Modeling the Quantum Nature of Atomic Nuclei by Imaginary Time Path Integrals and Colored Noise

Michele Ceriotti

Laboratory of Computational Science and Modeling, IMX,
École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland

E-mail: michele.ceriotti@epfl.ch

The vast majority of atomistic simulations of matter treat the nuclei as classical point particles, evolving in time according to Hamiltonian dynamics, and following classical Boltzmann statistics. As a matter of fact, this is quite a harsh approximation when dealing with hydrogen-containing compounds. Hydrogen nuclei – as well as other light elements – exhibit significant deviations from classical behavior up to and above room temperature. Here I will present an overview of path integral methods, that can be used to include nuclear quantum effects in atomic-scale models. I will cover the basic theory, discuss some of the practicalities in the implementation and use, and finally introduce the most recent advances towards making these techniques less computationally demanding by using colored (correlated) stochastic dynamics.

Simulations that describe matter at the level of individual atoms offer a very high level of accuracy, transferability and predictive power. In order to make them practically feasible, a number of approximations are often introduced that trade off the accuracy in describing some physical effects in exchange for a reduced complexity and computational cost. Perhaps the most widely adopted approximation is the decoupling of the electronic structure problem from that of the statistical and dynamical behavior of the atomic nuclei. This can take the form of Born-Oppenheimer approximation[1] – where the ground-state electronic structure problem is solved for a given configuration of the nuclei – or can be realized by modeling the interaction between the atoms using an empirical force field that represent effectively (and inexpensively) the potential energy surface for the atoms treated as point particles.

The Born-Oppenheimer approximation is generally very satisfactory, except when the system evolves in an electronic excited state or for a few ultra-fast chemical reactions. However, it only consists in a factorization of the combined electronic-nuclear wavefunction, and in principle the nuclei should be treated as quantum particles. The vast majority of atomistic simulations are performed with an additional approximation, that is to treat the nuclei as classical particles that evolve in time following Hamilton's equations and that are subject to Boltzmann, classical statistics. These are certainly reasonable approximations at high temperature, and when dealing with heavy nuclei. If however one compares the thermal energy $k_B T$ and the quantum of harmonic energy $\hbar\omega$ for a molecular vibration of frequency ω at temperature T, it will become clear that for many compounds $\hbar\omega/k_B T \gg 1$ even well above room temperature, which casts some shadows on the consequences of neglecting the quantum nature of the nuclear degrees of freedom in simulations.

There are several examples of the impact of the quantum mechanical behavior of nuclei on experimental observables. The heat capacity of substances deviates from the Dulong-Petit prediction of $3k_B T$ per atom (that corresponds to classical statistics for a harmonic crystal), in particular for stiff bonds (as in diamond) or for hydrogen-containing com-

pounds. The kinetic energy distribution of atomic nuclei, as measured by neutron Compton scattering, differs dramatically from the Maxwell-Boltzmann distribution[2]. Reaction rates at low temperature do not follow an Arrhenius behavior. The stability of different compounds or phases varies with isotope composition, and one can for instance estimate (extrapolating the values measured for ^1H_2O, ^2H_2O, ^3H_2O) that the pH of water would be around 8.5 if nuclei behaved classically. Some of these phenomena – isotope substitution effects in particular – simply cannot be observed in the absence of a quantum mechanical treatment of the nuclear degrees of freedom, while others entail a deviation of computed properties from their experimental counterparts. Neglecting nuclear quantum effects (NQEs) is particularly detrimental in the case of *ab initio* molecular dynamics, where the ground-state electronic structure problem is solved on the fly, and the nuclei evolve on the bare Born-Oppenhemier potential energy surface. Simulations employing empirical force fields can include NQEs indirectly, by fitting the parameters of the inter-atomic potential to experimental observables, or more directly using approximate techniques such as Feynman-Hibbs effective potentials[3].

Solving the Schrödinger equation for the nuclei is impractical except for very simple systems. Here we will discuss how the imaginary-time path integral formalism[4–7] can be used to evaluate accurately NQEs in complex condensed-phase applications[8–10]. We will focus on static, equilibrium properties, but will briefly mention extensions to the path integral formalism that can be used to treat approximately quantum dynamics[11,12]. We will focus on the case in which different nuclei can be treated as distinguishable particles, which is often true except for cases at cryogenic temperatures. Particle exchange statistics can be included within a path integral formalism, but at the cost of considerable complication and an increase of the computational cost[5,13].

1 Imaginary-Time Path Integrals

The path integral formulation of quantum mechanics makes it possible to express all the quantities that describe a physical system in terms of exponential averages of an appropriate action integral over the possible paths joining two points in phase space – much like the minimum action principle makes it possible to formulate classical mechanics as the minimization of the action over a tentative path[4]. Furthermore, it makes it possible to express the quantum mechanical partition function at inverse temperature $\beta = 1/k_B T$ [a]

$$Z = \mathrm{Tr}\, e^{-\beta \hat{H}}$$

as the *path integral*

$$Z = \oint \mathcal{D}\left[q\left(\tau\right)\right] e^{-\frac{1}{\hbar} \int_0^{\beta\hbar} \left[\frac{1}{2} m \dot{q}(\tau)^2 + V(q(\tau))\right] d\tau}. \tag{1}$$

The symbol $\oint \mathcal{D}\left[q\left(\tau\right)\right] \cdot$ is a functional integral over all the possible *closed* paths in configuration space, weighed with the exponential of an action-like integral over the path. In the following we will discuss how to give a practical definition of Eq. (1), and how to use this formalism to compute experimental observables including NQEs.

---

[a]We consider for simplicity the case of a single particle with position q and mass m in an external potential V

1.1 Trotter Factorization

Start by writing the partition function in the position representation

$$Z = \int dq_1 \left\langle q_1 \left| e^{-\beta \hat{H}} \right| q_1 \right\rangle .$$

The Hamiltonian can be written as the sum of a potential and kinetic energy terms, $\hat{H} = \hat{V} + \hat{T}$, and the position ket is an eigenstate of the potential energy, so that $e^{-\beta \hat{V}} |q_1\rangle = e^{-\beta V(q_1)} |q_1\rangle$. Unfortunately, one cannot factor $e^{-\beta \hat{H}}$ into the product $e^{-\beta \hat{V}} e^{-\beta \hat{T}}$, because potential and kinetic energy are not commuting operators. However, the error in doing such a factorization decreases when $\beta \hat{H}$ becomes small. So, one could write[b]

$$e^{-\beta \hat{H}} = \left(e^{-\beta \hat{H}/P} \right)^P \approx \left(e^{-\beta_P \hat{V}/2} e^{-\beta_P \hat{T}} e^{-\beta_P \hat{V}/2} \right)^P + \mathcal{O}\left(\beta_P^2 \right),$$

which becomes exact in the $P \to \infty$ limit. Note that we have also introduced the shorthand $\beta_P = \beta/P$. One can show that the partition function converges to the exact quantum mechanical result with a leading error of $\mathcal{O}\left(\beta^2/P^2 \right)$, and in practice for a system with a maximum frequency ω_{max} one needs a number of imaginary time slices that is at least a small multiple of $\beta \hbar \omega_{max}$.

