

# High Order Path Integrals Made Easy

Venkat Kapil,<sup>1</sup> Jörg Behler,<sup>2</sup> and Michele Ceriotti<sup>1, a)</sup>

<sup>1)</sup>*Laboratory of Computational Science and Modelling, Institute of Materials, Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland*

<sup>2)</sup>*Lehrstuhl für Theoretische Chemie, Ruhr-Universität Bochum, Bochum, Germany*

(Dated: 3 January 2017)

The precise description of quantum nuclear fluctuations in atomistic modelling is possible by employing path integral techniques, which involve a considerable computational overhead due to the need of simulating multiple replicas of the system. Many approaches have been suggested to reduce the required number of replicas. Among these, high-order factorizations of the Boltzmann operator are particularly attractive for high-precision and low-temperature scenarios. Unfortunately, to date several technical challenges have prevented a widespread use of these approaches to study nuclear quantum effects in condensed-phase systems. Here we introduce an inexpensive molecular dynamics scheme that overcomes these limitations, thus making it possible to exploit the improved convergence of high-order path integrals without having to sacrifice the stability, convenience and flexibility of conventional second-order techniques. The capabilities of the method are demonstrated by simulations of liquid water and ice, as described by a neural-network potential fitted to dispersion-corrected hybrid density functional theory calculations.

## I. INTRODUCTION

Molecules and materials that contain light nuclei - most notably hydrogen - exhibit considerable deviations from classical behavior, which are most pronounced at cryogenic temperatures but extend up to and even above room temperature<sup>1-3</sup>. Examples of such nuclear quantum effects (NQE) include a heat capacity that deviates dramatically from the Dulong-Petit limit<sup>4</sup>, equilibrium fractionation of isotopomers between different phases of a given compound<sup>5</sup> or different molecular sites<sup>6</sup>, non-Maxwell-Boltzmann distribution of particle velocities<sup>7</sup> as well as dynamical properties that differ from the predictions obtained from classical molecular mechanics<sup>8</sup>.

The interest in modelling NQEs in atomistic simulations has been growing constantly over the past years - not only because faster computers and more efficient algorithms have made such tasks more accessible, but also because of a paradigm change regarding the employed inter-atomic potentials. In particular, *ab initio* simulations solving explicitly the electronic structure problem at increasingly accurate levels of theory “on-the-fly”<sup>9</sup> as well as next-generation potentials<sup>10-12</sup>, that are designed to reproduce *ab initio* reference data rather than experiments, have made the need of modelling NQEs an urgent matter. This is because in the absence of empirical fitting parameters, which is a mandatory condition for predictive simulations, these *ab initio* studies employ the Born-Oppenheimer energy surface, without any term that could implicitly compensate for the lack of zero-point energy and tunnelling. Therefore, for predictive simulations it is insufficient to focus on the inter-atomic forces alone. NQEs also need to be included explicitly on a sound physical basis to achieve the highest possible accuracy.

The techniques of choice for treating the quantum nature of nuclei are path integral molecular dynamics (PIMD), and path integral Monte Carlo (PIMC) simulations<sup>13</sup>. Using an elegant mapping of the quantum mechanical partition function onto the classical partition function of an extended system composed of several replicas of the atomic configuration, path integral methods make it possible to treat exactly the quantum statistics of distinguishable<sup>14</sup> (and indistinguishable<sup>13,15</sup>) particles, however at a much larger cost compared to a classical simulation.

In the past few years, several approaches have been proposed to reduce the overhead of such simulations by accelerating the convergence with the number of replicas. This goal has been achieved by computing expensive parts of the potential on a reduced number of replicas<sup>16,17</sup>, by using a thermostat described by a generalized Langevin equation (GLE) to artificially generate the proper quantum fluctuations<sup>18-20</sup>, or by using a higher-order expansion of the quantum partition function - so-called “high-order path integral techniques”<sup>21-24</sup>. This latter approach is very appealing, particularly if low-temperature or high-accuracy are sought, since the convergence improves with the number of replicas  $P$  from  $P^{-2}$  of traditional methods to  $P^{-4}$ . Unfortunately, the higher order expansion introduces some cumbersome terms in the forces that depend formally on the Hessian of the physical potential. Therefore, rather than integrating the high-order PIMD equations of motion directly, research has focused on re-weighting schemes<sup>25-27</sup>, which are however affected by statistical inefficiency that worsens as system size increases<sup>28</sup>. A truncated cumulant expansion has recently shown considerable promise, although *ad hoc* estimators need to be devised specifically for different system properties<sup>29</sup>.

In the present work we will demonstrate that performing full high-order PIMD can be achieved using a symplectic finite-difference integrator. In this approach there is only a modest computational overhead, that is quickly

---

<sup>a)</sup>Electronic mail: michele.ceriotti@epfl.ch

paid off when high-accuracy or low-temperature simulations are to be performed. Moreover, we will show that it is possible to combine high-order path integrals with colored-noise acceleration techniques, although in practice there is only a small advantage relative to GLE techniques applied on top of second-order PIMD. We will demonstrate the capabilities of the method using liquid water described by a neural network (NN) potential<sup>10,30</sup> as an example.

## II. METHODS

### A. Second-order and fourth-order path integrals

The most effective framework for treating the quantum mechanical behavior of distinguishable particles is based on the path integral formalism, that maps the quantum mechanical partition function  $Z = \text{tr}[e^{-\beta\hat{H}}]$  at the inverse temperature  $\beta = \frac{1}{k_B T}$ , onto a classical partition function in an extended “ring polymer” phase space. This mapping corresponds to the application of the identity  $e^{-\beta\hat{H}} = \left[e^{-\beta\hat{H}/P}\right]^P$ , followed by a high-temperature expansion of the Boltzmann operator  $e^{-\beta_P\hat{H}}$  – where we have introduced the shorthand notation  $\beta_P = \beta/P$ . The most commonly used approach relies on a Trotter decomposition in terms of the potential operator  $\hat{V}$  and the kinetic energy operator  $\hat{T}$ ,

$$e^{-\beta\hat{H}} \approx \left[e^{-\beta_P\frac{\hat{V}}{2}} e^{-\beta_P\hat{T}} e^{-\beta_P\frac{\hat{V}}{2}}\right]^P, \quad (1)$$

which leads to an expansion of the (low temperature) Boltzmann operator that is accurate up to second order in  $\beta_P$ . It is easy to show that the application of the Trotter splitting yields a classical-like partition function. For a system of  $N$  distinguishable particles with masses  $\{m_i\}$  evolving under a potential  $V(\mathbf{q}_1, \dots, \mathbf{q}_P)$  at an inverse temperature of  $\beta$ , the ring polymer Hamiltonian reads  $\mathcal{H}_P^{\text{tr}}(\mathbf{p}, \mathbf{q}) = \mathcal{H}_P^0(\mathbf{p}, \mathbf{q}) + \mathcal{V}_P(\mathbf{q})$ . Here, the free ring polymer Hamiltonian is

$$\mathcal{H}_P^0(\mathbf{p}, \mathbf{q}) = \sum_{i=0}^{N-1} \sum_{j=0}^{P-1} \left( \frac{[\mathbf{p}_i^{(j)}]^2}{2m_i} + \frac{1}{2} m_i \omega_P^2 [\mathbf{q}_i^{(j)} - \mathbf{q}_i^{(j+1)}]^2 \right) \quad (2)$$

and the physical potential term  $\mathcal{V}_P(\mathbf{q})$  is just a sum over the potential  $V$  evaluated for the various replicas

$$\mathcal{V}_P(\mathbf{q}) = \sum_{j=0}^{P-1} V(\mathbf{q}_1^{(j)}, \dots, \mathbf{q}_N^{(j)}). \quad (3)$$

The “beads” in the ring polymer are connected cyclically (i.e.  $j + P \equiv j$  in Eq. (2)) by springs of frequency  $\omega_P = 1/\beta_P$ . Sampling has to be performed at the inverse temperature  $\beta_P$ .

High-order path integral schemes rely on a more accurate decomposition of the density matrix than in Eq. (1),

which however requires including terms that depend on the commutator between  $\hat{V}$  and  $\hat{T}$ . Several of these schemes have been proposed<sup>21</sup>, all however having similar advantages and shortcomings. We will focus in particular on the Suzuki-Chin (SC) decomposition<sup>22,23</sup>, that is accurate up to fourth order in  $\beta_P$ , and reads

$$e^{-\beta\hat{H}} \approx \left[ e^{-\beta_P\frac{\hat{V}_e}{3}} e^{-\beta_P\hat{T}} e^{-\beta_P\frac{4\hat{V}_o}{3}} e^{-\beta_P\hat{T}} e^{-\beta_P\frac{\hat{V}_e}{3}} \right]^{\frac{P}{2}}, \quad (4)$$

where

$$\hat{V}_e = \hat{V} + \frac{\alpha}{6} \beta_P^2 [\hat{V}, [\hat{T}, \hat{V}]], \quad (5)$$

$$\hat{V}_o = \hat{V} + \frac{(1-\alpha)}{12} \beta_P^2 [\hat{V}, [\hat{T}, \hat{V}]]. \quad (6)$$

$\alpha \in [0, 1]$  is an arbitrary parameter and can be adjusted to improve the convergence for a given problem. It seems, however, that no generally-applicable prescription for its choice can be obtained. In the present study, for reasons that will become apparent later on, we always used  $\alpha = 0$ . The main advantage of the SC scheme is that any structural observable can be computed seamlessly by averaging over the even beads in the path, without the complex correction terms that often enter estimators in other high-order schemes.

