Accelerated path-integral simulations using ring-polymer interpolation

Samuel J. Buxton[1] and Scott Habershon[1, a)]
Department of Chemistry and Centre for Scientific Computing, University of Warwick, Coventry, CV4 7AL, United Kingdom

Path-integral (PI) molecular simulations can be used to calculate exact quantum statistical mechanical properties for complex systems containing many interacting atoms and molecules. The limiting computational factor in a PI simulation is typically the evaluation of the potential energy surface (PES) and forces at each ring-polymer "bead"; for an n-bead ring-polymer, a PI simulation is typically n times greater than the corresponding classical simulation. To address the increased computational effort of PI simulations, several approaches have been developed recently, most notably based on the idea of ring-polymer contraction (RPC) which exploits either the separation of the PES into short-range and long-range contributions or the availability of a computationally-inexpensive PES which can be incorporated to effectively smooth the ring-polymer PES; neither approach is satisfactory in applications to systems described by *ab initio* PESs. In this Article, we describe a new method, ring-polymer interpolation (RPI), which can be used to accelerate PI simulations without any prior assumptions about the PES. In simulations of liquid water under ambient conditions, where quantum effects are known to play a subtel role in influencing experimental observables such as diffusion coefficients and radial distribution functions, we find that RPI can accurately reproduce the results of fully-converged PI simulations, albeit with far fewer PES evaluations; this approach therefore opens up the possibility of large-scale PI simulations on *ab initio* PESs at lower computational effort than current methods.

I. INTRODUCTION

Path-integral (PI) simulations enable the exact calculation of time-independent quantum properties in general molecular systems.[1–17] In the path integral formulation of quantum statistical mechanics,[1] each quantum particle in the system is mapped onto a classical n-bead harmonic ring-polymer; exploiting this isomorphism, sampling the classical configurational space of the ring-polymer by either Monte Carlo (PIMC) or molecular dynamics (PIMD) enables determination of static properties such that nuclear quantum effects such as zero-point energy (ZPE) conservation and tunnelling are exactly accounted for (at least in the $n \to \infty$ limit). PIMD simulations are particularly appealing due to the fact that many of the strategies developed to enable efficient sampling in classical MD simulations, including improved thermostats[18] and multiple time-step methods,[19,20] can be straightforwardly implemented within the PI framework. As a result, PIMD simulations have been employed to investigate systems ranging from structure in liquid and solid water phases[13,16,17,21–27] to free energies in enzyme-catalyzed proton transfer.[28–31] More recently, PIMD-based strategies have been proposed which enable calculation of approximate dynamic (time-dependent) properties; these approaches, including ring-polymer molecular dynamics (RPMD[16,25,26,29,32–41]), centroid molecular dynamics (CMD[22,42–47]) and semiclassical instanton theory[48,49] now provide a useful toolbox for interrogating the influence of quantum effects in complex condensed-phase dynamics.

Computationally, the most demanding aspect of PI simulations is the evaluation of the potential energy surface (PES) and resultant forces at each configuration around the n-bead ring-polymer; this usually results in a computational expense which is n times greater than the corresponding classical simulation. Although path integral simulations are formally exact in the $n \to \infty$ limit, in practice convergence of calculated properties is typically achieved by choosing the number of ring-polymer "beads" such that $\beta\hbar\omega_{\max}/n \ll 1$, where $\beta = 1/(k_{\mathrm{B}}T)$ and $\omega_{\max}$ is the characteristic highest physical frequency in the system. In liquid water at ambient temperature it is common to select $n = 32$, reflecting the fact that the high-frequency intramolecular O-H vibrations ($\omega_{\max} \simeq 3600\,\mathrm{cm}^{-1}$) have large ZPE relative to the available thermal energy ($k_{\mathrm{B}}T \simeq 200\ \mathrm{cm}^{-1}$ at $T = 298$ K).[16,41] If the PES describing the system is computationally-inexpensive (*e.g.* a simple empirical force-field and/or a small system size), then the additional cost associated with path integral simulations is of little consequence, particularly if one can exploit the implicit parallelism of path integral methods. However, when PES evaluation is computationally-expensive (*e.g.* using *ab initio* methods such as density functional theory (DFT)), then the associated cost of path integral simulations relative to classical simulations can be prohibitive.

Several techniques have been developed to address the challenge associated with the computational expense of PI simulations. The ring-polymer contraction (RPC[5,50]) scheme, outlined in more detail below, exploits the separation of the PES into components which have "high" and "low" associated frequencies; given that the number of ring-polymer beads required to converge quantum observables in path integral simulations depends strongly on the frequencies of the associated mo-

a)Electronic mail: S.Habershon@warwick.ac.uk

tions, as noted above, this separation allows one to employ different numbers of ring-polymer beads to evaluate the different contributions to the PES. When applied to simple empirical force-fields, such the q-TIP4P/F model[16] employed below, this RPC scheme enables one to use a "contracted" ring-polymer comprising a few (∼6) ring-polymer beads to evaluate the "slow" long-range PES components arising from Coulomb and Lennard-Jones dispersion forces, whereas the full complement of ring-polymer beads (∼32) must be used to assess the PES component corresponding to the high-frequency intramolecular motions. Importantly however, because the evaluation of the long-range components of typical empirical force-fields is the most time-consuming part of a molecular simulation, this RPC scheme allows an overall reduction in time required to evaluate forces in PIMD. The original RPC scheme was further refined with the introduction of electrostatic RPC,[50] whereby the evaluation of the Coulombic contribution to the empirical PES was further accelerated by exploiting a range-separation of the ring-polymer forces; this electrostatic RPC scheme ultimately enables PIMD simulations which are only a factor of around three slower than the corresponding classical MD simulation. A clear demonstration of the utility of this methodology was in the determination of the quantum-mechanical melting point of q-TIP4P/F water in direct coexistence PIMD simulations; here, systems comprising 696 molecules were simulated using PIMD for time-periods of up to 10 ns.[16]