One can then introduce $P - 1$ closure relations $\int dq_j |q_j\rangle \langle q_j|$, obtaining

$$\begin{aligned} Z \approx Z_P = \int dq_1 \ldots dq_P &\left[\left\langle q_1 \left| e^{-\beta_P V(q_1)/2} e^{-\beta_P \hat{T}} e^{-\beta_P V(q_2)/2} \right| q_2 \right\rangle \ldots \right. \\ &\left. \ldots \left\langle q_P \left| e^{-\beta_P V(q_P)/2} e^{-\beta_P \hat{T}} e^{-\beta_P V(q_1)/2} \right| q_1 \right\rangle \right] . \end{aligned} \tag{2}$$

The terms with the potential energy are just scalar values, that can be brought outside the quantum mechanical brackets. One is then left with a series of terms corresponding to the off-diagonal elements of the kinetic energy operator, that are readily evaluated by transforming in the momentum representation:

$$\begin{aligned} \left\langle q_i \left| e^{-\beta_P \hat{T}} \right| q_j \right\rangle &= \int dp \left\langle q_i \left| e^{-\beta_P \hat{T}} \right| p \right\rangle \langle p | q_j \rangle = \\ &= \frac{1}{2\pi\hbar} \int dp\, e^{-\beta_P p^2/2m} e^{ip(q_i - q_j)/\hbar} = \\ &= \frac{1}{2\pi\hbar} \sqrt{\frac{2\pi m}{\beta_P}} e^{-\frac{1}{2}\beta_P m \omega_P^2 (q_i - q_j)^2} \end{aligned} \tag{3}$$

where we have used $\langle p|q \rangle = e^{-ipq/\hbar}/\sqrt{2\pi\hbar}$, performed the integral over the momentum and introduced the spring constant $\omega_P = 1/\beta_P \hbar$. Plugging Eq. (3) into Eq. (2) one finally obtains the path integral configuration partition function

$$Z_P = \left(\frac{m}{2\pi\hbar^2 \beta_P} \right)^{P/2} \int dq_1 \ldots dq_P e^{-\beta_P \sum_{i=1}^{P} \left[V(q_i) + \frac{1}{2} m \omega_P^2 (q_i - q_{i+1})^2 \right]} \tag{4}$$

where cyclic boundary conditions are implied in the sum, $i + P \equiv i$.

---

[b]Note that we use the Trotter splitting $e^{A+B} \approx e^{A/2} e^B e^{A/2}$, which has a lower error than the asymmetric splitting $e^A e^B$.

Ignoring for a second the configuration-independent pre-factor, let us discuss the connection between Eq. (4) and the Feynmann path integral (1). Consider q_i to be a discrete sample from a continuous path, taken at $\tau_i = \beta\hbar i/P$. Then one can see the sum in the exponential as a discretization of a Riemann integral, and $(q_{i+1} - q_i)/(\tau_{i+1} - \tau_i)$ as a finite-difference approximation to $\dot{q}(\tau_i)$

$$\frac{\beta}{P} \sum_{i=1}^{P} \frac{(\tau_{i+1} - \tau_i)}{\beta\hbar/P} \left[V(q_i) + \frac{1}{2}m \frac{(q_i - q_{i+1})^2}{(\tau_{i+1} - \tau_i)^2} \right] \approx \frac{1}{\hbar} \int_0^{\beta\hbar} d\tau \left[V(q(\tau)) + \frac{1}{2}m\dot{q}(\tau)^2 \right].$$

The multiple integrals over the q_j coordinates are the discrete equivalent of the path integral $\oint \mathcal{D}[q(\tau)]$. Even though most path integral simulations can be implemented and understood without reference to the formulation in terms of a functional integral, it is useful to keep this limit in mind, particularly when writing estimators for physical observables that are usually better behaved when they can be expressed as the discretized version of a corresponding path integral[14].
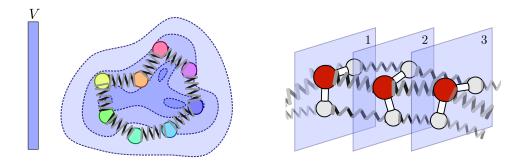

Figure 1: (Left panel) Cartoon representation of a classical ring polymer corresponding to the discretized path integral partition function (4). (Right panel) In a multi-atom setting, the ring polymer metaphor can be somewhat misleading. The path integral partition function is best seen as a sequence of imaginary-time slices: atoms within each replica interact with the physical potential, and the spring terms connect corresponding atoms in adjacent slices.

Eq. (4) corresponds precisely to the *classical* partition function of a cyclic polymer composed of P atoms, each of which is subject to the potential V and of a harmonic attractive interaction with its nearest neighbors. This isomorphism motivates the common practice of referring to the set of replicas for one atom as a "ring polymer" or a "necklace" and to each replica as a "bead". While this is a very suggestive metaphor, and can also be extended to give a pictorial representation of the Monte Carlo moves that are introduced to treat particle exchange effects[5], one has to keep in mind that in a real system it is better to regard the path integral partition function as describing a collection of "parallel universes", with atoms interacting with each other within each imaginary time slice and the kinetic part of the action corresponding to springs that connect each atom to its counterpart in the two adjacent time slices (see Figure (1)).

1.2 Estimators

Having defined an isomorphism between the quantum partition function of a system of distinguishable nuclei and the ring-polymer configurational partition function (4), one may proceed to evaluate experimental observables. To do so, one has to introduce appropriate *estimators*, functions of the coordinates of the ring polymer that correspond to physical quantities. The simplest case is that of the potential energy, or of any observable $\hat{A}(q)$ that depends solely on the atomic positions, such as a bond length, a radial distribution function, or the relative stability of two molecular configurations. The quantum mechanical expectation value can be written as $\langle A \rangle = \mathrm{Tr}\left[\hat{A}e^{-\beta\hat{H}}\right] / \mathrm{Tr}\, e^{-\beta\hat{H}}$. The operator $\hat{A}$ can be kept on the left, so that after the Trotter factorization and the splitting of the integral it appears close to the $\langle q_1 |$ in Eq. (2), yielding a term $A(q_1)$. Since all the replicas are equivalent, one may as well average over the value of the observable computed on all the beads, so that the expectation value reads

$$\langle A \rangle_P = \frac{\int \mathrm{d}q_1 \ldots \mathrm{d}q_P e^{-\beta_P \sum_{i=1}^{P}\left[V(q_i)+\frac{1}{2}m\omega_P^2(q_i-q_{i+1})^2\right]} \frac{1}{P}\sum_{i=1}^{P} A(q_i)}{\int \mathrm{d}q_1 \ldots \mathrm{d}q_P e^{-\beta_P \sum_{i=1}^{P}\left[V(q_i)+\frac{1}{2}m\omega_P^2(q_i-q_{i+1})^2\right]}}. \tag{5}$$

This average can be computed easily by sampling the ring polymer configurations consistently with the ring polymer energy $\sum_{i=1}^{P}\left[V(q_i)+\frac{1}{2}m\omega_P^2(q_i-q_{i+1})^2\right]$ at the inverse temperature β_P, with Monte Carlo or (as it will be discussed in Section 2) molecular dynamics, and accumulating statistics for each replica. Note that this does *not* necessarily mean that averages will converge faster then if one was sampling just one replica, as in classical sampling: different path integral replicas are typically highly correlated with one another, which means that in most cases very little is gained by the average in (5).