Following the same procedure as in the Trotter case, one can obtain a classical partition function based on the splitting (4), which is accurate up to fourth order in  $\beta_P$ . The Suzuki-Chin ring-polymer Hamiltonian  $\mathcal{H}_P^{\text{sc}}(\mathbf{p}, \mathbf{q}) = \mathcal{H}_P^0 + \mathcal{V}_P^{\text{sc}}(\mathbf{q})$  contains a modified potential term that acts differently on odd and even beads,

$$\mathcal{V}_P^{\text{sc}}(\mathbf{q}) = \sum_{j=0}^{P-1} \left( w_j V(\mathbf{q}^{(j)}) + \sum_{i=0}^{N-1} \frac{w_j d_j}{m_i \omega_P^2} |\mathbf{f}_i^{(j)}|^2 \right), \quad (7)$$

where  $\mathbf{f}_i^{(j)} = -\partial V(\mathbf{q}^{(j)}) / \partial \mathbf{q}_i^{(j)}$  is the physical force acting on the  $i$ -th atom in the  $j$ -th replica, and the scaling factors for odd and even beads are given by

$$\begin{aligned} w_j &= 2/3, & d_j &= \alpha/6 & j &\text{ is even,} \\ w_j &= 4/3, & d_j &= (1-\alpha)/12 & j &\text{ is odd.} \end{aligned} \quad (8)$$

Estimators for the Suzuki-Chin propagator can be derived by two routes. “Thermodynamic” (TD) estimators are obtained by applying thermodynamics identities to the SC ring polymer partition function, while “operator” (OP) estimators result from carrying out the splitting operation on an expression that already contains the quantum mechanical operator. The latter class of estimators is generally simpler to derive and evaluate, but in the present manuscript we will compare potential and kinetic energy operators computed using both routes. Expressions for both TD and OP-method estimators have been derived and reported several times<sup>25,31</sup>, but we also list them in Appendix C for the sake of completeness.

Evaluating the modified potential (7) only requires knowledge of the first derivative of the physical potential, which in a PIMD simulation has to be computed to

evolve the dynamics. However, evaluating the force associated with  $\mathcal{V}_P^{\text{sc}}$  is not so trivial, as it contains second derivatives of the potential,

$$\begin{aligned} \mathbf{f}_i^{\text{sc},(j)} &= w_j \mathbf{f}_i^{(j)} + \frac{2d_j}{\omega_P^2} \tilde{\mathbf{f}}_i^{(j)} = \\ &= w_j \mathbf{f}_i^{(j)} + \frac{2w_j d_j}{\omega_P^2} \sum_{k=0}^{N-1} \frac{\partial^2 V(\mathbf{q}^{(j)})}{\partial \mathbf{q}_i^{(j)} \partial \mathbf{q}_k^{(j)}} \frac{\mathbf{f}_k^{(j)}}{m_k}. \end{aligned} \quad (9)$$

## B. Finite-differences Suzuki-Chin PIMD

One should notice, however, that the expression for  $\tilde{\mathbf{f}}$  involves the second derivative of  $V$  *projected on the mass-scaled force*. As it has been recognized in the context of high-order path integral Monte Carlo<sup>32,33</sup>, and similarly to what has been done for instance in electronic structure theory<sup>34</sup>, it is possible to evaluate this kind of projected second derivatives by finite differences (FD), using either a symmetric estimator

$$\tilde{\mathbf{f}}_i^{(j)} = \lim_{\epsilon \rightarrow 0} \frac{1}{2\epsilon\delta} \left[ \mathbf{f}_i^{(j)}(\mathbf{q}^{(j)} + \epsilon\delta \mathbf{u}^{(j)}) - \mathbf{f}_i^{(j)}(\mathbf{q}^{(j)} - \epsilon\delta \mathbf{u}^{(j)}) \right] \quad (10)$$

or an asymmetric FD formula

$$\tilde{\mathbf{f}}_i^{(j)} = \lim_{\epsilon \rightarrow 0} \frac{1}{\epsilon\delta} \left[ \mathbf{f}_i^{(j)}(\mathbf{q}^{(j)} + \epsilon\delta \mathbf{u}^{(j)}) - \mathbf{f}_i^{(j)}(\mathbf{q}^{(j)}) \right], \quad (11)$$

where

$$\mathbf{u}_i^{(j)} = \frac{\mathbf{f}_i^{(j)}}{m_i}, \quad \delta = \left[ \frac{1}{NP} \sum_k \frac{\mathbf{f}_k^{(j)} \cdot \mathbf{f}_k^{(j)}}{m_k^2} \right]^{-1/2} \quad (12)$$

indicate a displacement that is parallel to the mass-scaled force, and a normalization coefficient so that  $\epsilon$  corresponds to the root mean square displacement applied on each atom when computing the derivative.

The expression in Eq. (10) can be used seamlessly to propagate the equations of motion and to evaluate the estimators for thermodynamic and structural properties listed in Appendix C. The crucial aspect that makes this procedure viable is that using any of the two finite-difference estimators for the derivative yields a rigorously time-reversible and symplectic integrator, when combined with a symmetric Trotter split velocity-Verlet integrator (see Appendix A). As a result, the scheme is stable even for relatively large values of the finite-difference step, which is advantageous e.g. when evaluating the forces in *ab initio* calculations, where residual errors in the convergence of the self-consistent solution to the electronic structure problem inevitably lead to noisy forces. It is also important to note that, when setting  $\alpha = 0$ ,  $\tilde{\mathbf{f}}^{(j)}$  only needs to be evaluated for odd beads, so that in practice the forward-backward evaluation of the derivatives costs as much as one full force evaluation, making this approach twice as expensive as a Trotter PIMD simulation with the same number of beads. As

we will show later, the asymmetric finite-difference estimator appears to be only marginally less stable than its symmetric counterpart, which means that with a judicious choice of  $\epsilon$  one can reduce the number of force evaluations by 50%.<sup>35</sup> This relatively small overhead can be reduced even further by using a multiple-time step (MTS) integrator for  $\tilde{\mathbf{f}}$  (see Appendix B), and is quickly compensated by the much faster asymptotic convergence. Our finite-difference Suzuki-Chin scheme is already advantageous at room temperature, and its lead becomes substantial for low-temperature or high-accuracy studies. Similar expressions can be easily derived in the context of other high-order factorizations such as that introduced by Takahashi and Imada<sup>21</sup>, and it is possible that perturbed path estimators<sup>29</sup> could be derived on top of a full fourth-order path integral Hamiltonian, providing even faster convergence to quantum expectation values. The availability of projected force derivatives also facilitates the implementation of the fourth-order version of estimators for the heat capacity<sup>26</sup> and for isotope fractionation ratios<sup>36</sup>. Finally, further dramatic speed-ups can be obtained whenever one can apply range-separation techniques such as ring-polymer contraction<sup>16</sup>, since we have made sure that our implementation in i-PI<sup>37</sup> is fully compatible with that of conventional real and imaginary-time multiple time stepping<sup>38</sup>.

## C. A generalized Langevin equation for high-order path integrals

Having access to direct sampling of  $\mathcal{H}_P^{\text{sc}}$  opens up the possibility of combining high-order path integrals with a generalized Langevin equation acceleration. In the Trotter case, the normal mode (NM) eigenvectors of the Hamiltonian for a harmonic potential  $V(q) = m\omega^2 q^2/2$  do not depend on the frequency  $\omega$  itself, that only leads to a shift to the NM frequencies. This makes it possible to apply sophisticated thermostating strategies, with different GLEs applied to individual NM coordinates – all without the need of knowing the NM decomposition of the physical potential<sup>20</sup>. Unfortunately, this is not the case for the Suzuki-Chin Hamiltonian. However, since the NM transformation remains an orthogonal transformation, it is possible to apply a single GLE to all Cartesian (or Trotter NM) coordinates which gives the same effect as applying such GLE onto the proper SC NMs. Starting from this observation, one can – with considerable effort, see Appendix D – derive a frequency-dependent effective temperature  $T^*(\omega)$ , that enforces different fluctuations on different ring-polymer vibrations so as to obtain converged quantum expectation values for *any* OP-method estimator of position-dependent properties in the harmonic limit, and for any number of beads. Contrary to the Trotter case, where one can further tune ring-polymer fluctuations to accelerate the convergence of the centroid-virial kinetic energy estimator, this is not possible here, so we can expect that the convergence of

the quantum kinetic energy will be less efficient than with the PIGLET approach<sup>20</sup>. GLE parameters enforcing the desired temperature curve for this Suzuki-Chin GLE (SC+GLE) approach have been obtained following the fitting protocol discussed in Ref.<sup>39</sup>, and are available for download from an on-line repository<sup>40</sup>.