Unfortunately, both the original RPC scheme and its electrostatic variant cannot be applied directly to more general PESs such as those arising in DFT or other *ab initio* calculations, the reason being that both schemes rely on the ability to decompose the PES into independent contributions which can be identified as either "high" or "low" frequency (or short-range and long-range); such a separation is not straightforward in *ab initio*-based PESs. As a result, Markland and Marsalek[20] and, independently, Kühne and coworkers,[51] have proposed a RPC-like scheme which relies on the availability of a "reference" PES which is broadly similar to the PES under direct investigation but much less computationally-expensive to evaluate. Here, the reference PES is evaluated on the entire n-bead ring-polymer while the "full" PES of interest is evaluated on a contracted ring-polymer with a smaller number of beads n_c; the underlying assumption at play is that the reference PES faithfully captures the rapidly-varying part of the full PES, such that the slowly-varying remainder between the reference and full PES can be evaluated using a reduced number of beads. In the work of Markland and Marsalek, a density functional tight binding (DFTB) model was employed as the reference PES while DFT was employed as the full PES to model liquid water; using $n = 32$ ring-polymer beads to evaluate the reference DFTB PES, it was found that correct reproduction of the expected DFT PIMD properties, including radial distribution function (RDF) and average proton kinetic energy, required a contracted ring-polymer comprising around $n_c = 6$ beads to evaluate the full DFT PES, thereby representing a computational saving of roughly a factor of five relative to a full $n = 32$ DFT PIMD simulation. However, in favourable cases, this work also found that using $n_c = 1$ contracted ring-polymer beads for evaluation of the full DFT PES also gave reasonably good reproduction of quantum properties, suggesting that this reference PES RPC scheme enables quantum simulations at near-classical cost. A similar conclusion was reached by Kühne and coworkers,[51] although in this case the reference PES was selected to be a simple fixed-charge empirical model, similar to q-TIP4P/F,[16] which was force-matched to a DFTB PES in a preliminary step; this approach enables the reference PES calculations to simultaneously exploit the original and electrostatic RPC schemes.

While the two reference PES RPC schemes highlighted above are undoubtedly successful in reducing the computational cost relative to full PIMD simulation on *ab initio* PESs, their reliance on the availability of an inexpensive yet reasonably accurate reference model suggests that there is room for further improvement. In this Article, we describe a new method which circumvents the necessity of a reference PES completely, enabling direct PIMD simulations on generic PESs at a fraction of the computational expense relative to the full PIMD simulation. Our scheme, referred to as ring-polymer interpolation (RPI), employs Gaussian process regression (GPR[52–59]) to evaluate the forces and potential energy on the n-bead ring-polymers using only a small number of direct PES evaluations at each time-step; in this paper, we show that our RPI scheme is trivial to implement and systematically converges towards the exact PIMD simulation results for test cases including liquid water under ambient conditions[16] and liquid *para*-hydrogen at low temperatures.[35] Overall, RPI provides a straightforward approach to performing accurate and efficient PIMD simulations on arbitrary PESs without a reference PES.

The remainder of this Article is organized as follows. First, we outline RPC schemes proposed to date and detail our new RPI approach. Then, we compare and contrast RPC methods and RPI in PIMD simulations of liquid *para*-hydrogen and liquid water. Finally, we conclude by highlighting several routes available to further improve our RPI scheme.

II. THEORY

In this Section, we briefly outline the PIMD simulation approach for calculating static (time-independent) quantum properties. To set the context for our new approach, we then outline the original RPC approach and the recent development of the reference PES RPC methodology. Finally, we describe our new RPI methodology for efficient PIMD simulations on arbitrary PESs without the need for a reference potential.

A. Path-integral molecular dynamics

In the PI approach to quantum statistical mechanics, each *quantum* particle is mapped onto a *classical* n-bead ring-polymer; the classical statistical mechanics of the ring-polymer corresponds exactly to the quantum statistical mechanics of the original system, enabling determination of time-dependent properties while exactly accounting for the role of quantum fluctuations.[1,60,61]

Here, we assume we have a system comprising N atoms described by a Hamiltonian which is a sum of kinetic and potential terms, $\hat{H} = \hat{T} + \hat{V}$; as is most common in PI simulations, exchange effects are neglected. The standard canonical PI approach begins with the quantum thermal partition function,

$$Z = \mathrm{Tr}\left[e^{-\beta\hat{H}}\right], \tag{1}$$

where $\beta = 1/(k_{\mathrm{B}}T)$, T is temperature, and $\hat{H}$ is the Hamiltonian operator for the system of interest. By evaluating the trace of Eq. 1 in a basis of position eigenstates and exploiting the well-known symmetric Trotter splitting,

$$e^{-\beta_n\hat{H}} = \lim_{n\to\infty} e^{-\beta_n\hat{V}/2}e^{-\beta_n\hat{T}}e^{-\beta_n\hat{V}/2} \tag{2}$$

where $\beta_n = \beta/n$, it is straightforward to show that the quantum partition function can be written as,

$$Z = \lim_{n\to\infty} \frac{1}{(2\pi\hbar)^f} \int d^f\mathbf{r}\, d^f\mathbf{p}\, e^{-\beta_n H_n(\mathbf{r},\mathbf{p})}. \tag{3}$$

Here, $\mathbf{r}$ and $\mathbf{p}$ are, respectively, the positions and momenta of a set of $N \times n$ particles, $f = 3Nn$, and the ring-polymer Hamiltonian $H_n(\mathbf{r},\mathbf{p})$ is given by

$$H_n(\mathbf{r},\mathbf{p}) = H_0(\mathbf{r},\mathbf{p}) + \sum_{i=1}^{n} V(\mathbf{r}^{(i)}), \tag{4}$$

and the free ring-polymer Hamiltonian is

$$H_0(\mathbf{r},\mathbf{p}) = \sum_{i=1}^{n}\sum_{j=1}^{N}\left[\frac{|\mathbf{p}_j^{(i)}|^2}{2m_j} + \frac{1}{2}m_j\omega_n^2(\mathbf{r}_j^{(i)} - \mathbf{r}_j^{(i-1)})^2\right]. \tag{5}$$

In Eqs. 4 and 5, m_j is the mass of particle j, $\omega_n = 1/(\beta_n\hbar)$ and $V(\mathbf{r}^{(i)})$ is the PES of the system evaluated on bead i. For clarity, note that $\mathbf{r}_j^{(i)}$ is the position of particle j in the ith ring-polymer bead.

The ring-polymer Hamiltonian of Eq. 4 defines a system in which each quantum particle has been replaced by a classical P-bead ring-polymer; by sampling the extended phase-space of the classical system, quantum statistical properties may be evaluated according to,

$$\begin{aligned}\langle A\rangle &= \mathrm{Tr}\left[e^{-\beta\hat{H}}\hat{A}\right] \\ &= \lim_{n\to\infty}\frac{1}{(2\pi\hbar)^f}\int d^f\mathbf{r}\, d^f\mathbf{p}\, e^{-\beta_n H_n(\mathbf{r},\mathbf{p})}A_n(\mathbf{r}),\end{aligned} \tag{6}$$

where the ring-polymer average of the operator $\hat{A}$ is

$$A_n(\mathbf{r}) = \frac{1}{n}\sum_{i=1}^{n} A(\mathbf{r}^{(i)}). \tag{7}$$

The Hamiltonian of Eq. 4 can be used to generate equations-of-motion for the ring-polymer positions and momenta, such that quantum thermal averages can be calculated using Eq. 6; this is the basis of the PIMD approach to calculating quantum properties in complex systems. In passing, we note that PIMD is only applicable in the calculation of static (time-independent) properties; however, the last two decades has witnessed the development of PI-based methods, including ring-polymer molecular dynamics (RPMD[24,26,32–37,39,41,42,48,62]), centroid molecular dynamics (CMD[43–47]) and, most recently, Matsubara dynamics[42,63] which can be used to approximate quantum-mechanical time-dependent properties such as time-correlation functions. Both the RPC methodologies and our new RPI approach are, in general, equally applicable to these dynamic simulation methods, although we focus here on PIMD simulations for clarity of presentation.