While it is simple to write an estimator for observables that depend only on the positions, it is generally more complex to extract momentum-dependent quantities. A good example is that of the total energy, that contains both the position-dependent potential energy, but also a kinetic energy term. The simplest form of an estimator for the total energy of the system can be obtained recalling the thermodynamic relation between the partition function and the mean energy $\langle E \rangle = -Z^{-1}\partial Z/\partial\beta$. Once applied to Eq. (2), this reads

$$\langle E \rangle = \frac{\int \mathrm{d}q_1 \ldots \mathrm{d}q_P e^{-\beta_P \sum_{i=1}^{P}\left[V(q_i)+\frac{1}{2}m\omega_P^2(q_i-q_{i+1})^2\right]} E^{\mathrm{TD}}(q_1,\ldots q_P)}{\int \mathrm{d}q_1 \ldots \mathrm{d}q_P e^{-\beta_P \sum_{i=1}^{P}\left[V(q_i)+\frac{1}{2}m\omega_P^2(q_i-q_{i+1})^2\right]}},$$

the ensemble average of the so-called *thermodynamic* (or *primitive*) energy estimator[c]

$$E^{\mathrm{TD}}(q_1,\ldots q_P) = \frac{P}{2\beta} - \frac{1}{P}\sum_{i=1}^{P}\frac{1}{2}m\omega_P^2(q_i-q_{i+1})^2 + \frac{1}{P}\sum_{i=1}^{P}V(q_i).$$

One can immediately recognize the estimator for the potential energy, and infer that the thermodynamic estimator for the kinetic energy alone reads

$$T^{\mathrm{TD}}(q_1,\ldots q_P) = \frac{P}{2\beta} - \frac{1}{P}\sum_{i=1}^{P}\frac{1}{2}m\omega_P^2(q_i-q_{i+1})^2. \tag{6}$$

---

[c]When doing the derivation keep in mind that $\omega_P = 1/\beta_P\hbar$ and that the constant scaling of Z_P also depends on β

Note that this estimator does not depend solely on the distribution of individual beads, but also on the cross-correlations between different replicas in the ring polymer. This is a general feature for non-local estimators that also contain a kinetic energy (or momentum) contribution.

Unfortunately, the thermodynamic kinetic energy estimator (6) is not very efficient, because its variance grows with the number of beads[15] as P/β^2. This means that computing the average to a given accuracy becomes more difficult as the number of replicas is increased, making the simulation even more computationally demanding.

Luckily, it is possible to exploit the virial theorem and do an integration by parts to derive the so-called centroid-virial kinetic energy estimator

$$T^{\mathrm{CV}}\left(q_1, \ldots q_P\right) = \frac{1}{2\beta} + \frac{1}{2P} \sum_{i=1}^{P} (q_i - \bar{q}) \frac{\partial V}{\partial q_i}, \qquad \bar{q} = \frac{1}{P} \sum_{i=1}^{P} q_i, \tag{7}$$

that does not exhibit this pathological behavior of the fluctuations[15]. This case is a typical example of a recurring theme in path integral methods: one can write estimators that yield the same average value, but have very different statistical convergence properties. See for instance Ref.[16] for a discussion of efficient estimators for the heat capacity, Ref.[14] for an estimator of the distribution of particle momenta, and Refs.[17,18] for a comparison of different estimators for the isotope fractionation ratio – the relative propensity of the isotopes of the same element for different stable phases.

1.3 High-order Path Integrals

We have seen that Eq. (2) can be interpreted as a discretized form of a line integral over closed paths in configuration space. The error arising from using a finite P is effectively a discretization error, so one might wonder if it is possible to increase the order of convergence by employing a different summation rule. The crux is the error arising from splitting the exponential of the Hamiltonian, neglecting the commutator $\left[\hat{T}, \hat{V}\right]$. Hence, one can hope to increase the order of convergence by including extra terms that also depend on this commutator. A considerable amount of research has been devoted to this topic, in part also because of a connection with algorithms to propagate Hamiltonian dynamics in real time[19–22,16,23]. Among the many factorizations that have been proposed, one of the simplest and most successful is the Suzuki-Chin propagator

$$e^{-2\beta\hat{H}/P} = e^{-\frac{1}{3}\beta_P \hat{V}_e} e^{-\beta_P \hat{T}} e^{-\frac{2}{3}\beta_P \hat{V}_o} e^{-\beta_P \hat{T}} e^{-\frac{1}{3}\beta_P \hat{V}_e} + \mathcal{O}\left(\beta_P^5\right), \tag{8}$$

in which one introduces two distinct modified potential energy operators that act on the odd and the even beads:

$$\hat{V}_e = \hat{V} + \frac{\alpha}{6m\omega_P^2} \left|\frac{\partial V}{\partial q}\right|^2, \qquad \hat{V}_o = \hat{V} + \frac{1-\alpha}{12m\omega_P^2} \left|\frac{\partial V}{\partial q}\right|^2.$$

The parameter $\alpha \in [0, 1]$ can be tuned to optimize the prefactor for the discretization error. The square modulus of the force can typically be computed without additional effort, in particular when using molecular dynamics to sample the ring polymer partition function, as we will discuss in Section 2. However, the factorization (8) becomes inconvenient precisely when one wants to integrate the dynamics that arises from $\hat{V}_{e,o}$: the derivative of

$|\partial V/\partial q|^2$ contains the second derivative of the potential, which often is considerably more computationally demanding than the force. In these cases, one can get around this difficulty by sampling the Trotter partition function, and using statistical reweighing to recover sampling consistent with the Suzuki-Chin modified potential[16, 24–26]. This technique works very well for small clusters; however, the difference between the Trotter and the Suzuki-Chin potential energy is size-extensive, which means that it becomes progressively less efficient to include reweighing factors in the averages as the number of atoms is increased, somewhat limiting the applicability of this approach[25].

2   Path Integral Molecular Dynamics

As it has been already discussed above, the ring polymer partition function (2) only needs to depend explicitly on coordinates, and it can be sampled by Monte Carlo techniques. While Monte Carlo moves can be very effective in sampling phase space – and for instance they are used extensively in techniques to include particle exchange statistics[27, 13] – they typically do not exploit the possibility of obtaining the inter-atomic forces with little overhead over the calculation of the potential energies. In many cases, particularly when one does not want to develop custom-tailored Monte Carlo moves for the specific system at hand, it can be much more effective to sample the Boltzmann distribution by integrating in time Hamilton's equations

$$\dot{q} = \frac{\partial H}{\partial p} = \frac{p}{m}, \qquad \dot{p} = -\frac{\partial H}{\partial q} = -\frac{\partial V}{\partial q}. \tag{9}$$

In practice, it is easy to see that the constant pre-factor in Eq. (2) corresponds to a Gaussian integral over a set of auxiliary variables that can be taken to be the conjugate momenta to q_i's, so that one can equivalently write the ring polymer partition function as

$$Z_P = \frac{1}{(2\pi\hbar)^P} \int \mathrm{d}p_1 \ldots \mathrm{d}p_P \int \mathrm{d}q_1 \ldots \mathrm{d}q_P e^{-\beta_P \sum_{i=1}^{P} \left[V(q_i) + \frac{p_i^2}{2m} + \frac{1}{2}m\omega_P^2 (q_i - q_{i+1})^2 \right]}. \tag{10}$$

Note that the momenta p_i are exclusively sampling devices, *and are in no way related to the physical momentum*. Among other things, this means that one could change the inertial mass in the $p_i^2/2m$ term to be something different from the physical mass, with no other effect than changing the partition function by an immaterial, temperature-independent scaling.

2.1   Implementation: Normal Modes Propagator

While the underlying idea behind path integral molecular dynamics (PIMD) is very simple, there are some technical aspects that should be considered in order to obtain an efficient implementation. Start by splitting the path integral Hamiltonian into a "free-particle" component, and one that depends on the physical potential:

$$H_P = H_0 + V_P = \sum_{i=1}^{P} \left[\frac{p_i^2}{2m} + \frac{1}{2}m\omega_P^2 (q_i - q_{i+1})^2 \right] + \sum_{i=1}^{P} V(q_i).$$

7

The free ring-polymer Hamiltonian H_0 is just a multi-dimensional harmonic oscillator, and therefore it can be diagonalized exactly by the unitary transformation

$$\tilde{p}_j = \sum_i p_i C_{ij}, \quad \tilde{q}_j = \sum_i q_i C_{ij}, \quad C_{ij} = \sqrt{\frac{2}{P}} \cdot \begin{cases} 1/\sqrt{2} & j = 0 \\ \cos 2\pi ij/P & j < P/2 \\ (-1)^i/\sqrt{2} & j = P/2 \\ \sin 2\pi ij/P & j > P/2 \end{cases}.$$

$\tilde{q}_j$ and $\tilde{p}_j$ are the position and momentum of the j-th free-particle normal mode, which has an associated frequency

$$\omega_j = 2\omega_P \sin j\pi/P. \tag{11}$$

Considering that for a physical system that contains a normal mode of frequency $\omega_{\max}$ one has to use at least $2\beta\hbar\omega_{\max}$ replicas, a converged PIMD calculation will involve normal-mode vibrational frequencies larger than $4\omega_{\max}$. This implies that in principle one should use a much smaller time step to integrate the equations of motion for PIMD compared to conventional classical molecular dynamics, adding further computational burden to the technique.