### III. NEURAL NETWORK WATER: A BENCHMARK

For a comprehensive benchmark study of the methods discussed in the previous section we will use simulations of water, a prototypical system for the investigation of nuclear quantum effects. For this purpose we use a neural network (NN) potential fitted to *ab initio* calculations performed with the B3LYP hybrid density functional<sup>41</sup> and the D3 dispersion corrections by Grimme<sup>42</sup>, as implemented in CP2K<sup>43</sup>. The potential is fully reactive, i.e. it allows for the possibility of bond breaking and formation, and has recently been shown to provide an excellent description of nuclear quantum effects in water, as probed by isotope fractionation and the nuclear quantum kinetic energy<sup>44</sup>, at the same time allowing us to obtain thorough sampling. The potential was evaluated using a NN implementation<sup>45</sup> for LAMMPS<sup>46</sup>. Unless otherwise specified, each result we report involved a trajectory of at least 200 ps for a supercell containing 128 molecules at the experimental density. We enforced constant-temperature sampling at  $T = 300$  K using a PILE-G scheme<sup>47</sup> with  $\gamma_k = \omega_k/2$ , and a weak, global thermostat on the centroid – so that effectively canonical-sampling runs correspond to the thermostatted ring-polymer molecular dynamics (TRPMD) protocol<sup>48</sup> suitable to discuss quantum dynamical properties. In order to probe the behavior of our approach in a lower-temperature regime, we also performed simulations of a 96-molecules box of hexagonal ice at  $T = 100$  K. For colored-noise simulations we used the PIGLET thermostat<sup>20</sup> for Trotter PI, and the SC+GLE strategy discussed above for SC PIMD. These calculations will be a challenging test case for our techniques, because the reactive nature of the NN potential allows for quantum fluctuations of the hydrogen bond probing the strongly anharmonic regions in the potential energy surface of water<sup>49</sup>.

#### A. Stability of the finite-difference scheme

A possible problem that one has to be aware of when using force evaluation schemes based on discrete approximations of the derivatives is that in many cases – most notably for *ab initio* simulations – imperfect convergence of self-consistency schemes can introduce numerical noise. In particular, when using a small displacement in a finite difference scheme, the signal-to-noise ratio degrades, which can lead to instabilities in the integration of the equations of motion. Therefore it is important to test how sensitive are the results to the specific value of the

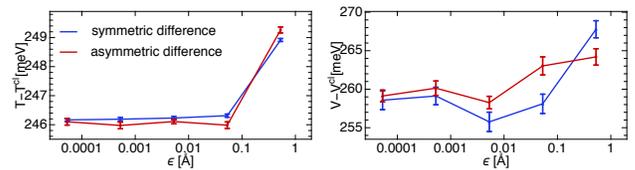


FIG. 1. Expectation values of the quantum contributions per molecule to the potential energy  $V$  as well as to the kinetic energy  $T$ , as a function of the finite-difference displacement  $\epsilon$ , for a SC-PIMD simulation of liquid water at 300 K performed with 16 beads. The two sets of points correspond to the symmetric (blue) and the asymmetric (red) finite-difference integrators. Error bars indicate the statistical error, which is of the order of 1% for the potential and of the order of 0.1% for the kinetic energy. The superscript “cl” indicates the classical component. Note that given the definition of the displacement vector  $\mathbf{u}$  in Eq. (12), the finite-difference step  $\epsilon$  indicates the root-mean-squared displacement of an atom during the evaluation of the derivative.

atomic displacement. Figure 1 shows that the SC integrators we introduce here, due to their time-reversibility and symplectic properties, show exceptionally good stability, with no appreciable effect of the root mean squared atomic displacement on the quantum expectation values of the potential and kinetic energies for  $\epsilon \lesssim 0.1\text{Å}$ . Over this broad range of displacements there is no appreciable difference between the symmetric and the forward FD estimator, so that the latter should be used whenever one does not parallelize fully the calculation over the  $P$  beads.

#### B. Convergence of energy estimators

The most straightforward measure for the convergence of a PIMD method to the quantum limit is given by the potential and kinetic energy estimators. Figure 2 shows such convergence tests, comparing Trotter and SC path integrals with and without colored noise. Results are in line with the expectations. Fourth-order PIMD gives a much improved asymptotic convergence, without the statistical instabilities observed in re-weighting strategies<sup>28</sup> and giving with  $P = 16$  results that are superior to Trotter PI with  $P = 32$ . The number of evaluation of  $\tilde{\mathbf{f}}_i^{(j)}$  can be reduced with a MTS scheme (see Appendix B). Even by computing the SC force as often as every  $M = 2$  steps, it can be clearly seen that also at room temperature our finite-differences implementation of high-order path integrals provides higher accuracy at a smaller cost than Trotter PIMD. As shown in figure 3, the improvement becomes even more significant as the temperature is lowered. In a simulation of hexagonal ice at  $T = 100\text{K}$ , SC PIMD reaches an error of a few meV per molecule when  $P = 48$ . When using Trotter PIMD, one would need to use more than 128 beads to obtain a similar accuracy.

GLEs improve significantly the convergence of both standard and fourth-order PIMD, giving potential ener-

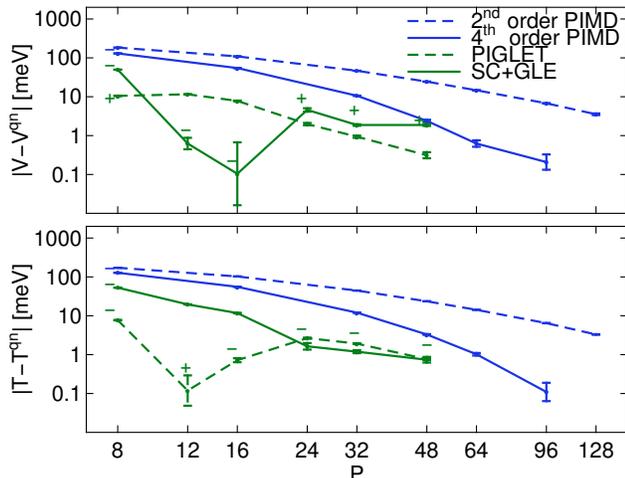


FIG. 2. Error per molecule on the value of the potential energy  $V$  as well as the kinetic energy  $T$ , as a function of the number of beads  $P$ , for a simulation of liquid water at 300 K performed with second and fourth-order PIMD, and with the corresponding colored-noise methods, namely PIGLET (second order) and SC+GLE (fourth order). We only report here the OP-method estimators – see Figure 8 for a comparison with the TD estimators. The fully-converged value is taken to be SC PIMD with  $P = 48$ , and errors are plotted on a log-log scale to highlight the faster convergence of fourth-order methods.

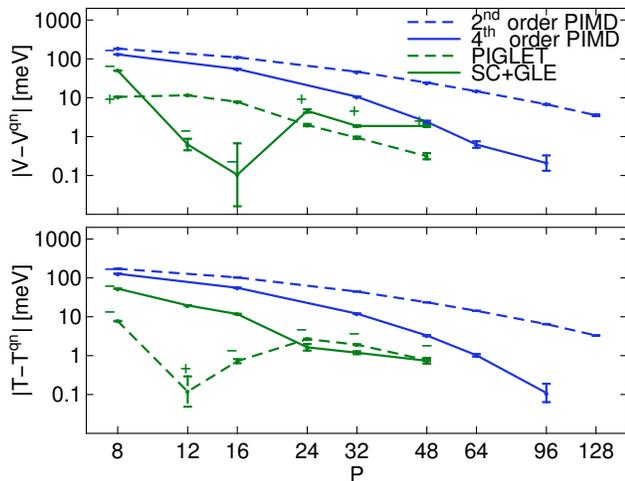


FIG. 3. Error per molecule on the value of the potential energy  $V$  as well as the kinetic energy  $T$ , as a function of the number of beads  $P$ , for a simulation of ice at 100 K performed with second and fourth-order PIMD, and with the corresponding colored-noise methods, namely PIGLET (second order) and SC+GLE (fourth order). We only report here the OP-method estimators – see Figure 8 for a comparison with the TD estimators. The fully-converged value is taken to be SC PIMD with  $P = 128$ , and errors are plotted on a log-log scale to highlight the faster convergence of fourth-order methods.

gies that are within a few percent of the converged results with as few as 4-6 beads for water, and 16-24 beads for ice. Although the GLE-thermostatted results are better than the canonically sampled PI simulations for all values of  $P$ , we observe that the convergence of GLE techniques is non-monotonic, similar to what was observed in simulations of small molecules at ultra-low temperature<sup>50</sup>. It appears that the convergence of SC+GLE is not better than that obtained by PIGLET, which underscores the fact that the limiting factor for convergence of GLE schemes has more to do with zero-point energy leakage between different modes than with the asymptotic convergence of the PI section of the method. SC+GLE results are more sensitive to the coupling strength of the colored noise than in the case of PIGLET, probably due to the more complex form of the full path integral Hessian in the harmonic limit (see Appendix D).

