and M. Parrinello, Phys. Rev. Lett. **103**, 30603 (2009).

[41] H. Dammak, Y. Chalopin, M. Laroche, M. Hayoun, and J.-J. Greffet, Phys. Rev. Lett. **103**, 190601 (2009).

[42] M. Ceriotti, G. Bussi, and M. Parrinello, J. Chem. Theory Comput. **6**, 1170 (2010).

[43] M. Ceriotti, *A novel framework for enhanced molecular dynamics based on the generalized Langevin equation*, Ph.D. thesis, ETH Zürich (2010).

[44] M. Ceriotti, https://epfl-cosmo.github.io/gle4md/ (2016).

[45] M. Ceriotti, https://github.com/epfl-cosmo/gle4md (2016).

[46] M. Ceriotti, J. More, and D. E. Manolopoulos, Comp. Phys. Comm. **185**, 1019 (2014).

[47] CP2k Developers Team, www.cp2k.org (2016).

[48] J. VandeVondele, M. Krack, F. Mohamed, M. Parrinello, T. Chassaing, and J. Hutter, Comput. Phys. Commun. **167**, 103 (2005).

[49] J. Hutter, M. Iannuzzi, F. Schiffmann, and J. VandeVondele, WIREs Comput. Mol. Sci. **4**, 15 (2014).

[50] P. R. Schreiner, S. Kim, H. F. Schaefer III, and P. von Ragué Schleyer, J. Chem. Phys. **99**, 3716 (1993).

[51] M. P. Deskevich, A. B. McCoy, J. M. Hutson, and D. J. Nesbitt, J. Chem. Phys. **128**, 094306 (2008).

[52] S. Jang, S. Jang, and G. A. Voth, J. Chem. Phys. **115**, 7832 (2001).

[53] T. M. Yamamoto, J. Chem. Phys. **123**, 104101 (2005).

[54] M. Buchowiecki and J. Vaníček, Chem. Phys. Lett. **588**, 11 (2013).

[55] O. Marsalek, P.-Y. Chen, R. Dupuis, M. Benoit, M. Méheut, Z. Bačić, and M. E. Tuckerman, J. Chem. Theory Comput. **10**, 1440 (2014).

[56] M. Ceriotti, G. a. R. Brain, O. Riordan, and D. E. Manolopoulos, Proc. Roy. Soc. A **468**, 2 (2011).

[57] S. Habershon and D. E. Manolopoulos, J. Chem. Phys. **131**, 244518 (2009).

[58] L. Walewski, H. Forbert, and D. Marx, Mol. Phys. **111**, 2555 (2013).

[59] L. Walewski, H. Forbert, and D. Marx, Comput. Phys. Commun. **185**, 884 (2014).

[60] L. Walewski, H. Forbert, and D. Marx, J. Chem. Phys. **140**, 144305 (2014).