Luckily, one can exploit the possibility of diagonalizing the free ring-polymer Hamiltonian to avoid reducing the time step. This can be achieved by doing a so-called staging transformation[28], which will not be discussed here, or by performing the integration in the normal modes basis. In the latter case, one can for instance manipulate the inertial mass of the conjugate momenta in the normal modes representation, so as to artificially slow down the dynamics of the fast ring-polymer normal modes without changing the sampling properties. In most cases, however, this is not necessary: one can perform a multiple time step procedure[29] to integrate the normal modes analytically based on the evolution of a free ring polymer and then include the physical potential using a velocity Verlet algorithm[30].

This is based on the symmetric Trotter splitting of the Liouville operator for H_P into the part related to H_0 and that related to the physical potential $\mathcal{L} = \mathcal{L}_0 + \mathcal{L}_V$. Considering the evolution of the ring polymer over a time step Δt

$$e^{-\Delta t \mathcal{L}} \approx e^{-\Delta t \mathcal{L}_V/2} e^{-\Delta t \mathcal{L}_0} e^{-\Delta t \mathcal{L}_V/2}.$$

This corresponds to the following recipe for the evolution across a time step:

$$\begin{aligned} p_i &\leftarrow p_i - \frac{\Delta t}{2}\frac{\partial V(q_i)}{\partial q_i} \\ \tilde{p}_j &\leftarrow \sum_i p_i C_{ij} \qquad \tilde{q}_j \leftarrow \sum_i q_i C_{ij} \\ \begin{pmatrix} \tilde{p}_j \\ \tilde{q}_j \end{pmatrix} &\leftarrow \begin{pmatrix} \cos\omega_j\Delta t & -m\omega_j\sin\omega_j\Delta t \\ [1/m\omega_j]\sin\omega_j\Delta t & \cos\omega_j\Delta t \end{pmatrix} \begin{pmatrix} \tilde{p}_j \\ \tilde{q}_j \end{pmatrix} \\ p_i &\leftarrow \sum_j C_{ij}\tilde{p}_j \qquad q_i \leftarrow \sum_j C_{ij}\tilde{q}_j \\ p_i &\leftarrow p_i - \frac{\Delta t}{2}\frac{\partial V(q_i)}{\partial q_i}. \end{aligned} \tag{12}$$

In practice, one could reduce the number of normal modes transformations (that can be performed by fast Fourier transform, and therefore are not dramatically demanding anyway) by keeping the momenta in the normal modes representation, and transforming the physical forces in the normal modes basis. See for instance Ref.[31] for a discussion of Eq. (12) in a many-atoms context.

It is also worth mentioning that whenever the inter-atomic potential can be split in a part that varies on a short length scale (such as intra-molecular bends and stretches) and a long-range scale (such as electrostatics), the normal-modes representation can also be exploited within a *ring-polymer contraction* scheme[32,33], that reduces the cost of the simulation by computing the long-range part of the potential on a reduced number of replicas.

2.2 Efficient Stochastic Thermostatting

Molecular dynamics generates trajectories that are consistent with a Boltzmann distribution. However, they do not yield ergodic sampling, because they conserve the total energy and so a single MD simulation cannot account for the thermal fluctuations that are characteristic of a constant-temperature ensemble, and that are needed to converge averages based on the partition function (10). This fact has been recognized for a long time, and led to the development of modified dynamical equations that describe the heat exchange with a reservoir, and generate ergodic canonical trajectories[34–40]. In principle, any of these techniques could be applied to PIMD, which is just classical molecular dynamics in an extended phase space. Perhaps, the simplest technique to be used in conjunction with the normal-modes propagation in Eq. (12) is one that applies Langevin dynamics in the normal modes representation, so that the friction can be tuned to match the critical-damping value for the free ring-polymer normal modes. We introduce a Langevin term in the equation of motion for the momenta, corresponding to a further term $\mathcal{L}_\gamma$ in the Liouvillian, that can therefore be split as

$$e^{-\Delta t \mathcal{L}} \approx e^{-\Delta t \mathcal{L}_\gamma/2} e^{-\Delta t \mathcal{L}_V/2} e^{-\Delta t \mathcal{L}_0} e^{-\Delta t \mathcal{L}_V/2} e^{-\Delta t \mathcal{L}_\gamma/2}.$$

This splitting corresponds to applying the following step

$$\dot{\tilde{p}}_j \leftarrow e^{-\gamma_j \Delta t/2} \tilde{p}_j + \sqrt{(1 - e^{-\gamma_j \Delta t}) m/\beta_n} \xi_j$$

twice per time step, immediately before and immediately after the Hamiltonian propagator (12). ξ_j's are uncorrelated Gaussian random numbers, such that $\langle \xi_j(t)\xi_{j'}(t') \rangle = \delta_{jj'}\delta(t - t')$. Note that in the most naive implementation this step requires also to transform forth and back to normal modes representation, a problem that is easily circumvented by propagating the momenta in the normal modes form. The friction can be taken to be $\gamma_j = \lambda \omega_j$, where ω_js are the free ring polymer normal mode frequencies, and λ is a scaling that can be set to one to have critical damping of harmonic oscillations, or to a smaller value when one needs a more gentle thermostatting. The centroid mode (that has zero frequency in the free particle limit) can be treated separately, using either a white-noise Langevin thermostat or stochastic velocity rescaling[39,31].

2.3 Approximate Quantum Dynamics

Even though, strictly speaking, path integral molecular dynamics is just a sampling technique to obtain static, equilibrium averages, it can be used as the basis to compute ap-

proximate time correlation functions that include some of the nuclear quantum effects on dynamical properties. Examples of time-dependent properties that can be estimated in this way include diffusion coefficients, reaction rates, vibrational spectra. The two main techniques are centroid molecular dynamics (which amounts at performing microcanonical molecular dynamics on the centroid potential of mean force[11,41]) and ring-polymer molecular dynamics (which corresponds to molecular dynamics on the ring polymer potential energy surface, using the physical masses in the definition of the Hamiltonian[42,43,12]). The use of both these techniques can be partially justified based on how they can capture quantum mechanical behavior in some limits, or as approximations to more rigorous formulations of quantum dynamics[44–48]. They typically behave similarly for dynamical properties that evolve on a long time scale, but exhibit evident artifacts for short-time dynamics, e.g. for the stretching modes in IR spectra[49]. Interestingly, using physical masses but thermostatting the internal modes of the ring polymer – an approach that is half-way between the two techniques – eliminates the most apparent artifacts[50,51], providing a reliable albeit approximate method to probe the impact of quantum nuclei on vibrational properties.

3   Accelerating Convergence with Colored Noise

An alternative approach to accelerate the convergence of physical observables in path integral molecular dynamics combines a simulation with a small number of replicas with correlated-noise Langevin dynamics[52,53]. The idea is that a non-equilibrium Langevin dynamics generates frequency-dependent fluctuations, that can be used to mimic the effect of zero-point energy on a multi-dimensional harmonic oscillator, without the need of knowing the normal modes frequencies or eigenvectors. By combining colored noise with path integral molecular dynamics one can obtain a simulation protocol that, in the harmonic limit, yields exact results for any number of beads. As the number of replicas is increased, this protocol converges by construction to PIMD, and so the method can be made as accurate as one wishes, and in most cases has smaller errors than plain PIMD for any number of replicas.