Although the possibility of combining high-order path integrals with correlated noise sampling might be beneficial in some specific cases – for instance when computing structural properties at ultra-low temperature – it seems that the best course of action should be to use PIGLET whenever an accuracy of a few percent is sufficient, and resort to SC PIMD with conventional thermostating whenever one wants (a) to reach the ultimate level of convergence, (b) to use sampling techniques (e.g. replica exchange) for which it is necessary to have a well-defined functional form for the phase-space density, or (c) to compute complicated estimators whose convergence is not accelerated by GLEs.

### C. Radial distribution functions

The radial pair correlation functions  $g(r)$  represent the most frequently used indicators of the structure of water. Figure 4 shows the convergence of a few key features in the O-O, O-H and H-H correlation functions. As it has already been noted<sup>51</sup>, for Trotter PIMD there is an interesting non-monotonic convergence behavior of the  $g_{OO}$  distribution function, that gets less structured when going from classical to 2 and 4 beads, and then becomes more structured when it approaches convergence. Such a trend can be seen as a manifestation of the competition between quantum effects in different vibrational modes, that progressively converge as the number of replicas is increased. Overall, the convergence of the radial distribution functions with  $P$  is fully consistent with what is observed for the energy estimators. The more strongly quantized degrees of freedom – such as the O-H stretch – show the slowest convergence, and the most dramatic improvements with SC path integrals and colored-noise techniques. GLE methods give very good agreement with 6-8 beads, but if a very high accuracy is required – as it is often needed in case of radial distribution functions, for which changes in the peak shapes of a few percent can be significant – SC PIMD with 16 beads gives the best performance/cost ratio.

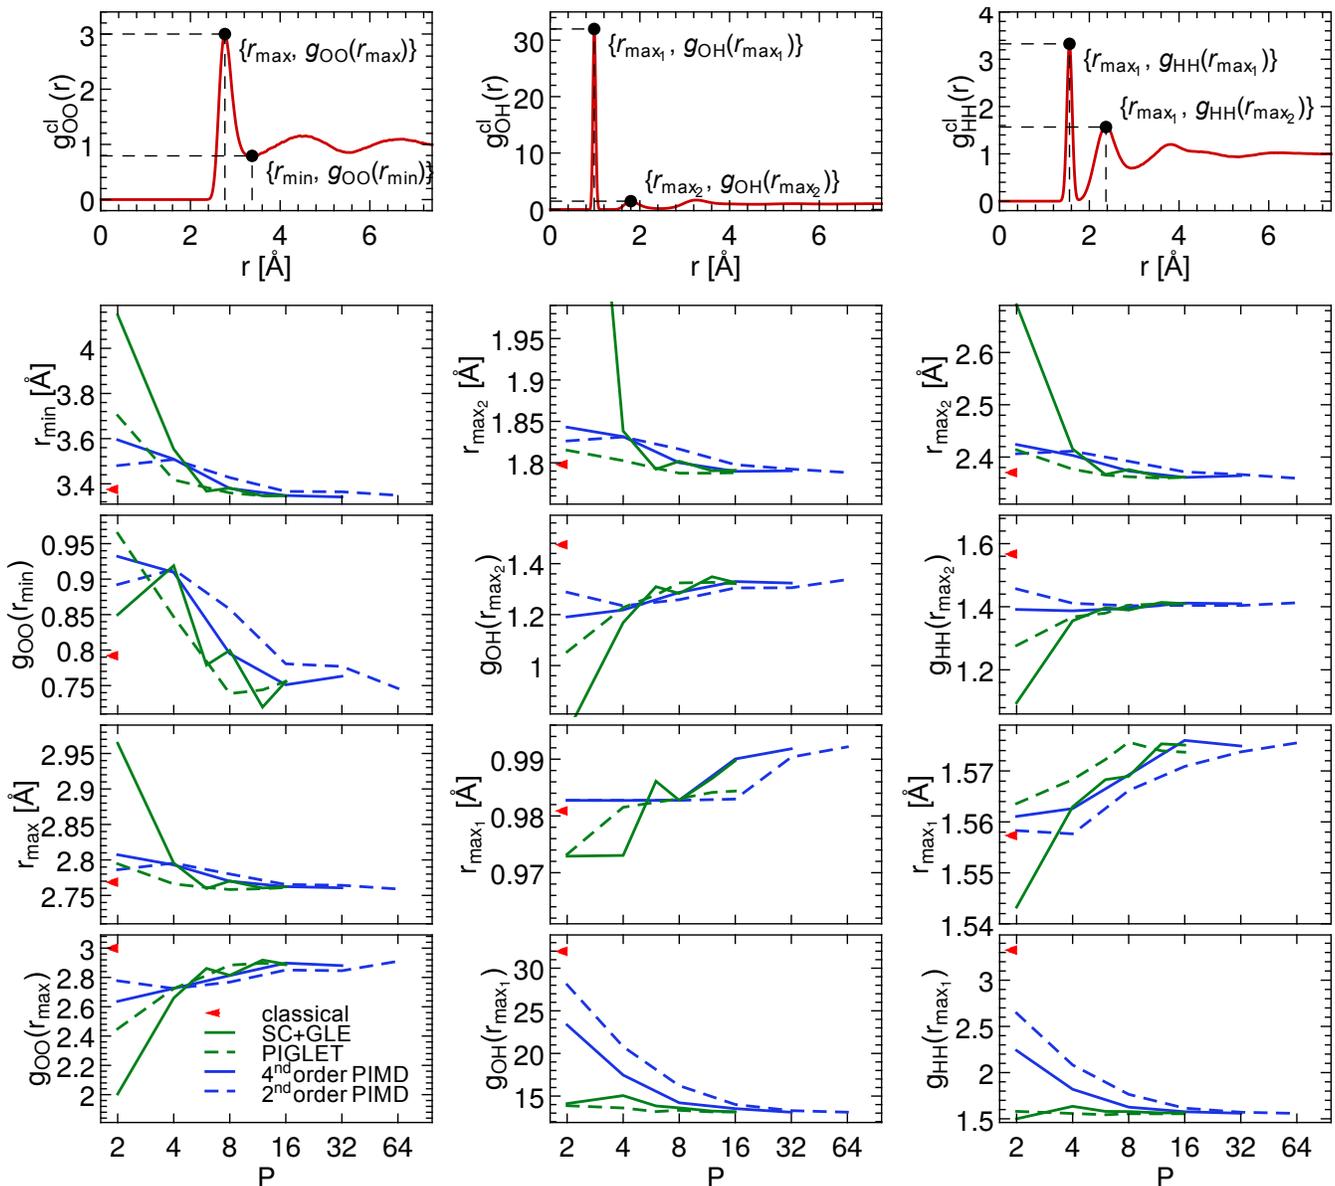


FIG. 4. The top panels show the radial distribution functions for O-O (left), O-H (center), and H-H (right) for the classical simulations along with some characteristic points whose convergence as a function of the number of replicas  $P$  is plotted in the four lower panels for each case. The red arrows show the corresponding values for a purely classical simulation.

#### D. H-bond fluctuations

One of the most remarkable effects of quantum fluctuations in room-temperature water is the occurrence of transient self-dissociation events, in which a quantum fluctuation momentarily brings a proton closer to the acceptor oxygen atom than to the oxygen it is covalently bound to<sup>49</sup>. The extent of these fluctuations is a particularly challenging quantity to compute, because of the small fraction of particles that undergo such broad excursions at any given time, the strong anharmonicity of the potential in this region, and the dependence on the level of electronic structure theory<sup>52</sup>. Figure 5 shows

the probability of having a proton mid-way between the donor and acceptor oxygen relative to the probability of the most common value of the proton transfer coordinate  $\nu$ . The convergence is very slow for all methods, with the exception of SC PIMD - although for  $P \leq 4$  fourth-order methods give dramatic over-estimation of these fluctuations. Colored-noise methods do accelerate convergence, but tend to yield too high fluctuations. For  $P = 6$ , PIMD would underestimate the fluctuations by a factor of 5, whereas PIGLET provides a too high value by a factor of 2. SC+GLE improves the convergence relative to PIGLET - an advantage that is however less significant when one considers the increase in

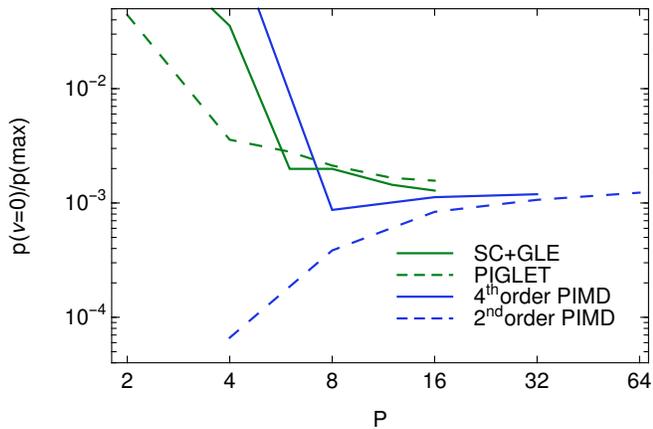


FIG. 5. Relative probability for observing proton delocalization over a H-bond.  $p(\nu)$  is the probability density relative to the proton-transfer coordinate  $\nu$ , and the plot reports the ratio between  $p(0)$  and the most likely H-bond configuration  $p(\max)$ , as a function of the number of replicas  $P$ . Note that the convergence is slower than for the energy in Fig. 2.

computational cost. Even when predicting strongly anharmonic fluctuations, GLE techniques make it possible to reach semi-quantitative accuracy quickly, and fourth-order path integrals are useful to reach full convergence in the asymptotic regime.