B. Ring-polymer contraction

At this point it is worth emphasising the additional computational cost of PIMD relative to standard classical MD simulations. As shown above, PIMD requires n evaluations of the PES at each time-step in order to determine the forces acting on each of the n ring-polymer beads representing the quantum particles; in the usual case when evaluation of the PES and forces is the most time-consuming part of the simulation, this suggests that PIMD simulations are around a factor of n times more computationally-expensive than classical MD. While parallel computing offers one route to minimizing the impact of this additional expense, an alternative is to seek new algorithms which exploit the underlying physical features of the problem to reduce the number of force evaluations at each time-step; RPC is one route to addressing this goal.

The underlying assumption of the RPC scheme is that the PES can be split into identifiable parts associated with "low" and "high" frequency motion:

$$V(\mathbf{r}) = V_{\mathrm{l}}(\mathbf{r}) + V_{\mathrm{h}}(\mathbf{r}). \tag{8}$$

As a result, the PES sampled by the ring-polymer in

PIMD can be written as

$$\sum_{i=1}^{n} V(\mathbf{r}^{(i)}) = \sum_{i=1}^{n} V_l(\mathbf{r}^{(i)}) + \sum_{i=1}^{n} V_h(\mathbf{r}^{(i)}). \tag{9}$$

As a concrete example, $V_l(\mathbf{r})$ might correspond to the intermolecular component of a typical empirical force-field, comprising Lennard-Jones and point-charge Coulomb interactions, whereas $V_h(\mathbf{r})$ might represent the intramolecular PES, perhaps comprising bond-stretching and bond-angle bending contributions; this decomposition has been exploited in simulations of the SPC/F and q-TIP4P/F water models.[5,16]

The characteristic vibrational frequency, ω, in a given system provides a rule-of-thumb in determining the required number of ring-polymer beads required to obtain converged quantum statistical-mechanical properties; typically, n is chosen such that $\beta\hbar\omega_{\max}/n \ll 1$. This immediately suggests that the low-frequency contribution to the PES ($V_{\rm l}(\mathbf{r})$) requires fewer ring-polymer beads for converged evaluation than the high-frequency contribution ($V_{\rm h}(\mathbf{r})$). RPC exploits this fact by evaluating $V_{\rm l}(\mathbf{r})$ on a "contracted" ring-polymer containing $n' < n$ beads, rather than the full n-bead ring-polymer. To achieve this the ring-polymer is first transformed into a representation comprising the normal modes of the free ring-polymer Hamiltonian of Eq. 5. Subsequently, the $n - n'$ highest-frequency normal modes are removed and the inverse Fourier transformation back to real-space is performed, resulting in a n'-bead ring-polymer; the net transformation is

$$\mathbf{r}_j^{(i')} = \sum_{i=1}^{n} T_{i'i}\mathbf{r}_j^{(i)}, \tag{10}$$

where $\mathbf{r}_j^{(i')}$ is the position of particle j in replica i' in the contracted ring-polymer, and the elements $T_{i'i}$ are known functions arising from the normal-mode transformation of the free ring-polymer Hamiltonian.[5] Once the contracted ring-polymer coordinates have been generated, the contribution of $V_l(\mathbf{r})$ to the *total* potential energy of the full n-bead ring-polymer system is approximated as

$$\sum_{i=1}^{n} V_l(\mathbf{r}) \simeq \frac{n}{n'} \sum_{i'=1}^{n'} V_l(\mathbf{r}^{(i')}), \tag{11}$$

while the forces on the full n-bead ring-polymer can be recovered from the contracted PES by application of the chain-rule. Importantly, the "low-frequency" contribution to the PES, most commonly identified as the long-range intermolecular interaction terms, is usually the most computationally-expensive to evaluate. As a result, evaluating $V_l(\mathbf{r})$ on a sub-set of the n ring-polymer beads offers a direct improvement in computationally efficiency; as one might expect, practical assessment of RPC demonstrates that the resulting simulations are around a factor of n/n' faster than the corresponding n-bead PIMD simulation. Furthermore, the convergence of RPC with respect to n' has also been clearly demonstrated, for example by analyzing quantum kinetic energies, potential energies and structural properties for liquid water.[5]

To conclude this section, we note that RPC can be further refined in systems containing point-charge or dipolar electrostatic interactions.[50,64] Here, the electrostatic contributions are themselves "range-separated", with the short-range contribution typically evaluated on a small number of ring-polymer beads while the long-range contribution is evaluated only once, at the centroid (centre-of-mass) of the ring-polymer. This general strategy has been demonstrated for both the SPC/F empirical force-field, containing point-charge Coulombic interactions,[50] as well as the TTM3-F model,[64], possessing Thole-type polarisability; overall, this electrostatic RPC approach enables PIMD simulations which are roughly a factor of ten times faster than the standard n-bead PIMD simulation.

C. Ring-polymer contraction with a reference potential

While RPC, as described above, is certainly successful in reducing the computational cost of PIMD simulations, it has one important disadvantage; RPC exploits the separation of the full PES $V(\mathbf{r})$ into contributions which can be identified as having slowly- and rapidly-varying components. In the case of empirical force-fields, such as SPC/F, q-TIP4P/F and TTM3-F water models, this decomposition is straightforward. However, in the case of PESs derived from *ab initio* simulations, such as density functional theory (DFT), a trivial PES decomposition is not immediately available.

To address this challenge, Markland and Marsalek[20] and, independently, Kühne and coworkers,[51] proposed a scheme based on using a reference PES, $V_{\rm ref}(\mathbf{r})$. In particular, one rewrites the potential energy contribution to the ring-polymer Hamiltonian as

$$\begin{aligned}\sum_{i=1}^{n} V(\mathbf{r}^{(i)}) &= \sum_{i=1}^{n} \left[V(\mathbf{r}^{(i)}) + V_{\rm ref}(\mathbf{r}^{(i)}) - V_{\rm ref}(\mathbf{r}^{(i)}) \right], \\ &= \sum_{i=1}^{n} V_{\rm ref}(\mathbf{r}^{(i)}) + \sum_{i=1}^{n} \left[V(\mathbf{r}^{(i)}) - V_{\rm ref}(\mathbf{r}^{(i)}) \right], \\ &\simeq \sum_{i=1}^{n} V_{\rm ref}(\mathbf{r}^{(i)}) + \frac{n}{n'} \sum_{i'=1}^{n'} \left[V(\mathbf{r}^{(i')}) - V_{\rm ref}(\mathbf{r}^{(i')}) \right],\end{aligned} \tag{12}$$

The third line of Eq. 12 defines the key approximation introduced in this reference-PES-based RPC scheme. Here, the reference PES $V_{\rm ref}(\mathbf{r})$ is evaluated on the full n-bead ring-polymer; the difference term in Eq. 12 is, however, evaluated on a n'-bead contracted ring-polymer, with the coordinates of the contracted replicas determined in the same manner as in the original RPC scheme.