3.1   A Brief Introduction to Colored Noise

Stochastic differential equations (SDEs) combine a random/noisy term together with deterministic prescriptions for the time evolution of a system[54]. The prototypical example for the application of SDEs to a physical problem is the Langevin equation, that was originally introduced as a model for Brownian motion[55], and has been used extensively to describe the coupling between an open system and an external bath[56]. Langevin dynamics introduce a friction term $-\gamma p$ and a Gaussian random force ξ on top of Hamiltonian dynamics

$$
\begin{aligned}
\dot{q} &= p/m \\
\dot{p} &= -V'(q) - \gamma p + \sqrt{2m\gamma/\beta}\xi(t).
\end{aligned}
\tag{13}
$$

The balance between the noisy force (which is uncorrelated in time, $\langle \xi(t)\xi(t')\rangle = \delta(t - t')$) and the friction guarantees canonical sampling at inverse temperature β.

If the problem of integrating out the degrees of freedom associated with the bath is considered carefully[57,58], one will find that in the general case it is not possible to represent
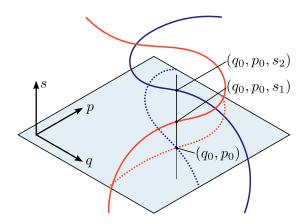
Figure 2: Two Markovian trajectories in the extended (q, p, s) space will appear as history-dependent, non-Markovian trajectories when projected into the (q, p) subspace.

the effect of the bath by equations as (13), but that the friction and the noisy force have to be associated to a finite memory kernel, leading to the non-Markovian dynamics

$$\dot{q} = p/m$$
$$\dot{p} = -V'(q) - \int_{-\infty}^{t} K(t - s)p(s)\mathrm{d}s + \sqrt{2m/\beta}\zeta(t). \tag{14}$$

The autocorrelation of the noisy force $H(t) = \langle \zeta(t)\zeta(0) \rangle$ must be related to the memory of the friction $K(t)$ by a fluctuation-dissipation relation $H(t) = K(t)$ to guarantee constant-temperature sampling. Besides giving the possibility of modeling the statistical mechanics and the dynamical behavior of an open system[59,60], a generalized Langevin equation of this form gives a great deal of freedom for manipulating the sampling properties of a molecular dynamics simulation by changing the form of the memory kernels $K(t)$ and $H(t)$. However, integrating an equation of motion of this form is not very practical, since one would need to generate Gaussian random numbers with prescribed correlation[61], to store the past trajectory of $p(t)$ and to perform an integral over the past history to compute the friction term.

Fortunately, one can reverse the reasoning that leads to Eqs. (14) starting from Markovian equations for the combined system/bath ensemble to show that a non-Markovian dynamics for (q, p) can be obtained by supplementing the physical variables with a vector $\mathbf{s}$ of n_s fictitious momenta[62–64], and integrating the Markovian equations

$$\dot{q} = p/m$$
$$\begin{pmatrix} \dot{p} \\ \dot{\mathbf{s}} \end{pmatrix} = \begin{pmatrix} -V'(q) \\ \mathbf{0} \end{pmatrix} - \begin{pmatrix} a_{pp} & \mathbf{a}_p^T \\ \bar{\mathbf{a}}_p & \mathbf{A} \end{pmatrix} \begin{pmatrix} p \\ \mathbf{s} \end{pmatrix} + \sqrt{m} \begin{pmatrix} b_{pp} & \mathbf{b}_p^T \\ \mathbf{b}_p & \mathbf{B} \end{pmatrix} \begin{pmatrix} \boldsymbol{\xi} \end{pmatrix}, \tag{15}$$

in the extended phase space comprising $(q, p, \mathbf{s})$. The additional degrees of freedom encode the memory of the system relative to the (q, p) physical variables - e.g. two trajectories

11

starting at the same (q, p) position (that would necessarily evolve in the same way in a Markovian context) can now lead to different time evolution, if they correspond to different values of the s momenta (Figure 2). Let $\mathbf{A}_p$ and $\mathbf{B}_p$ be the full $(n_s + 1) \times (n_s + 1)$ matrices in Eq. (15). The fluctuation-dissipation relation that guarantees canonical sampling is then $\mathbf{A}_p + \mathbf{A}_p^T = \beta^{-1} \mathbf{B}_p \mathbf{B}_p^T$. Eq. (15) corresponds to a non-Markovian dynamics (14) with $K(t) = 2a_{pp}\delta(t) - \mathbf{a}_p^T e^{-|t|\mathbf{A}} \bar{\mathbf{a}}_p$.

Formulating the Generalized Langevin Equation (GLE) in this Markovian form is not only convenient for the sake of integrating the dynamics on a computer. If one considers a harmonic model, for which the force $V'(q) = m\omega^2 q$ is linear, the whole set of dynamical equations for $(q, p, \mathbf{s})$ constitutes a linear, Markovian stochastic differential equation (an Ornstein-Uhlenbeck process[54]) that can be solved analytically. The opportunity of computing exactly and inexpensively the response of a normal mode of frequency ω to a given set of GLE parameters $\mathbf{A}_p$ and $\mathbf{B}_p$ makes it possible to optimize iteratively the parameters to fulfill the desired sampling properties (e.g. ergodicity, small disturbance on the dynamics, etc.) as a function of the normal mode frequency.

A crucial feature of Eqs. (15) that makes this approach very useful is that, when applying identical equations with independent random terms to a set of arbitrarily-coupled degrees of freedom, the dynamical behavior is invariant to a unitary transformation of the coordinates to which the stochastic dynamics is applied. This property, that is a consequence of the linear nature of the SDE and of the Gaussian statistics of $\boldsymbol{\xi}$, means that one can apply the GLEs to the Cartesian coordinates of a set of atoms and obtain the same response as if the GLEs had been applied to the normal modes of the system. Hence, one can optimize $\mathbf{A}_p$ and $\mathbf{B}_p$ based on a simple, one-dimensional harmonic oscillator model, but the predictions will be verified in a realistic physical model without the need of knowing explicitly the normal modes basis. This idea has been applied to obtain efficient sampling of many frequencies at the same time, and to avoid disrupting adiabatic decoupling in Car-Parrinello molecular dynamics[40,65], but also to thermostat efficiently PIMD[31] and to stabilize resonances in multiple time step integration[66].

3.2 Quantum thermostats

Generalized Langevin dynamics that satisfy the fluctuation-dissipation relation can be used to alter the dynamical properties of a trajectory, and also to make sampling more efficient in PIMD[31]. By releasing the constraint $\mathbf{A}_p + \mathbf{A}_p^T = \beta^{-1} \mathbf{B}_p \mathbf{B}_p^T$, one is effectively in a non-equilibrium scenario, that can be understood in terms of simultaneous coupling to multiple heat baths at different temperature. The practical effect of this simulation protocol is that a steady state will be reached, in which normal modes of different frequency will have fluctuations that are consistent with different values of the temperature. This peculiar non-equilibrium GLE is particularly useful to model inexpensively the quantum nature of atomic nuclei.