### E. Vibrational density of states

This far we have focused exclusively on static, time-independent properties. The path integral formalism is of a statistical mechanical nature, and strictly speaking no dynamical observable can be inferred. That said, several methods inspired by PIMD (such as centroid molecular dynamics, CMD<sup>53,54</sup> and ring polymer molecular dynamics, RPMD<sup>8,55</sup>) have been proposed to approximately estimate diffusion coefficients, vibrational spectra and other time-dependent quantities. For these benchmarks we will focus on thermostatted RPMD (TRPMD)<sup>48</sup>, a simple approach that can be seen as combining elements of CMD and RPMD, alleviating some of their most severe artifacts<sup>56</sup>, at the price however of a broadening of high-frequency peaks<sup>57</sup>. The idea is just to attach a Langevin thermostat to ring-polymer modes, with a damping coefficient adjusted to be proportional to the frequency of the mode in the free-particle limit. For  $V = 0$  there is no difference between the second and fourth-order Hamiltonians, and consequently the TRPMD approach can be applied in exactly the same way to a fourth-order simulation.

Generally, one performs (T)RPMD using a number of replicas that is sufficient to converge satisfactorily the static properties to their quantum values. In Figure 6 we investigate the convergence of the vibrational density of states (velocity-velocity correlation spectrum) of

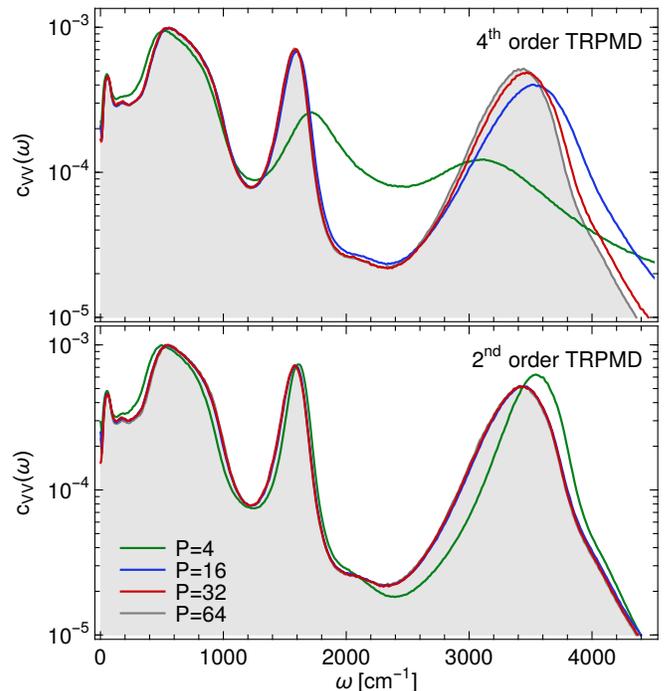


FIG. 6. Vibrational densities of states (Fourier-transforms of the velocity-velocity correlation functions) for TRPMD simulations of liquid water at 300 K, using a second-order Hamiltonian (lower panel) and a fourth-order Hamiltonian (upper panel). Simulations with  $P = 4, 16, 32$  (green, blue, red) are compared with a fully-converged Trotter TRPMD simulation (gray, shaded).

water with increasing numbers of replicas. Interestingly, the  $c_{vv}(\omega)$  converges faster than structural properties. When using  $\mathcal{H}_P^v$ ,  $P = 16$  is sufficient to obtain a vibrational spectrum that is indistinguishable from the fully-converged limit. On the other hand, convergence for  $\mathcal{H}_P^{sc}$  is dramatically slowed down. This is consistent with what observed in Ref.<sup>31</sup> for a harmonic potential and the closely-related case of Takahashi-Imada path integrals: the physical vibration is shifted to higher values, and the discrepancy decays slowly, as  $1/P^2$ . Even at  $P = 32$  one can observe a significant blue shift and broadening of the OH stretch peak relative to fully converged Trotter TRPMD. Although TRPMD based on fourth-order path integrals eventually converges to the same spectrum as conventional Trotter TRPMD, it does so very slowly, and so there is no advantage in applying fourth-order factorizations to approximate quantum dynamics. The SC scheme could however be used to accelerate the convergence of the mean field centroid force in the case of fully adiabatic centroid molecular dynamics.<sup>25</sup>

## IV. CONCLUSIONS

In the present work we have shown that high-order path integral partition functions can be sampled effi-

ciently by molecular dynamics using a finite-differences evaluation of the second derivatives of the potential energy. This approach makes it possible to use these methods without incurring sampling problems that are associated with statistical reweighting.

We benchmarked our method for the paradigmatic example of liquid water, using a Neural Network potential to model inter-atomic forces at the DFT level, that can also describe extreme anharmonicities in the stretch mode of a H-bonded OH. We show that our finite-differences SC integrator is indeed extremely stable, and that, even in a room temperature regime, it enables computational savings – particularly when combined with a multiple time stepping scheme. An explicit MD integrator for high-order PIMD makes it possible to combine a Suzuki-Chin factorization with colored noise, and approximate quantum dynamics schemes. Unfortunately, we find that a combined SC+GLE approach does not lead to significant improvements over its Trotter (PIGLET) counterpart in this temperature regime, due to the stronger coupling between different ring polymer vibrations in the full SC Hamiltonian, and that a “fourth-order RPMD” scheme has poor performances for dynamical properties. The possibility of performing molecular dynamics with fourth order path integrals, with arbitrary potentials and a moderate computational overhead, provides an additional tool for assessing accurately the impact of quantum fluctuations of nuclei in the condensed phase, and might open the way to new approaches to reduce even further the computational cost of this kind of simulations at and below room temperature.

## ACKNOWLEDGMENTS

We acknowledge financial support by the Swiss National Science Foundation (project ID 200021-159896), and computational time from CSCS under the project IDs s466, s553, and s618. JB is grateful for funding by the DFG cluster of excellence RESOLV (EXC 1069) and for a DFG Heisenberg fellowship (Be3264/6-1). Discussions with Tobias Morawietz are gratefully acknowledged.

### Appendix A: Symplectic behavior of the finite-differences SC integrator

One of the crucial properties that make the scheme we introduce in this work viable is the fact that the velocity-Verlet integrator remains symplectic even when using a finite-differences scheme to compute the SC forces. To illustrate this, let us consider a simplified one-dimensional example that captures the essence of the method. The propagation over a time step  $\Delta t$  of the SC force alone

corresponds to the steps

$$\begin{aligned}\tilde{p} &= p - \Phi(q) \\ q' &= q + \Delta t \tilde{p} \\ p' &= \tilde{p} - \Phi(q'),\end{aligned}\tag{A1}$$

where we introduced the shorthand

$$\Phi(q) = \frac{\Delta t}{2} \left[ c_1 V(q) + c_2 \frac{V(q + \epsilon f(q)) - V(q - \epsilon f(q))}{2\epsilon} \right].\tag{A2}$$

The expression (A1) (as well as the corresponding expression for the forward FD integrator) is clearly time-reversible due to the fact that  $\Phi$  only depends on the instantaneous value of  $q$ . For the same reason, the integrator is symplectic, i.e. the determinant of the Jacobian of the propagator  $|J| = \partial q'/\partial q \partial p'/\partial p - \partial q'/\partial p \partial p'/\partial q$  has a unit value. It is straightforward to see that this is the case, by noting that

$$\begin{aligned}\frac{\partial q'}{\partial q} &= 1 - \Delta t \Phi'(q) \\ \frac{\partial p'}{\partial p} &= 1 - \Delta t \Phi'(q') \\ \frac{\partial q'}{\partial p} &= \Delta t \\ \frac{\partial p'}{\partial q} &= -\Phi'(q) - \Phi'(q') [1 - \Delta t \Phi'(q)]\end{aligned}\tag{A3}$$

### Appendix B: Multiple-time step Suzuki-Chin PIMD

Applying a multiple time step (MTS) procedure to a SC PIMD simulation is not completely trivial. In the presence of a splitting of the physical potential  $V$  into a short-range (cheap)  $V_{\text{sr}}$  and long-range (expensive)  $V_{\text{lr}}$  potential, the forces arising from the two terms get nonlinearly coupled. Differentiating the  $|\mathbf{f}_{\text{sr}} + \mathbf{f}_{\text{lr}}|^2$  term leads to mixed-derivatives containing the short-range Hessian projected on the long-range force and the long-range Hessian projected on the slow force. In order to obtain significant savings, one would have to keep in the outer loop all the terms that contain long-range forces, including those mixed with the short-range Hessian. These terms would fluctuate on a fast time scale, thus reducing the stability range of the outer time step and limiting the achievable computational savings.