The assumption which underlies this reference RPC

(rRPC) scheme is that the *difference* potential $V(\mathbf{r}) - V_{\mathrm{ref}}(\mathbf{r})$ is slowly-varying, in the same way that the intermolecular PES was assumed to be slowly-varying in the original RPC scheme. The upshot of this rRPC scheme is that the computationally-expensive full PES $V(\mathbf{r})$ is evaluated on just $n' < n$ ring-polymer beads while the less demanding reference PES $V_{\mathrm{ref}}(\mathbf{r})$ is evaluated on the full n-bead ring-polymer. To date, simulations of liquid water and the protonated water dimer, both described using DFT to calculate the full PES, have demonstrated the validity of this assumption, enabling PIMD simulations for a fraction of the cost of the standard n-bead approach.[20]

An important assumption of the rRPC methodology is that a reference PES is available for the system at hand which is simultaneously inexpensive to evaluate *and* provides a reasonable level of reproduction of the properties of the system under investigation; in previous applications, both DFTB and a force-matched empirical PES have been used as reference PESs with clear success.[20,51] However, being tied to the availability of a reference PES is clearly undesirable; for example, it adds another layer of complexity to code management, with not one but two PESs required for evaluation at different points, and there is no guarantee that the reference PES will be sufficiently accurate to model more complex chemical reactions. As a result, it is appealing to investigate alternative strategies which circumvent the necessity of a reference PES; the strategy we outline here aims to do just this.

D. A new approach: ring-polymer interpolation

We now present a new approach to accelerating PIMD simulations on general PESs; in particular, we do not assume anything about the form of the underlying PES, and neither do we rely on the availability of a reference PES.

The underlying idea behind our approach is illustrated in Fig. 1. Here, we show the (shifted) PES imaginary-time autocorrelation function[8,38,65] for a PIMD simulation ($n = 32$) of liquid water at 298 K, described by the q-TIP4P/F empirical force-field.[16] The shifted imaginary-time autocorrelation function can be defined as,

$$C_j^{\mathrm{im}} = \langle V_1 V_j \rangle - \langle V_1 V_{(1+\frac{n}{2})} \rangle, \tag{13}$$

where

$$\langle V_1 V_j \rangle = \lim_{n\to\infty} \frac{1}{(2\pi\hbar)^f} \int d^f\mathbf{r}\, d^f\mathbf{p}\, e^{-\beta_n H_n(\mathbf{r},\mathbf{p})} V(\mathbf{r}^{(1)}) V(\mathbf{r}^{(j)}). \tag{14}$$

Here $V(\mathbf{r}^{(k)})$ refers to the value of the PES evaluated at bead k; because this is a static property (independent of real time), such imaginary-time correlation functions can be calculated exactly in PIMD simulations. The function C_j^{im} expresses the average correlation in PES values as one steps around the n-bead ring-polymer in the PIMD simulation, and is defined such that it approaches zero as one steps towards the ring-polymer bead which sits diametrically opposite a chosen reference bead (in the case of an even number of ring-polymer beads, as considered here, the index of the antipode bead is $1 + \frac{n}{2}$). The key observation which is relevant to this work is that there is a significant degree of correlation in the PES values as one steps around the ring; for example, we find that the correlation function at $j = 5$ is still non-zero, indicating that the PES value of this ring-polymer bead is correlated (related to) the PES value at bead $j = 1$. This suggests that it may not be necessary to perform independent evaluations of the PES on all beads; instead, the extent of correlation in PES values can be exploited. This is the basis of our ring-polymer interpolation (RPI) method.

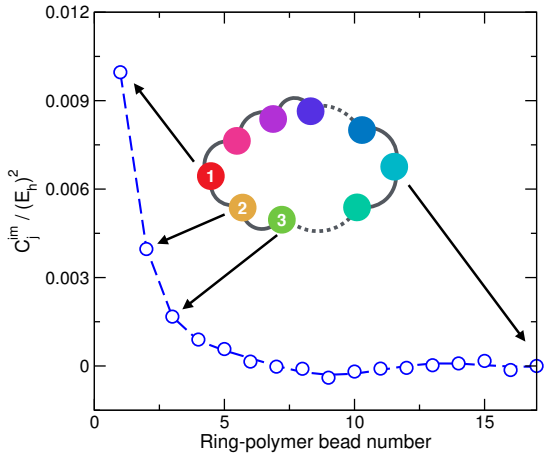

FIG. 1

Our approach employs the idea of interpolation to approximate the PES around the ring-polymers sampled during PIMD simulations. Consider a PIMD simulation sampling an n-bead ring-polymer in an $3N$-dimensional system; each of the n ring-polymer beads can be consecutively labelled by an integer $1 \leq j \leq n$ and, because of the cyclic nature of the ring-polymer, any bead may be selected as $i = 1$. From the n-bead ring-polymer beads, a smaller subset $n' < n$ of ring-polymer beads are selected and the full PES is evaluated on these n' beads; the PES and forces on the remaining $n - n'$ ring-polymer beads can then be recovered by direct interpolation using the known PES values at the n' beadslambda. Based on the results of Fig. 1, we note that this interpolation can be performed in the one-dimensional space defined simply by the bead indices.

In principle, any interpolation method can be used to recover the PES on the $n - n'$ ring-polymer beads; in this Article, we choose to use Gaussian Process regression (GPR[57,59,66–68]), primarily because of its simplicity and flexibility. GPR is well-known to the machine-learning community, and has found extensive use in computational chemistry; for example, recent work has high-

lighted development of accurate PESs by applying Gaussian Process to *ab initio* electronic structure calculations for systems such as bulk silicon[58] and liquid water.[57] In the context of this work, we assume that we have evaluated full PES values at n' selected beads; using these values, the PES at any bead k in the ring-polymer is then approximated in GPR as,

$$V(\mathbf{r}^{(k)}) \simeq \sum_{i=1}^{n'} w_i e^{-\alpha(k-\lambda_i)^2}, \tag{15}$$

where λ_i labels the ring-polymer indices of the selected subset of n' beads and α is a width parameter. The expansion weights, w_i, are determined by requiring that the PES values are correctly reproduced at the n' beads; it is straightforward to show that this requires solution of a $n' \times n'$ linear equation,