To see how, consider a one-dimensional harmonic oscillator of frequency ω, sampled canonically at inverse temperature $\beta = 1/k_B T$. The phase-space distribution of position and momentum are Gaussian distributions $\rho(p) \propto \exp{-p^2/2\sigma_p^2}$ and $\rho(q) \propto \exp{-q^2/2\sigma_q^2}$, irrespective of the underlying classical or quantum description. However, the classical and quantum distributions differ because of the magnitude of fluctuations: for a classical oscillator $\sigma_p^2 = m/\beta$ and $\sigma_q^2 = 1/\beta m\omega^2$, while for a quantum oscilla-
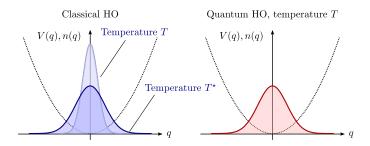
Figure 3: The phase-space distribution of positions for a harmonic oscillator at finite temperature T is a Gaussian, both in classical and quantum mechanical treatments. The only difference between the two cases is the amplitude of the distribution. Quantum fluctuations can be mimicked in a classical context by increasing the temperature $T^\star$, that however will be frequency dependent.

tor $\sigma_p^2 = m\frac{\hbar\omega}{2}\coth\frac{\beta\hbar\omega}{2}$ and $\sigma_q^2 = \frac{\hbar}{2m\omega}\coth\frac{\beta\hbar\omega}{2}$. A classical oscillator would exhibit the very same behavior as the quantum mechanical one if it were modeled at an effective temperature

$$T^\star(\omega) = \frac{\langle p^2 \rangle}{mk_B} = \frac{m\langle q^2 \rangle}{k_B} = \frac{\hbar\omega}{2k_B}\coth\frac{\hbar\omega}{2k_BT}. \tag{16}$$

such that the fluctuations match those predicted by quantum mechanics at temperature T (see Figure 3).

The problem with this idea is that the effective temperature $T^\star(\omega)$ depends on the oscillator frequency as well as on the physical target temperature. In the case of a multi-dimensional oscillator (that can be taken to be a decent model for a solid at low temperature) one would need to associate a different temperature to each normal mode. If one wanted to do so using conventional thermostatting, one would need to know the normal modes frequencies and phonon displacement patterns, and apply tailored white-noise thermostats at different temperatures working in the normal modes representation. The advantage of a non-equilibrium GLE formulation, instead, is that the dynamical response of the system determines the fluctuations along different directions automatically, without the need of knowing the normal modes of the system being studied. If one can fit a set of $\mathbf{A}_p$ and $\mathbf{B}_p$ parameters that enforce the frequency-dependent temperature (16), for any frequency within a range that encompasses the vibrational modes relevant for the system at hand, then it suffices to apply the same GLE to each Cartesian degree of freedom. The quantum $T^\star(\omega)$ curve is then enforced automatically, giving quantum fluctuations at the cost of conventional molecular dynamics.

This "quantum thermostat" (QT) idea[67,65] works surprisingly well also for strongly anharmonic potentials, and the main limitation when applying it to real systems does not depend much on failure to describe strongly anharmonic behaviour. One can see that large deviations from quantum behavior are caused by zero-point energy leakage, i.e. the fact that due to weak anharmonic couplings energy flows from the fast normal modes, that are thermostatted at high $T^\star$ to account for the large zero-point energy, to slow normal modes that are nearly classical[65]. This energy flow was not accounted for when designing $T^\star(\omega)$, and so there will be a (significant) deviation between the desired quasi-harmonic quan-

tum fluctuations and the actual fluctuations. This problem is common to semi-classical methods to treat quantum dynamics[68], and has been also recognized in other stochastic approaches to obtain approximate quantum effects[69,70]. Rather than trying to remedy this problem by exploiting information on the anharmonic couplings – which would be an ad-hoc, non-transferable solution, requiring in-depth knowledge of the system – one can control zero-point energy leakage by exploiting the tunability of the GLE thermostats, enforcing a strong coupling across the whole frequency range so as to counterbalance effectively the zero-point energy leakage. This approach improves significantly the performance of the quantum thermostat when applied to anharmonic problems[65,71], and made it possible to describe semi-quantitatively the role of NQEs in several real applications[72,73].

3.3 Combining Generalized Langevin Equations and PIMD

The approximations behind the quantum thermostat and related semi-classical methods are basically uncontrolled, and very hard to gauge unless it is possible to perform a harmonic analysis. Therefore, the quantum thermostat can be regarded as an inexpensive technique to assess qualitatively the importance of NQEs, but it is not recommend if one wants to infer quantitative conclusions. One could then imagine to combine colored-noise and path integral molecular dynamics: the former is only exact in the harmonic limit, while the latter converges systematically but often requires a large number of replicas and is therefore computationally demanding.

The crux is designing a GLE thermostat that enforces exact quantum fluctuations in the harmonic limit *for any number of replicas*, even in cases where PIMD alone would be far from converged. Such a PI+GLE method inherits from the quantum thermostat the property of being exact for harmonic problems, and naturally converge to (Boltzmann-sampled) PIMD when the number of beads is large enough to have a converged result in the absence of correlated noise.

In order to work out the properties of the GLE that would achieve this goal, one can proceed in the same way as with the quantum thermostat, only considering that now, in the presence of a harmonic potential of frequency ω, the ring-polymer normal mode frequencies will be changed from the free-particle value (11) and become $\omega_j = \sqrt{\omega^2 + 4P^2 \sin^2 j\pi/P}$. These are the frequencies that will be picked up by the colored-noise dynamics. Introducing a frequency-dependent *configurational* temperature $T^*(\omega) = \langle q^2 \rangle (\omega) m \omega^2 / k_B$ (momentum fluctuations are not important per se in a PIMD framework), one gets the requirement for having quantum fluctuations of the beads to be

$$
\begin{aligned}
\frac{m\omega^2}{k_B T} \langle q^2 \rangle &= \frac{m\omega^2}{P k_B T} \sum_i \langle q_i^2 \rangle = \frac{m\omega^2}{P k_B T} \sum_j \langle \tilde{q}_j^2 \rangle \\
&= \frac{1}{P} \sum_j \frac{T^*(\omega_j)/T}{\omega_j^2/\omega^2} = \frac{\hbar\omega}{2k_B T} \coth \frac{\hbar\omega}{2k_B T}.
\end{aligned}
\tag{17}
$$

Since ω_j depends on the physical frequency ω, Eq. (17) must be seen as a functional equation that defines the $T^*(\omega)$ curve. As shown in Ref.[52], Eq. (17) can be conveniently written as a function of the a-dimensional parameter $x = \beta\hbar\omega/2$, that expresses qualitatively how much an oscillator deviates from a classical behavior, and can be solved numerically with
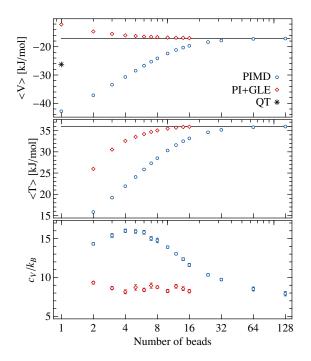
Figure 4: The average value of the potential energy, the virial kinetic energy and the constant-volume heat capacity for a simulation of a flexible water model[74] at T=298 K, plotted as a function of the number of beads. The results obtained with conventional PIMD and PI+GLE are compared, and the value of V obtained with the original quantum thermostat (QT) is also reported. (Adapted from Ref.[52])

an iterative technique. As the number of path integral replicas is increased, the curve remains constant up to a larger value of x, that corresponds to the fact that PI+GLE behaves as conventional PIMD with Boltzmann sampling of the ring-polymer Hamiltonian for oscillators with larger and larger frequency. This implies that PI+GLE is bound to converge to the exact quantum averages, just because in the large P limit it converges to PIMD.

Figure 4 shows the convergence with number of beads of potential and kinetic energy for a quantum simulation of an empirical water model[74] at room temperature, comparing plain PIMD and PI+GLE. Colored noise accelerates dramatically the convergence of observables to the quantum expectation values, and the possibility of converging results systematically makes it possible to assess the error. A careful examination of Figure 4 shows that the mean kinetic energy $\langle T \rangle$ converges somewhat more slowly than $\langle V \rangle$.