There is however another aspect for which a MTS procedure can be beneficial in the context of SC PIMD. The force-dependent term in the SC Hamiltonian is scaled by a pre-factor that becomes smaller as  $P$  is increased. One can then apply a MTS splitting in which the (weighted) Trotter force  $w_j \mathbf{f}^{(j)}$  is applied in the inner loop, and the hard-to-compute  $w_j d_j \tilde{\mathbf{f}}^{(j)}$  term is applied in the outer loop, every  $M$  steps. As shown in Figure 7, this method becomes more and more accurate as SC-PI approaches full convergence – which is the regime in which high-order path integrals give the greatest advantage. Up to

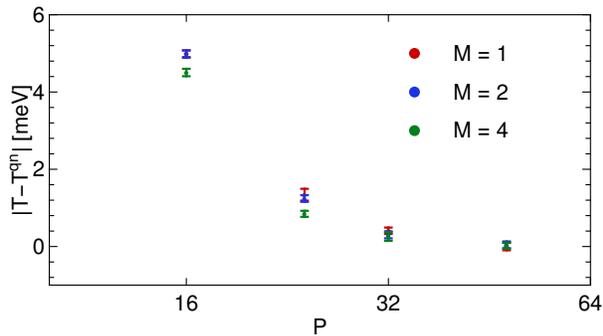


FIG. 7. Convergence of the quantum kinetic energy per molecule for a simulation of liquid water at 300K as a function of  $P$ , using a finite-difference SC integrator and a MTS scheme. The first derivatives of the physical potential have been evaluated every  $\Delta t = 0.25\text{fs}$ , whereas the finite-difference term has been evaluated once  $M$  time steps.

an outer time step of about 1fs (corresponding to an overhead of about 10% relative to Trotter PI with the same number of beads, when using a forward-FD scheme) the inaccuracy due to the MTS splitting is smaller than the residual finite- $P$  convergence error. Pushing  $M$  to even higher values leads to more pronounced errors, without significantly reducing the computational cost.

### Appendix C: Estimation of equilibrium averages

Let us first consider the thermodynamic expression for position-dependent operators, that can be derived using the identity

$$\langle X \rangle = -\frac{1}{\beta Z^{\text{sc}}} \frac{\partial Z^{\text{sc}}(V + \lambda X)}{\partial \lambda} \Big|_{\lambda=1} \quad (\text{C1})$$

which in case of the potential energy gives

$$\mathcal{V}_P^{\text{TD}}(\mathbf{q}) = \frac{1}{P} \sum_{j=0}^{P-1} w_j V(\mathbf{q}^{(j)}) + \frac{2}{P} \sum_{j=0}^{P-1} \sum_{i=0}^{N-1} \frac{w_j d_j}{m_i \omega_P^2} \left| \mathbf{f}_i^{(j)} \right|^2. \quad (\text{C2})$$

Note that this is very similar, but not identical, to the potential term  $\mathcal{V}_P^{\text{sc}}$  entering the SC Hamiltonian. OP-method estimators exploit the fact that the SC factorization can be seen as an imaginary-time propagator that goes over two replicas at a time – so that even beads sample the proper quantum mechanical distribution of  $\mathbf{q}$ . As a consequence, position-dependent OP estimators are very simple, and take the general form

$$\mathcal{X}_P^{\text{OP}}(\mathbf{q}) = \frac{2}{P} \sum_{j \in \text{Even}} X(\mathbf{q}^{(j)}). \quad (\text{C3})$$

For instance, for the potential this yields

$$\mathcal{V}_P^{\text{OP}}(\mathbf{q}) = \frac{2}{P} \sum_{j \in \text{Even}} V(\mathbf{q}^{(j)}). \quad (\text{C4})$$

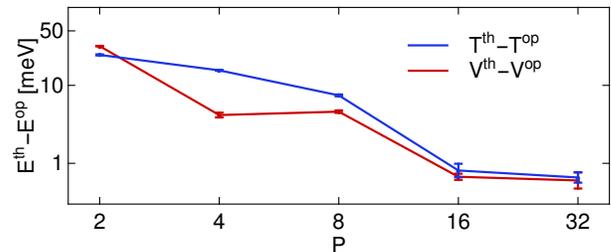


FIG. 8. Difference between the TD and OP fourth-order estimators of the potential and kinetic energy per molecule for a simulation of room-temperature liquid water, as a function of the number of replicas.

Moving on to the case of the kinetic energy estimators, for the TD-method the identity we obtain is

$$\langle \mathcal{T} \rangle = \frac{1}{\beta Z^{\text{sc}}} \sum_{i=1}^N m_i \frac{\partial Z^{\text{sc}}}{\partial m_i}, \quad (\text{C5})$$

which can be written as

$$\begin{aligned} \mathcal{T}_P^{\text{TD}}(\mathbf{q}) &= \frac{3N}{2\beta} + \frac{1}{2P} \sum_{j=0}^{P-1} \sum_{i=0}^{N-1} \left( \mathbf{q}_i^{(j)} - \bar{\mathbf{q}}_i \right) \cdot \mathbf{f}_i^{\text{sc},(j)} \\ &+ \frac{1}{P} \sum_{j=0}^{P-1} \sum_{i=0}^{N-1} \frac{w_j d_j}{m_i \omega_P^2} \left| \mathbf{f}_i^{(j)} \right|^2. \end{aligned} \quad (\text{C6})$$

where

$$\bar{\mathbf{q}}^{(j)} = \frac{1}{P} \sum_{j=0}^{P-1} \mathbf{q}_i^{(j)} \quad (\text{C7})$$

is the centroid of the ring polymer. The second-derivative component within the SC force  $\mathbf{f}_i^{\text{sc},(j)}$  can be computed using the same finite-difference expression (10) that is used to propagate the dynamics. The OP estimator for the kinetic energy is effectively equivalent to the usual centroid virial kinetic energy estimator, evaluated on just the even beads:

$$\mathcal{T}_P^{\text{OP}}(\mathbf{q}) = \frac{3N}{2\beta} + \frac{1}{P} \sum_{j \in \text{Even}} \sum_{i=0}^{N-1} \left( \mathbf{q}_i^{(j)} - \bar{\mathbf{q}}^{(j)} \right) \cdot \mathbf{f}_i^{(j)}, \quad (\text{C8})$$

Note that in general the possibility of computing OP estimators, that do not contain  $\mathbf{f}_i^{\text{sc},(j)}$ , is an important feature of SC path integrals, as that avoids computing second derivatives of the potential in re-weighted schemes. As we have seen, the finite-difference approach we use here makes it trivial to compute the TD estimators, so one can verify which flavor converges more rapidly to the quantum limit and/or cross-validate simulation results. In the case of room-temperature liquid water, we find that there is little difference between the convergence of TD and OP, and that the two estimators become indistinguishable by the time full convergence is achieved (see Figure 8).

## Appendix D: Effective temperature curves for SC+GLE

In the harmonic limit, for a physical potential of frequency  $\omega$ , the frequencies  $\omega_k(\omega)$ , and the eigenvectors  $\mathbf{u}^{(k)}(\omega)$  of the SC Hamiltonian can be obtained by diagonalizing the dynamical matrix  $D_{jj'}$ , that reads

$$D_{jj'} = \begin{cases} 2\omega_P^2 + \omega^2 w_j \left(1 + 2d_j \left(\frac{\omega}{\omega_P}\right)^2\right) & j = j' \\ -\omega_P^2 & j = j' \pm 1 \\ 0 & \text{otherwise} \end{cases}$$

where the cyclic boundary conditions  $j + P \equiv j$  are implied. The dynamical matrix, and as a consequence the derivations that follow, depend on the choice of the parameter  $\alpha$  (see Eq. (8)). We will continue from here on assuming  $\alpha = 0$ , but the generalization to an arbitrary  $\alpha$  is straightforward.

A further complication stems from the fact that the estimator for the fluctuations of  $q$  is not just an average over the coordinates of all beads. In fact, one has to choose whether to design the GLE so as to speed up the convergence of either the TD or the OP estimator for  $\langle q^2 \rangle$ . Given its simplicity, and the direct connection with all structural observables, we opted for the latter choice, that gives

$$\langle q^2 \rangle = \frac{2}{P} \sum_{j \in \text{even}} \left\langle [q^{(j)}]^2 \right\rangle = \frac{1}{P} \sum_k U_k \left\langle [\tilde{q}^{(k)}]^2 \right\rangle \quad (\text{D1})$$

where

$$U_k = 2 \sum_{j \in \text{even}} \left| u_j^{(k)} \right|^2 \quad (\text{D2})$$

gives the weight of the  $k$ -th mode on the displacement of the even beads.