$$\mathbf{A}\mathbf{w} = \mathbf{g}, \tag{16}$$

where

$$A_{ij} = \delta_{ij}\sigma^2 + e^{-\alpha(\lambda_i-\lambda_j)^2}, \tag{17}$$

and

$$g_i = V(\mathbf{r}^{(i)}), \tag{18}$$

Here, σ^2 can be viewed as either a regularisation parameter or as representing the error in the PES evaluations at the n' beads, and the vector $\mathbf{g}$ contains the n' PES values. Given the PES values at the n' beads, solution of Eq. 16 is straightforward using standard linear algebra packages.[69] Alternatively, from explicit inversion of $\mathbf{A}$, the set of required weights is given by

$$w_i = \sum_{m=1}^{n'} A^{-1}_{im} V(\mathbf{r}^{\lambda_m}). \tag{19}$$

From Eq. 15, we then have

$$\begin{aligned} V(\mathbf{r}^{(k)}) &\simeq \sum_{i=1}^{n'}\sum_{m=1}^{n'} A^{-1}_{ik} V(\mathbf{r}^{\lambda_m}) e^{-\alpha(k-\lambda_i)^2} \\ &\simeq \sum_{m=1}^{n'} \bar{\omega}_{km} V(\mathbf{r}^{\lambda_m}), \end{aligned} \tag{20}$$

where

$$\bar{\omega}_{km} = \sum_{i=1}^{n'} A^{-1}_{ik} e^{-\alpha(k-\lambda_i)^2}. \tag{21}$$

Using Eqs. 20 and 21, we see that we can then write down the total potential energy of the full n-bead ring-polymer as

$$V_n(\mathbf{r}) = \sum_{k=1}^{n} V(\mathbf{r}^{(k)}) \simeq \sum_{k=1}^{n}\sum_{m=1}^{n'} \bar{\omega}_{km} V(\mathbf{r}^{\lambda_m}). \tag{22}$$

To obtain the forces on each of the n ring-polymer beads using the GPR interpolation PES of Eqs. 20 and 22, we can use the chain rule as follows:

$$\begin{aligned} \frac{\partial V_n}{\partial \mathbf{r}^{(j)}} &= \sum_{m=1}^{n'} \frac{\partial V_n}{\partial \mathbf{r}^{\lambda_m}} \frac{\partial \mathbf{r}^{\lambda_m}}{\partial \mathbf{r}^{(j)}}, \\ &= \sum_{m=1}^{n'} \left(\sum_{k=1}^{n} \bar{\omega}_{km} \frac{\partial V(\mathbf{r}^{\lambda_m})}{\partial \mathbf{r}^{\lambda_m}} \right) \frac{\partial \mathbf{r}^{\lambda_m}}{\partial \mathbf{r}^{(j)}}, \\ &= \sum_{m=1}^{n'} W_m \left(\frac{\partial V(\mathbf{r}^{\lambda_m})}{\partial \mathbf{r}^{\lambda_m}} \right) \left(\frac{\partial \mathbf{r}^{\lambda_m}}{\partial \mathbf{r}^{(j)}} \right), \end{aligned} \tag{23}$$

where

$$W_m = \sum_{k=1}^{n} \bar{\omega}_{km} \tag{24}$$

The final missing ingredient to obtain the required forces is the determination of the derivative of the positions of the selected n' beads with respect to the positions of the original n ring-polymer beads, as required in the final line of Eq. 23. This problem can itself be solved by using GPR, by expressing the position of each of the n' reference beads as an interpolation of the positions of the original n-bead ring-polymer. Noting that a separate interpolation will be different for each degree-of-freedom, we write

$$r^{\lambda_m}_{\eta} = \sum_{j=1}^{n} y_{\eta,j} e^{-\gamma(\lambda_m - j)^2}, \tag{25}$$

where η now defines the position component of GPR bead labelled λ_m, $y_{\eta,j}$ is the associated GPR weight function and γ is a further width parameter. As in Eq. 19, the weights can be written in terms of the positions of the n ring-polymer beads as

$$y_{\eta,j} = \sum_{i=1}^{n} C^{-1}_{ji} r^{(i)}_{\eta}, \tag{26}$$

where the kernel matrix $\mathbf{C}$ is of the same form as Eq. 17, albeit with different associated parameters,

$$C_{ij} = \delta_{ij}\sigma_r^2 + e^{-\gamma(i-j)^2}, \tag{27}$$

where σ_r is the assumed error parameter (or regularization parameter) for interpolation of positions. Combining

Eqs. 25 and 2, we then find

$$r_\eta^{\lambda_m} = \sum_{k=1}^{n} B_{mk} r_\eta^{(k)}, \tag{28}$$

where

$$B_{mk} = \sum_{j=1}^{n} C_{jk}^{-1} e^{-\gamma(j-m)^2}. \tag{29}$$

As a result of this additional interpolation, we find that the required derivative of the GPR bead positions with respect to the full ring-polymer bead positions is given as

$$\frac{\partial r_\eta^{\lambda_m}}{\partial r^{(j)}} = B_{mk}. \tag{30}$$

As a result, combining Eqs. 23 and 30, the forces on the full n-bead ring-polymer can be calculated from the GPR interpolation PES.

In summary, our RPI approach proceeds as follows at each time-step of a PIMD simulation:

1. Select the indices of the n' GPR reference points; this can be done by simply evenly distributing the n' reference points around n-bead ring-polymer.
2. For each DOF, calculate the required GPR interpolation weights using Eq. and calculate the interpolated positions using Eq. 25.
3. Evaluate the full PES at the n' GPR reference points.
4. Use GPR to evaluate the total PES on the ring-polymer (Eq. 22) and the forces on the n-bead ring-polymer (Eqs. 23 and 30).

This completes our description of our RPI approach. It is clear that RPI requires n' evaluations of the full PES at each PIMD time-step, compared to n evaluations required by full PIMD; as a result, one can expect that the computational effort of a RPI simulation should be roughly n'/n compared to that of full PIMD. RPI does require additional matrix equations to be solved to determine the weights for the PES and position interpolation but we note that, in typical PIMD simulations, the sizes of these matrices will typically be a few tens or less. As a final point, we emphasize that, unlike RPC, our RPI approach does not assume anything about the underlying PES of the system (beyond the usual smoothness assumption which is inherent to GPR); as a result, RPI as described here is directly applicable to any PES, including *ab initio* PESs and empirical force-fields.

III. RESULTS AND DISCUSSION

To assess the suitability of RPI as a method for accelerating PI simulations, we perform simulations of liquid water at 298 K, as described by the empirical q-TIP4P/F model.[16] This system has been employed extensively as a model for quantum effects in liquid water, ice and water clusters. We particularly focus on the convergence of quantum expectation values as a function of the number of ring-polymer beads employed in RPI and RPC simulations; comparison to full PIMD simulations provides a route to assessing efficiency and accuracy. Furthermore, RPC clearly represents the best current approach to accelerating convergence in such systems, thereby providing another convenient benchmark against which to assess RPI.