This is due to a specific shortcoming of PI+GLE, that becomes clear when one considers the expression for the centroid-virial kinetic energy estimator (7) in a one-dimensional harmonic potential, that reads:

$$
\begin{aligned}
\langle T \rangle &= \frac{1}{2\beta} + \frac{1}{2P}\omega^2 \sum_{i=0}^{P-1} \langle q_i^2 \rangle - \frac{1}{2}\omega^2 \langle \bar{q}^2 \rangle \\
&= \langle V \rangle + \frac{1}{2\beta} - \frac{1}{2}\omega^2 \langle \bar{q}^2 \rangle .
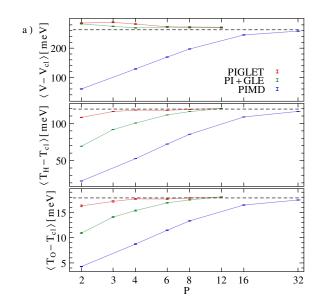\end{aligned}
\tag{18}
$$

15

Figure 5: The quantum contribution to the potential energy, and to the kinetic energy of hydrogen and oxygen atoms as computed by the centroid virial estimator for a simulation of a flexible water model[74] at T=298 K, plotted as a function of the number of beads. Note the much accelerated convergence rate of the kinetic energy when using PIGLET compared to PI+GLE. (Adapted from Ref.[53])

The quantum mechanical expectation values for potential and kinetic energy of a harmonic oscillator of frequency ω read

$$\langle V \rangle = \langle T \rangle = \frac{\hbar\omega}{4} \coth \frac{\beta\hbar\omega}{2}, \tag{19}$$

so it is not sufficient that the fluctuations of q are consistent with $\langle V \rangle = \frac{\hbar\omega}{4} \coth \frac{\beta\hbar\omega}{2}$, but it is also necessary to make sure that the fluctuations of the centroid satisfy $\frac{1}{2}\omega^2 \langle \bar{q}^2 \rangle = \frac{1}{2\beta}$. This observation is a sign of a general limitation of the PI+GLE scheme, that enforces only quantum fluctuations for the "marginal" distribution of the beads, which is sufficient to guarantee accelerated convergence of any observable that depends only on q but does not necessarily help converging more complex estimators that also depend on the correlations between different beads.

Fortunately, it is relatively easy to extend the PI+GLE idea to include further correlations. When PIMD is propagated in the normal modes representation as discussed in Section 2 one can apply different thermostats to the various normal modes, and tune them separately to warrant faster convergence of multiple estimators simultaneously. This idea is used for instance in the PIGLET scheme[53] to converge simultaneously structural observables and the centroid-virial kinetic energy estimator, obtaining a further improvement over the PI+GLE scheme, as shown in Figure 5.

16

4 Conclusive Remarks

Atomistic simulations of ever-increasing accuracy are making it possible to model molecules and materials with predictive accuracy. One of the fundamental physical effects that is still ignored in many simulations involves the deviations from classical behavior of the atomic nuclei, that is very pronounced, at room temperature and below, for the lightest elements – hydrogen in particular. The importance of nuclear quantum effects has been demonstrated very clearly, for instance, by pioneering path integral molecular dynamics studies of water and charged water defects[75, 10, 76, 77], and it is essential when it comes to compare with experiments – such as deep inelastic neutron scattering[2, 78, 14, 72, 79], or isotope fractionation measurements[80–83, 17, 84] – for which the quantum mechanical nature of nuclei is the very origin of the observed signal. Nevertheless in many modern simulations of water and H-containing compounds nuclear quantum effects are still ignored. This is at times justified: in neat water quantum effects along different molecular directions cancel out almost perfectly[74, 85, 86], and so overall NQEs on many thermodynamic quantities are small. However, simulations that aim at probing the dissociation of covalent O-H bonds[87, 88], studying self-ionization or re-combination of hydronium and hydroxide[89], and more in general quantitatively characterizing the behavior of hydrogen bonds in unusual environments (high pressure, confinement, interfaces, solvation shells of ions, . . .) should always at the very least assess whether NQEs do or do not play a major role. In many cases, the only excuse for not including NQEs is the large computational overhead involved in path integral simulations, combined with the technical complexity of the method and the fact that few up-to-date, well-maintained implementations exist in mainstream atomistic simulation packages. In this lecture we have provided a concise but complete introduction to the theory and the practical implementation of atomistic modeling of the quantum nature of atomic nuclei by path integral molecular dynamics. When one looks beyond its cumbersome notation, PIMD is just classical dynamics in an extended phase space, which means that all the tricks that are used in classical simulations (including e.g. thermostats for sampling the canonical ensemble[9], or barostats for simulations at constant-pressure[90]) can be applied transparently. What is more, simulating NQEs does not need to involve an increase of several orders of magnitude in computational effort. High-order factorizations of the quantum partition function lead to accelerated convergence in some circumstances, and correlated-noise thermostats give an accurate estimate of NQEs in water at room temperature with as little as six path integral beads.

The implementation burden that has been traditionally associated with path integral simulations is also being reduced by the introduction of modularly designed packages that deal with the PIMD aspects of the simulation while delegating to external codes the evaluation of energy and forces. An example of this concept is given by i-PI[88], an open-source[d] Python package that is already interfaced with the development versions of CP2K[91, 92], LAMMPS[93] and FHI-AIMS[94]. When one also considers that the evaluation of energy and forces over multiple PI replicas involves a trivial layer of concurrency, that is ideally adapted to massively parallel high-performance computing, one can conclude that time is ripe for making the modeling of nuclear quantum effects in hydrogen-containing liquids and solids routine.

---

[d] http://epfl-cosmo.github.io/gle4md/index.html?page=ipi

References

1. M. Born and R. Oppenheimer, *Zur Quantentheorie der Molekeln*, Ann. Phys., **389**, 457–484, 1927.
2. C. Andreani, D. Colognesi, J. Mayers, G. F. Reiter, and R. Senesi, *Measurement of momentum distribution of light atoms and molecules in condensed matter systems using inelastic neutron scattering*, Adv. Phys., **54**, 377–469, 2005.
3. Gregory A Voth, *Calculation of equilibrium averages with Feynman-Hibbs effective classical potentials and similar variational approximations*, Phys. Rev. A, **44**, 5302–5305, 1991.
4. R P Feynman and A R Hibbs, *Quantum Mechanics and Path Integrals*, McGraw-Hill, New York, 1964.
5. David Chandler and Peter G Wolynes, *Exploiting the isomorphism between quantum theory and classical statistical mechanics of polyatomic fluids*, J. Chem. Phys., **74**, 4078–4095, 1981.
6. M Parrinello and A Rahman, *Study of an F center in molten KCl*, J. Chem. Phys., **80**, 860, 1984.
7. B J Berne and D Thirumalai, *On the Simulation of Quantum Systems: Path Integral Methods*, Annu. Rev. Phys. Chem., **37**, 401–424, 1986.
8. Dominik Marx and Michele Parrinello, *Ab initio path integral molecular dynamics: Basic ideas*, J. Chem. Phys., **104**, 4077, 1996.
9. Mark E Tuckerman, Dominik Marx, Michael L Klein, and Michele Parrinello, *Efficient and general algorithms for path integral Car-Parrinello molecular dynamics*, J. Chem. Phys., **104**, 5579–5588, 1996.
10. D Marx, M E Tuckerman, J Hutter, and M Parrinello, *The nature of the hydrated excess proton in water*, Nature, **397**, 601–604, 1999.
11. Jianshu Cao and Gregory A Voth, *A new perspective on quantum time correlation functions*, J. Chem. Phys., **99**, 10070–10073, 1993.
12. Scott Habershon, David E Manolopoulos, Thomas E Markland, and Thomas F Miller, *Ring-polymer molecular dynamics: quantum effects in chemical dynamics from classical trajectories in an extended phase space.*, Annu. Rev. Phys. Chem., **64**, 387–413, 2013.
13. D M Ceperley, *Path integrals in the theory of condensed helium*, Rev. Mod. Phys., **67**, 279–355, 1995.
14. Lin Lin, Joseph A Morrone, Roberto Car, and Michele Parrinello, *Displaced Path Integral Formulation for the Momentum Distribution of Quantum Particles*, Phys. Rev. Lett., **105**, 110602, 2010.
15. M. F. Herman, E J Bruskin, and B J Berne, *On path integral Monte Carlo simulations*, J. Chem. Phys., **76**, 5150, 1982.
16. Takeshi M Yamamoto, *Path-integral virial estimator based on the scaling of fluctuation coordinates: Application to quantum clusters with fourth-order propagators*, J. Chem. Phys., **123**, 104101, 2005.
17. Michele Ceriotti and Thomas E Markland, *Efficient methods and practical guidelines for simulating isotope effects.*, J. Chem. Phys., **138**, 014112, 2013.
18. Bingqing Cheng and Michele Ceriotti, *Direct path integral estimators for isotope fractionation ratios.*, J. Chem. Phys., **141**, 244112, 2014.