The design of the effective-temperature curve  $T^*(\omega)$  can then proceed in a similar way to what was done in Refs.<sup>19,20</sup>: considering that for a classical oscillator  $\langle q^2 \rangle = 1/m\beta\omega^2$  one can write the functional equation

$$\frac{\hbar}{2\omega k_B} \coth \frac{\hbar\omega\beta}{2} = \frac{1}{P} \sum_k U_k(\omega) \frac{T^*(\omega_k(\omega))}{\omega_k(\omega)^2}. \quad (\text{D3})$$

Here we made explicit the dependence of the eigenvector coefficients and of the normal modes frequencies on the physical frequency of the underlying potential. Solution of Eq. (D3) can be obtained by singling out the lowest-lying NM, obtaining the iteration

$$T^*(\omega_0) = \frac{\omega_0^2}{U_0} \left[ \frac{P\hbar}{2\omega(\omega_0)k_B} \coth \frac{\hbar\omega(\omega_0)\beta}{2} - \sum_{k>0} U_k(\omega_0) \frac{T^*(\omega_k(\omega_0))}{\omega_k(\omega_0)^2} \right] \quad (\text{D4})$$

that can be made to converge with an appropriate mixing scheme<sup>19</sup> and with the starting condition

$$T^*(\omega_0) = \frac{\hbar\omega_0}{\sqrt{6}k_B} \coth \left( \frac{\beta\hbar\omega_0}{\sqrt{6}P} \right) \quad (\text{D5})$$

Yet another complication associated with using a GLE in connection with SC path integrals is that the lowest normal-mode frequency is not equal to the physical frequency, as in Trotter PIMD. For this reason, one needs to invert the  $\omega_0(\omega)$  relation to find what is the physical frequency that corresponds to the argument of  $T^*$  we are solving for. In a similar way, one can then obtain the frequencies of the higher NMs as a function of the lower frequency  $\omega_0$ , which eventually makes it possible to solve numerically the iteration in Eq. D4. In fact, it is possible to give a closed (albeit cumbersome) expression for such inverse relation<sup>58</sup>

$$\omega(\omega_0) = \frac{2^{2/3}A^{2/3} + \sqrt[3]{A}(4x_0^2 - 6P^2) + 4\sqrt[3]{2}x_0^2(3P^2 + 2x_0^2)}{8\beta\hbar\sqrt[3]{A}} \\ A = 27P^6 - 72P^2x_0^4 + 16x_0^6 + 3[81P^{12} - 432P^8x_0^4 + \\ + 384P^4x_0^8 - 384P^2x_0^{10}]^{1/2} \\ x_0 = \beta\hbar\omega_0. \quad (\text{D6})$$

## Appendix E: Accuracy of the neural-network fit

The high-dimensional neural-network potential we use in the present work employs the framework developed by Behler and Parrinello<sup>10,59,60</sup>. Our fit, which follows the NN water potential reported by Morawietz et al. for the case of generalized gradient approximation energetics<sup>30</sup>, has already been shown to provide excellent agreement with both the underlying B3LYP+D3 data it has been fitted to and with experimental quantities, when computing quantum mechanical observables such as the nuclear kinetic energy and isotope fractionation ratios<sup>44</sup>. Here we will just present a few additional diagnostics to demonstrate that it is similarly accurate when it comes to structural and dynamical observables. We used classical MD simulations, for which we could accumulate more than 60 ps of ab initio trajectory (4 independent runs of 15 ps each) with 64 water molecules at constant temperature and volume conditions, so we can perform a statistically meaningful comparison. We applied a weak global thermostat<sup>61</sup>, so as to also be able to extract information on dynamical properties.

As it can be seen in Figure 9, radial distribution functions (RDF) obtained from NN runs are completely indistinguishable from the reference ab initio simulation. It ought to be noted that although pair correlation functions are often used to demonstrate the reliability of an approximate simulation protocol, they constitute a very integrated measure, and success in reproducing RDFs is a necessary but not sufficient criterion to establish the equivalence of two inter-atomic potentials. The lower panels in Figure 9 show two more challenging tests: the histogram of potential energy (that also directly relates to the - classical - heat capacity) and the density of states, as given by the velocity-velocity correlation spectrum.

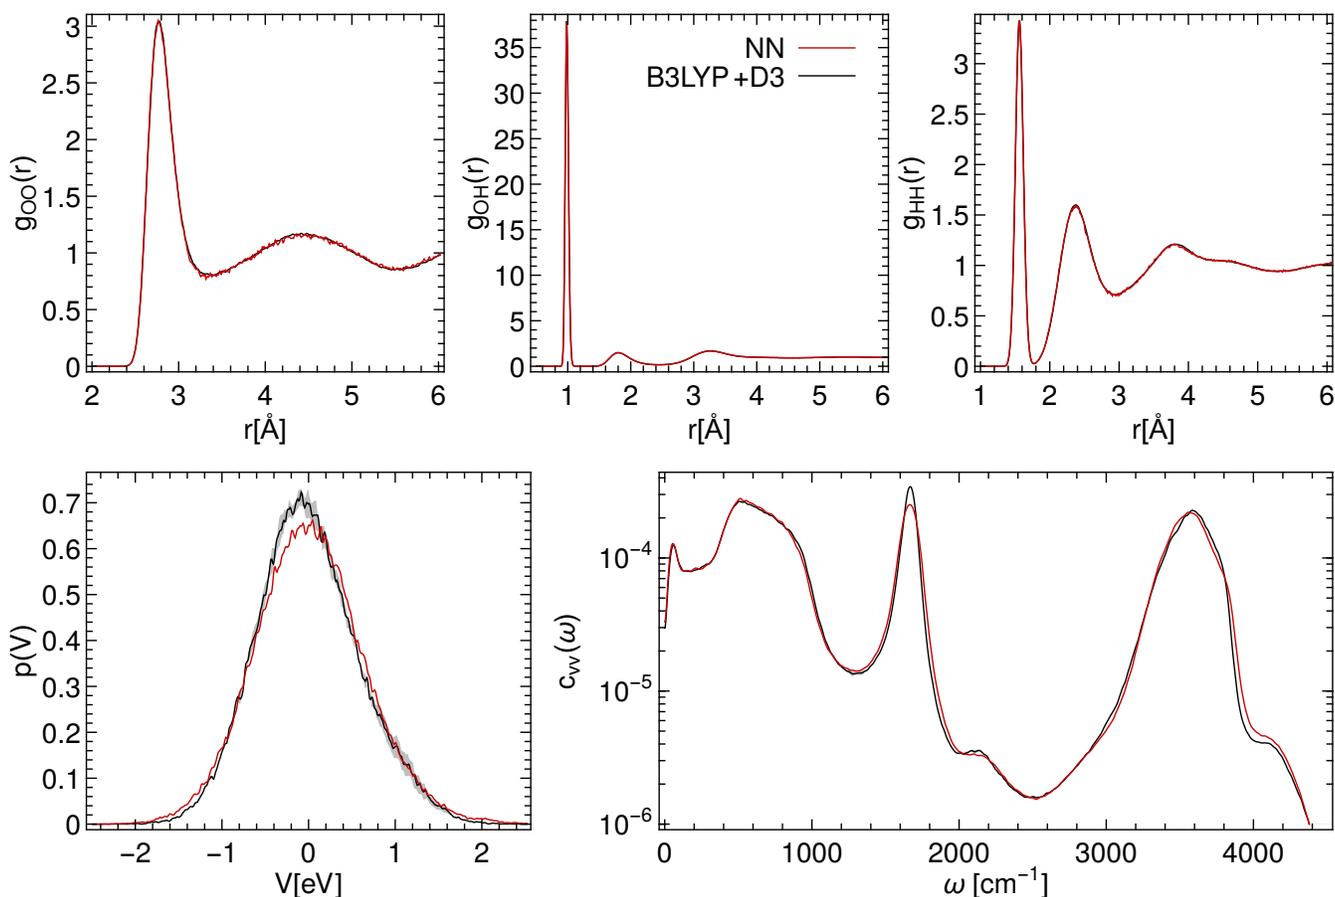


FIG. 9. Comparison between structural and dynamical properties from a *NVT* simulation of 64 water molecules at 300K using B3LYP+D3 (black) and the NN fit (red). The top panels show the pair correlation functions, the bottom-left panel shows the distribution of potential energy fluctuations around the average, and the bottom-right panel shows the vibrational density of states (velocity-velocity correlation function)

Although in these cases one can appreciate differences between the ab initio reference and the NN fit (with the latter showing slightly larger energy fluctuations, and a noticeably broader bend peak), the agreement is excellent. Together with the accurate reproduction of quantum kinetic energy and isotope fractionation ratios, these results give us great confidence in the quality of the NN potential. The possibility of studying larger boxes for longer simulation times outweighs by far the minute discrepancies that are seen for the most stringent tests of Fig. 9.