Before presenting RPI results, we first highlight our approach to determining the GPR parameters α and γ. First, to simplify matters, we assume that $\gamma = \alpha$, requiring a simple optimization of a single variable; this may, of course, not be the best choice in terms of ultimate accuracy, but a practical scheme for PI simulations should not require complex optimization of multiple parameters. Second, to determine the best α we adopt the simple approach of minimizing the root-mean-square error (RMSE) between the PES values given by RPI and the exact PES values on the full set of ring-polymer beads. In the examples discussed here, this RMSE is evaluated using 500 configurations taken from a short full PIMD simulation for the target system; evaluating the RMSE for different values of α enables one to select an appropriate value to perform larger RPI simulations.

In modelling liquid water, the same system set-up was used for PIMD simulations (systematically increasing the number of ring-polymer beads up to $n = 32$), RPI simulations (using a full set of $n = 32$ beads, but with varying number of GPR reference beads) and RPC simulations (again, using a full set of $n = 32$ beads, but varying the number of contracted beads). A system of 125 water molecules at a temperature of $T = 298$ K and density of $\rho = 0.997$ g cm^{-3} was equilibrated (with an Anderson thermostat) for ??? ps. After equilibration, static thermal averages were calculated in a further constant-NVT simulation of ??? ps. Furthermore, the quantum diffusion coefficient was calculated using RPMD.[16,25,26,29,32–41] Periodic boundary conditions were implemented using the minimum-image convention. The Ewald summation was used to calculate electrostatic interactions, and a cut-off of 9 Å was employed in the calculation of the short-range contribution to the Ewald summation energy and the Lennard-Jones term. Properties of interested were averaged over five independent calculations, providing error estimates.

The convergence of the thermally-averaged potential energy calculated by RPI with respect to the number of GPR reference beads is illustrated in Fig. 2; given that the number of GPR reference beads represents the number of full PES evaluations which must be performed during each simulation timestep, it is desirable that average observable values converge quickly to the correct result (here, taken to be the full PIMD simulation with $n = 32$ beads) with as few GPR points as possible. In

the case of the RPC method, we consider two alternatives; in the first case, labelled RPC(EI), the intramolecular interactions are calculated explicitly on *all* 32 ring-polymer beads, whereas the intermolecular PES contribution is calculated on the number of ring-polymer beads referenced on the x-axis. This RPC(EI) method is the original implementation of RPC, which exploits the fact that, for simple potentials of the form considered here, one can easily identify intramolecular and intermolecular contributions, enabling this 'trick' to be implemented efficiently; as we have noted above, this standard RPC(EI) method cannot be applied to non-separable potentials (*e.g.* DFT) without further modifications. The second RPC method we consider here, labelled as RPC(full), does not employ separation of intramolecular and intermolecular terms; instead, ring-polymer contraction is simply applied to the total PES. This RPC(full) approach is not condoned in any way, and is clearly not the way in which RPC should be applied, but, by comparing the convergence of RPC(full) and RPI, we can highlight the fundamentally different approaches taken in these methods.

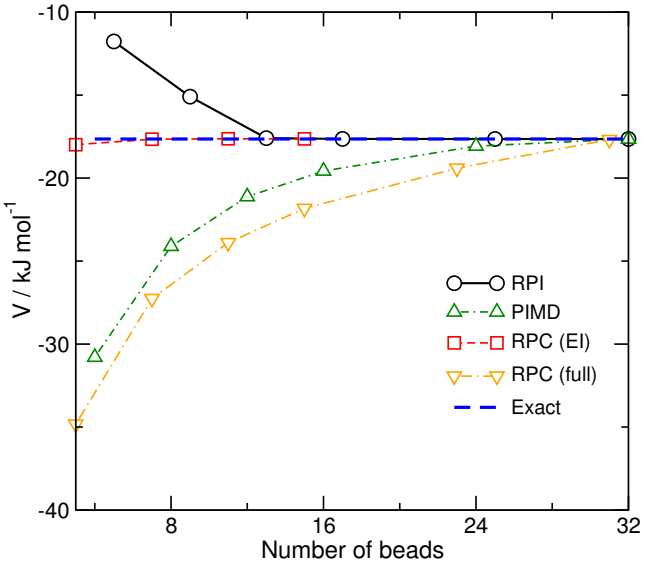

FIG. 2

Figure 2 clearly demonstrates that the new RPI method converges quickly on the correct thermally-averaged value of the potential energy as the number of GPR reference beads is increased; it is worth bearing in mind that *all* RPI simulations used $n = 32$ ring-polymer beads in total, and it is only the number of GPR reference beads which is increased in these RPI convergence tests. We find that, once the number of GPR reference beads, n', is great than about 13, one obtains essentially the exact, fully-converged $n = 32$-bead PIMD result for the average potential energy; in other words, RPI reduces the number of PES and force evaluations required to obtain exact results by a factor of around 2.5. This convergence property is much faster than in a standard PIMD simulation, where it is found that there is still a significant error in the average PES value when one uses $n = 24$ ring-polymer beads. Of course, the RPC(EI) method very rapidly converges on the correct answer for this potential, requiring only about seven ring-polymer beads for evaluation of the *intermolecular* terms (noting again that the *intramolecular* term is evaluated explicitly on 32 ring-polymer beads). Finally, as expected, applying RPC to the full PES is not very successful, and it is found that the convergence of the average potential energy is worse than the standard PIMD case. This behaviour arises because one is using PES evaluations on a contracted ring-polymer to approximate the PES on the full ring-polymer; in contrast, the RPI approach, using interpolation around the ring-polymer, is clearly capable of sufficiently approximation the full ring-polymer PES with just a few reference points.

The same trend in convergence of the different PIMD acceleration methods is evident in Fig. 3, which shows results for calculations of the average quantum kinetic energy of the system, calculated using the virial estimator.[70] As in Fig. 2, and as expected, RPC(EI) converges on the corret results using $n = 7$ ring-polymer beads for the intermolecular term and $n = 32$ beads for the intramolecular term, whereas the RPC(full) method again converges more slowly than even the standard PIMD method. In the case of the kinetic energy expectation value, we find that the exact value is obtained when using $n' \geq 17$, which is slightly worse than in the case of the potential energy (Fig. 2), but still results in a simulation which requires roughly half the number of PES and force evaluations to obtain the exact results when compared to PIMD. The difference in convergence between Figs. 2 and 3 is most likely due to the fact that potential energy values were used to optimize the value of the GPR width-parameter $\gamma = \alpha$, as described above, and then the optimized values for each n' were used to calculate all other properties. It seems likely that, if one were to optimise the GPR parameters to simultaneously match both kinetic energy and potential energy for selected configurations, the convergence of RPI illustrated in Fig. 3 should improve. Nevertheless, the results of both Figs. 2 and 3 clearly demonstrate that RPI can indeed converge on exact quantum result using around 50% of the PES and forces evaluations required by standard PIMD.

While calculating individual averaged values, such as quantum potential or kinetic energy, is a good demonstration that RPI is converging on the correct PIMD properties, a further test is in assessing whether the larger-scale structural properties of the system are correct. Figure 4 shows the O−−H radial distribution function (RDF), $g_{\mathrm{OH}}(r)$ calculated using RPI with increasing numbers of GPR reference beads; we chose to illustrate this particular RDF because the O−−H distance is strongly sensitive to the correct incorporation of nuclear quantum effects, so any errors in treatment of such fluctuations in the RPI simulations should be evident.

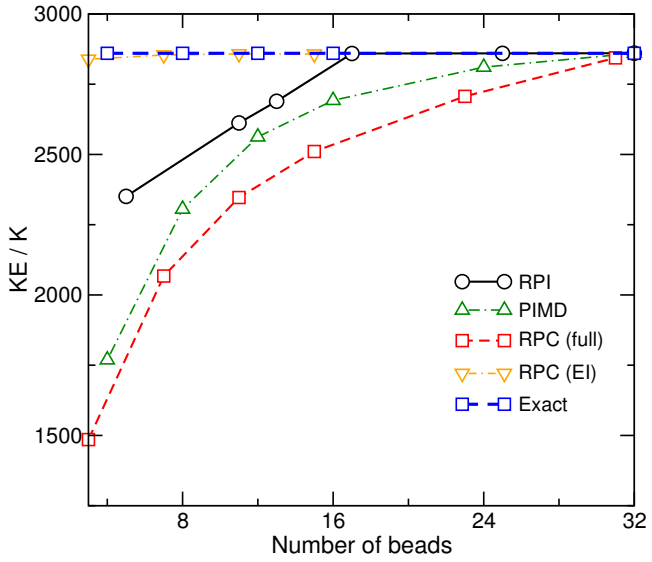

FIG. 3

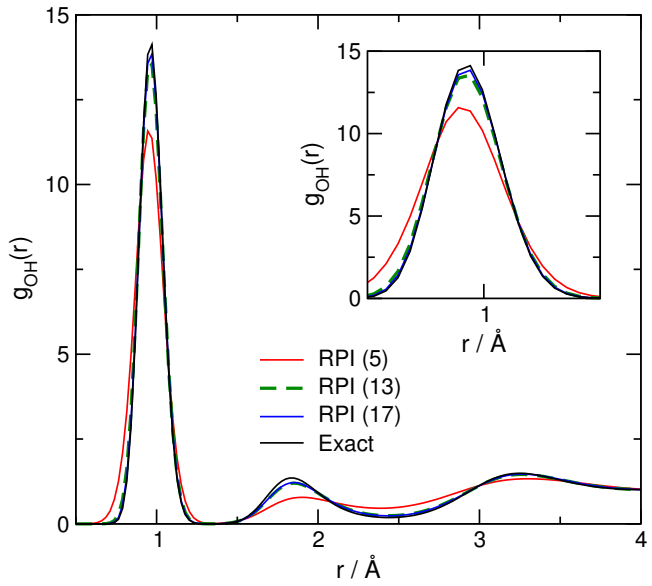

FIG. 4

We find that the RDF calculated by RPI is essentially converged using about $n' = 13$ GPR reference beads, with very small error relative to the exact PIMD result. This simulation clearly demonstrates that the atomic-level structure obtained in RPI simulations is the same as that obtained in a full PIMD simulation.

The overall conclusion of the results of Figs. 2-4 is that RPI can clearly converge on exact PIMD simulation results, yet reduces the number of required PES and force evaluations by at least a factor of two relative to converged PIMD. Perhaps more importantly, we emphasize that RPI is *directly* applicable to more complex PESs, including those generated by *ab initio* electronic structure methods. Compared to evaluation of the PES, the additional computational overhead required for RPI is very small, and the method itself only requires minor modifications to any standard PIMD code. As a final point, we note that RPI could also be used within the context of other RPC methods, notably the reference-potential based schemes proposed recently,[20,51] providing a hybrid method which simultaneously accelerates the overall PIMD scheme and the evaluation of the reference potential on the ring-polymer.

IV. CONCLUSIONS

In this Article, we have suggested a new approach to accelerating path-integral simulations. Instead of relying on evaluating the PES and forces on a contracted ring-polymer, as has been proposed previously, our RPI method instead uses the idea of interpolating the PES values around the ring-polymer, using PES evaluations at just a few selected positions around the ring-polymer. For liquid water at 298 K, described with the q-TIP4P/F empirical water model, we have found that RPI converges on exact PIMD results, but requires just around 50% of the total PES and force evaluations. In contrast to RPC-based methods, RPI is directly applicable to any PES, including *ab initio* methods, and requires only small modifications to any existing PIMD code. As noted above, we are now exploring how RPI could be combined with existing RPC methods to further accelerate convergence of properties in PIMD simulations.

ACKNOWLEDGEMENTS

SJB gratefully acknowledges award of a studentship from the EPSRC. Computing facilities were provided by the Scientific Computing Research Technology Platform (SC-RTP) of the University of Warwick. Data from Figs. ??....

[1] R. P. Feynman and A. R. Hibbs, *Quantum mechanics and path integrals* (McGraw-Hill, New York, 1965).
[2] D. Marx and J. Hutter, *Ab initio molecular dynamics: Basic theory and advanced methods* (Cambridge University Press, 2009).
[3] A. M. Reilly, S. Habershon, C. A. Morrison, and D. W. H. Rankin, J. Chem. Phys. **132**, 094502 (2010).
[4] A. M. Reilly, S. Habershon, C. A. Morrison, and D. W. H. Rankin, J. Chem. Phys. **132**, 134511 (2010).
[5] T. E. Markland and D. E. Manolopoulos, J. Chem. Phys. **129**, 024105 (2008).
[6] M. Ceriotti, M. Parrinello, T. E. Markland, and D. E. Manolopoulos, J. Chem. Phys. **133**, 124104 (2010).
[7] M. Ceriotti, D. E. Manolopoulos, and M. Parrinello, J. Chem. Phys. **134**, 084104 (2011).
[8] M. Parrinello and A. Rahman, J. Chem. Phys. **80**, 860 (1984).
[9] R. Ramírez, C. P. Herrero, A. Antonelli, and E. R. Hernández, J. Chem. Phys. **129**, 064110 (2008).
[10] J. Morales and K. Singer, Mol. Phys. **73**, 873 (1991).
[11] R. P. Steele, J. Zwickl, P. Shushkov, and J. C. Tully, J. Chem. Phys. **134**, 074112 (2011).
[12] J. R. Schmidt and J. C. Tully, J. Chem. Phys. **127**, 094103 (2007).