<sup>1</sup>S. Raugei and M. L. Klein, *J. Chem. Am. Soc.* **125**, 8992 (2003).

<sup>2</sup>M. A. Morales, J. M. McMahon, C. Pierleoni, and D. M. Ceperley, *Phys. Rev. Lett.* **110**, 065702 (2013).

<sup>3</sup>M. Ceriotti, W. Fang, P. G. Kusalik, R. H. McKenzie, A. Michaelides, M. A. Morales, and T. E. Markland, *Chem. Rev.* **116**, 7529 (2016).

<sup>4</sup>P. Debye, *Ann. Phys.* **344**, 789 (1912).

<sup>5</sup>R. A. Berner, S. T. Petsch, J. A. Lake, D. J. Beerling, B. N. Popp, R. S. Lane, E. A. Laws, M. B. Westley, N. Cassar, F. I. Woodward, and W. P. Quick, *Science* **287**, 1630 (2000).

<sup>6</sup>M. A. Webb, Y. Wang, B. J. Braams, J. M. Bowman, and T. F. Miller III, *Geochimica et Cosmochimica Acta* (2016).

<sup>7</sup>C. Andreani, D. Colognesi, J. Mayers, G. F. Reiter, and R. Senesi, *Adv. Phys.* **54**, 377 (2005).

<sup>8</sup>S. Habershon, D. E. Manolopoulos, T. E. Markland, and T. F. Miller, *Annual review of physical chemistry* **64**, 387 (2013).

<sup>9</sup>M. Del Ben, J. Hutter, and J. Vandevonede, *J. Chem. Theory Comput.* **8**, 4177 (2012).

<sup>10</sup>J. Behler and M. Parrinello, *Phys. Rev. Lett.* **98**, 146401 (2007).

<sup>11</sup>A. P. Bartók, M. C. Payne, R. Kondor, and G. Csányi, *Phys. Rev. Lett.* **104**, 136403 (2010).

<sup>12</sup>G. R. Medders, V. Babin, and F. Paesani, *J. Chem. Theory Comput.* **10**, 2906 (2014).

<sup>13</sup>D. M. Ceperley, *Rev. Mod. Phys.* **67**, 279 (1995).

<sup>14</sup>M. Parrinello and A. Rahman, *J. Chem. Phys.* **80**, 860 (1984).

<sup>15</sup>D. Chandler and P. G. Wolynes, *J. Chem. Phys.* **74**, 4078 (1981).

<sup>16</sup>T. E. Markland and D. E. Manolopoulos, *J. Chem. Phys.* **129**, 024105 (2008).

<sup>17</sup>T. E. Markland and D. E. Manolopoulos, *Chem. Phys. Lett.* **464**, 256 (2008).

<sup>18</sup>M. Ceriotti, G. Bussi, and M. Parrinello, *Phys. Rev. Lett.* **103**, 30603 (2009).

<sup>19</sup>M. Ceriotti, D. E. Manolopoulos, and M. Parrinello, *J. Chem. Phys.* **134**, 84104 (2011).

<sup>20</sup>M. Ceriotti and D. E. Manolopoulos, *Phys. Rev. Lett.* **109**, 100604 (2012).

<sup>21</sup>M. Takahashi and M. Imada, *J. Phys. Soc. Jap.* **53**, 3765 (1984).

<sup>22</sup>M. Suzuki, *Phys. Lett. A* **201**, 425 (1995).

- <sup>23</sup>S. A. Chin, Phys. Lett. A **226**, 344 (1997).
- <sup>24</sup>Y. Kamibayashi and S. Miura, The Journal of Chemical Physics **145**, 074114 (2016).
- <sup>25</sup>S. S. Jang and G. A. Voth, J. Chem. Phys. **115**, 7832 (2001).
- <sup>26</sup>T. M. Yamamoto, J. Chem. Phys. **123**, 104101 (2005).
- <sup>27</sup>O. Marsalek, P.-Y. Chen, R. Dupuis, M. Benoit, M. Méheut, Z. Bačić, and M. E. Tuckerman, J. Chem. Theory Comput. **10**, 1440 (2014).
- <sup>28</sup>M. Ceriotti, G. a. R. Brain, O. Riordan, and D. E. Manolopoulos, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences **468**, 2 (2011).
- <sup>29</sup>I. Poltavsky and A. Tkatchenko, Chem. Sci. , (2016).
- <sup>30</sup>T. Morawietz, A. Singraber, C. Dellago, and J. Behler, Proc. Nat. Acad. Sci. **113**, 8368 (2016).
- <sup>31</sup>A. Pérez and M. E. Tuckerman, J. Chem. Phys. **135**, 064104 (2011).
- <sup>32</sup>C. Predescu, Phys. Rev. E **70**, 066705 (2004).
- <sup>33</sup>M. Buchowiecki, J. Vaníček, and J. Vaníček, Chem. Phys. Lett. **588**, 11 (2013).
- <sup>34</sup>A. Putrino, D. Sebastiani, and M. Parrinello, J. Chem. Phys. **113**, 7102 (2000).
- <sup>35</sup>In this work we could afford parallelizing calculations over all of the beads. As a consequence, there would be no advantage in leaving half of the processors idle during the evaluation of  $\mathbf{f}^{(j)}$ . Therefore, we used throughout the symmetric FD expression.
- <sup>36</sup>B. Cheng and M. Ceriotti, J. Chem. Phys. **141**, 244112 (2014).
- <sup>37</sup>M. Ceriotti, J. More, and D. E. Manolopoulos, Comp. Phys. Comm. **185**, 1019 (2014).
- <sup>38</sup>V. Kapil, J. VandeVondele, and M. Ceriotti, J. Chem. Phys. **144**, 054111 (2016).
- <sup>39</sup>M. Ceriotti, G. Bussi, and M. Parrinello, J. Chem. Theory Comput. **6**, 1170 (2010).
- <sup>40</sup>M. Ceriotti, “GLE4MD,” <http://epfl-cosmo.github.io/gle4md> (2010).
- <sup>41</sup>A. D. Becke, J. Chem. Phys. **98**, 5648 (1993).
- <sup>42</sup>S. Grimme, J. Antony, S. Ehrlich, and H. Krieg, J. Chem. Phys. **132**, 154104 (2010).
- <sup>43</sup>J. VandeVondele, M. Krack, F. Mohamed, M. Parrinello, T. Chassaing, and J. Hutter, Comp. Phys. Comm. **167**, 103 (2005).
- <sup>44</sup>B. Cheng, J. Behler, and M. Ceriotti, J. Phys. Chem. Letters **7**, 2210 (2016).
- <sup>45</sup>A. Singraber, T. Morawietz, J. Behler, and C. Dellago, to be published.
- <sup>46</sup>S. Plimpton, J. Comp. Phys. **117**, 1 (1995).
- <sup>47</sup>M. Ceriotti, M. Parrinello, T. E. Markland, and D. E. Manolopoulos, J. Chem. Phys. **133**, 124104 (2010).
- <sup>48</sup>M. Rossi, M. Ceriotti, and D. E. Manolopoulos, J. Chem. Phys. **140**, 234116 (2014).
- <sup>49</sup>M. Ceriotti, J. Cuny, M. Parrinello, and D. E. Manolopoulos, Proc. Natl. Acad. Sci. USA **110**, 15591 (2013).
- <sup>50</sup>F. Uhl, D. Marx, and M. Ceriotti, J. Chem. Phys. **145**, 054101 (2016).
- <sup>51</sup>S. Ganeshan, R. Ramírez, and M. V. Fernández-Serra, Phys. Rev. B **87**, 134207 (2013).
- <sup>52</sup>L. Wang, M. Ceriotti, and T. E. Markland, J. Chem. Phys. **141**, 104502 (2014).
- <sup>53</sup>J. Cao and G. A. Voth, J. Chem. Phys. **99**, 10070 (1993).
- <sup>54</sup>J. Cao and G. A. Voth, J. Chem. Phys. **101**, 6168 (1994).
- <sup>55</sup>I. R. Craig and D. E. Manolopoulos, J. Chem. Phys. **121**, 3368 (2004).
- <sup>56</sup>A. Witt, S. D. Ivanov, M. Shiga, H. Forbert, and D. Marx, J. Chem. Phys. **130**, 194510 (2009).
- <sup>57</sup>M. Rossi, H. Liu, F. Paesani, J. Bowman, and M. Ceriotti, J. Chem. Phys. **141**, 181101 (2014).
- <sup>58</sup>G. Brain, *Higher Order Propagators in Path Integral Molecular Dynamics*, Ph.D. thesis, Part II Chemistry Thesis, Oxford University (2011).
- <sup>59</sup>J. Behler, J. Chem. Phys. **134**, 074106 (2011).
- <sup>60</sup>J. Behler, J. Phys.: Condens. Matter **26**, 183001 (2014).
- <sup>61</sup>G. Bussi, D. Donadio, and M. Parrinello, J. Chem. Phys. **126**, 14101 (2007).