[13]S. Habershon and D. E. Manolopoulos, J. Chem. Phys. **135**, 224111 (2011).
[14]S. Habershon, Phys. Chem. Chem. Phys. **16**, 9154 (2014).
[15]T. Spura, C. John, S. Habershon, and T. D. Kühne, Mol. Phys. **113**, 808 (2015).
[16]S. Habershon, T. E. Markland, and D. E. Manolopoulos, J. Chem. Phys. **131**, 024501 (2009).
[17]A. Wallqvist and B. Berne, Chem. Phys. Lett. **117**, 214 (1985).
[18]M. Ceriotti, M. Parrinello, T. E. Markland, and D. E. Manolopoulos, J. Chem. Phys. **133**, 124104 (2010).
[19]V. Kapil, J. VandeVondele, and M. Ceriotti, J. Chem. Phys. **144**, 054111 (2016).
[20]O. Marsalek and T. E. Markland, J. Chem. Phys. **144**, 054112 (2016).
[21]R. A. Kuharski and P. J. Rossky, J. Chem. Phys. **82**, 5164 (1985).
[22]L. H. de la Pena and P. G. Kusalik, J. Chem. Phys. **121**, 5992 (2004).
[23]L. H. de la Pena and P. G. Kusalik, J. Chem. Phys. **125**, 054512 (2006).
[24]S. Habershon, G. S. Fanourgakis, and D. E. Manolopoulos, J. Chem. Phys. **129**, 074501 (2008).
[25]S. Habershon and D. E. Manolopoulos, J. Chem. Phys. **131**, 244518 (2009).
[26]T. F. M. III and D. E. Manolopoulos, J. Chem. Phys. **123**, 154504 (2005).
[27]B. Chen, I. Ivanov, M. L. Klein, and M. Parrinello, Phys. Rev. Lett. **91**, 215503 (2003).
[28]C. Alhambra, J. Gao, J. C. Corchado, J. Villà, and D. G. Truhlar, J. Am. Chem. Soc. **121**, 2253 (1999).
[29]N. Boekelheide, R. Salomón-Ferrer, and T. F. Miller III, Proc. Nat. Acad. Sci. USA **108**, 16159 (2011).
[30]D. G. Truhlar, J. Gao, C. Alhambra, M. Garcia-Viloca, J. Corchado, M. L. Sánchez, and J. Villà, Acc. Chem. Res. **35**, 341 (2002).
[31]M. H. M. Olsson, P. E. M. Siegbahn, and A. Warshel, J. Am. Chem. Soc. **126**, 2820 (2004).
[32]I. R. Craig and D. E. Manolopoulos, J. Chem. Phys. **121**, 3368 (2004).
[33]I. R. Craig and D. E. Manolopoulos, J. Chem. Phys. **122**, 084106 (2005).
[34]I. R. Craig and D. E. Manolopoulos, J. Chem. Phys. **123**, 034102 (2005).
[35]T. F. Miller and D. E. Manolopoulos, J. Chem. Phys. **122**, 184503 (2005).
[36]B. J. Braams and D. E. Manolopoulos, J. Chem. Phys. **125**, 124105 (2006).
[37]A. R. Menzeleev and T. F. Miller, J. Chem. Phys. **132**, 034106 (2010).
[38]S. Habershon, B. J. Braams, and D. E. Manolopoulos, J. Chem. Phys. **127**, 174108 (2007).
[39]T. E. Markland, S. Habershon, and D. E. Manolopoulos, J. Chem. Phys. **128**, 194506 (2008).
[40]X. C. Huang, S. Habershon, and J. M. Bowman, Chem. Phys. Letters **450**, 253 (2008).
[41]S. Habershon, D. E. Manolopoulos, T. E. Markland, and T. F. Miller, Annu. Rev. Phys. Chem. **64**, 387 (2013).
[42]T. J. H. Hele, M. J. Willatt, A. Muolo, and S. C. Althorpe, J Chem Phys **142**, 191101 (2015).
[43]E. Geva, Q. Shi, and G. A. Voth, J. Chem. Phys. **115**, 9209 (2001).
[44]G. A. Voth, Adv. Chem. Phys. **93**, 135 (2007).
[45]S. Jang and G. A. Voth, J. Chem. Phys. **111**, 2371 (1999).
[46]J. Cao and G. A. Voth, J. Chem. Phys. **100**, 5106 (1994).
[47]F. Paesani and G. A. Voth, J. Chem. Phys. **132**, 014105 (2010).
[48]J. O. Richardson and S. C. Althorpe, J. Chem. Phys. **131**, 214106 (2009).
[49]A. R. Menzeleev, N. Ananth, and T. F. Miller, J. Chem. Phys. **135**, 074106 (2011).
[50]T. E. Markland and D. E. Manolopoulos, Chem. Phys. Letters **464**, 256 (2008).
[51]C. John, T. Spura, S. Habershon, and T. D. Kühne, Phys. Rev. E **93**, 043305 (2016).
[52]T. L. Fletcher and P. L. A. Popelier, J. Chem. Theory Comput. **12**, 2742 (2016).
[53]C. E. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning* (MIT Press, 2006).
[54]S. M. Kandathil, T. L. Fletcher, Y. Yuan, J. Knowles, and P. L. A. Popelier, J. Comput. Chem. **34**, 1850 (2013).
[55]M. J. L. Mills and P. L. A. Popelier, Comp. Theo. Chem. **975**, 42 (2011).
[56]M. J. L. Mills and P. L. A. Popelier, Theor. Chem. Acc. **131**, 1 (2012).
[57]A. P. Bartók, M. J. Gillan, F. R. Manby, and G. Csányi, Phys. Rev. B **88**, 054104 (2013).
[58]A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, Phys. Rev. Lett. **104**, 136403 (2010).
[59]A. P. Bartók and G. Csányi, Int. J. Quantum Chem. **115**, 1051 (2015).
[60]D. Chandler, *Introduction to Modern Statistical Mechanics* (Oxford University Press, Oxford, UK, 1987).
[61]M. E. Tuckerman, *Statistical Mechanics: Theory and molecular simulation* (Oxford University Press, 2012).
[62]T. J. H. Hele and S. C. Althorpe, J. Chem. Phys. **138**, 084108 (2013).
[63]T. J. H. Hele, M. J. Willatt, A. Muolo, and S. C. Althorpe, J. Chem. Phys. **142**, 134103 (2015).
[64]G. S. Fanourgakis, T. E. Markland, and D. E. Manolopoulos, J. Chem. Phys. **131**, 094102 (2009).
[65]G. Krilov and B. J. Berne, J. Chem. Phys. **111**, 9147 (1999).
[66]C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning* (The MIT Press, Cambridge, Massachusetts, 2006).
[67]J. P. Alborzpour, D. P. Tew, and S. Habershon, J. Chem. Phys. **145**, 174112 (2016).
[68]L. Mones, N. Bernstein, and G. Csányi, **12**, 5110 (2016).
[69]E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen, *LAPACK Users' Guide*, 3rd ed. (Society for Industrial and Applied Mathematics, Philadelphia, PA, 1999).
[70]M. F. Herman, E. J. Bruskin, and B. J. Berne, J. Chem. Phys. **76**, 5150 (1982).