19. M Takahashi and M Imada, *Monte Carlo calculation of quantum systems. II. Higher order correction*, J. Phys. Soc. Japan, **53**, 3765–3769, 1984.

20. X P Li and J Q Broughton, *High-order correction to the Trotter expansion for use in computer simulation*, J. Chem. Phys., **86**, 5094, 1987.

21. M Suzuki, *Hybrid exponential product formulas for unbounded operators with possible applications to Monte Carlo simulations*, Phys. Lett. A, **201**, 425–428, 1995.

22. Seogjoo Soonmin Jang and Gregory A Voth, *Applications of higher order composite factorization schemes in imaginary time path integral simulations*, J. Chem. Phys., **115**, 7832–7842, 2001.

23. Marcin Buchowiecki and Jirí Vanícek, *Monte Carlo evaluation of the equilibrium isotope effects using the Takahashi-Imada factorization of the Feynman path integral*, Chem. Phys. Lett., **588**, 11–16, 2013.

24. Alejandro Pérez and Mark E Tuckerman, *Improving the convergence of closed and open path integral molecular dynamics via higher order Trotter factorization schemes.*, J. Chem. Phys., **135**, 064104, 2011.

25. Michele Ceriotti, Guy a. R. Brain, Oliver Riordan, and David E. Manolopoulos, *The inefficiency of re-weighted sampling and the curse of system size in high-order path integration*, Proc. R. Soc. A Math. Phys. Eng. Sci., **468**, 2–17, 2011.

26. Ondrej Marsalek, Pei-Yang Chen, Romain Dupuis, Magali Benoit, Merlin Méheut, Zlatko Bačić, and Mark E. Tuckerman, *Efficient Calculation of Free Energy Differences Associated with Isotopic Substitution Using Path-Integral Molecular Dynamics*, J. Chem. Theory Comput., **10**, 1440–1453, 2014.

27. DM M Ceperley and EL L Pollock, *Path-integral computation of the low-temperature properties of liquid 4He*, Phys. Rev. Lett., **56**, 351–354, 1986.

28. Michiel Sprik, Michael Klein, and David Chandler, *Staging: A sampling technique for the Monte Carlo evaluation of path integrals*, Phys. Rev. B, **31**, 4234–4244, 1985.

29. M. Tuckerman, B. J. Berne, and G. J. Martyna, *Reversible multiple time scale molecular dynamics*, J. Chem. Phys., **97**, 1990, 1992.

30. Loup Verlet, *Computer "Experiments" on Classical Fluids. I. Thermodynamical Properties of Lennard-Jones Molecules*, Phys. Rev., **159**, 98–103, 1967.

31. Michele Ceriotti, Michele Parrinello, Thomas E Markland, and David E Manolopoulos, *Efficient stochastic thermostatting of path integral molecular dynamics.*, J. Chem. Phys., **133**, 124104, 2010.

32. Thomas E Markland and David E Manolopoulos, *A refined ring polymer contraction scheme for systems with electrostatic interactions*, Chem. Phys. Lett., **464**, 256, 2008.

33. Thomas E Markland and David E Manolopoulos, *An efficient ring polymer contraction scheme for imaginary time path integral simulations.*, J. Chem. Phys., **129**, 024105, 2008.

34. T Schneider and E Stoll, *Molecular-dynamics study of a three-dimensional one-component model for distortive phase transitions*, Phys. Rev. B, **17**, 1302–1322, 1978.

35. Hans C Andersen, *Molecular dynamics simulations at constant pressure and/or temperature*, J. Chem. Phys., **72**, 2384–2393, 1980.

36. H J C Berendsen, J P M Postma, W F Van Gunsteren, A DiNola, and J R Haak, *Molecular dynamics with coupling to an external bath*, J. Chem. Phys., **81**, 3684, 1984.

37. Shuichi Nosé, *A unified formulation of the constant temperature molecular dynamics methods*, J. Chem. Phys., **81**, 511, 1984.

38. W G Hoover, *Canonical dynamics: Equilibrium phase-space distributions*, Phys. Rev. A, **31**, 1695–1697, 1985.

39. G Bussi, D Donadio, and M Parrinello, *Canonical sampling through velocity rescaling*, J. Chem. Phys., **126**, 14101, 2007.

40. Michele Ceriotti, Giovanni Bussi, and Michele Parrinello, *Langevin Equation with Colored Noise for Constant-Temperature Molecular Dynamics Simulations*, Phys. Rev. Lett., **102**, 020601, 2009.

41. Jianshu Cao and Gregory A Voth, *The formulation of quantum statistical mechanics based on the Feynman path centroid density. IV. Algorithms for centroid molecular dynamics*, J. Chem. Phys., **101**, 6168–6183, 1994.

42. I R Craig and D E Manolopoulos, *Quantum statistics and classical mechanics: Real time correlation functions from ring polymer molecular dynamics*, J. Chem. Phys., **121**, 3368, 2004.

43. Rosana Collepardo-Guevara, Ian R Craig, and David E Manolopoulos, *Proton transfer in a polar solvent from ring polymer reaction rate theory.*, J. Chem. Phys., **128**, 144502, 2008.

44. Seogjoo Jang and Gregory a. Voth, *A derivation of centroid molecular dynamics and other approximate time evolution methods for path integral centroid variables*, J. Chem. Phys., **111**, 2371, 1999.

45. Bastiaan J Braams and David E Manolopoulos, *On the short-time limit of ring polymer molecular dynamics.*, J. Chem. Phys., **125**, 124105, 2006.

46. Jeremy O Richardson and Stuart C Althorpe, *Ring-polymer molecular dynamics rate-theory in the deep-tunneling regime: Connection with semiclassical instanton theory*, J. Chem. Phys., **131**, 214106, 2009.

47. Timothy J H Hele and Stuart C Althorpe, *Derivation of a true ($t \rightarrow 0^+$) quantum transition-state theory. I. Uniqueness and equivalence to ring-polymer molecular dynamics transition-state-theory.*, J. Chem. Phys., **138**, 084108, 2013.

48. Seogjoo Jang, Anton V. Sinitskiy, and Gregory a. Voth, *Can the ring polymer molecular dynamics method be interpreted as real time quantum dynamics?*, J. Chem. Phys., **140**, 154103, 2014.

49. Alexander Witt, Sergei D Ivanov, Motoyuki Shiga, Harald Forbert, and Dominik Marx, *On the applicability of centroid and ring polymer path integral molecular dynamics for vibrational spectroscopy.*, J. Chem. Phys., **130**, 194510, 2009.

50. Mariana Rossi, Michele Ceriotti, and David E Manolopoulos, *How to remove the spurious resonances from ring polymer molecular dynamics.*, J. Chem. Phys., **140**, 234116, 2014.

51. Mariana Rossi, Hanchao Liu, Francesco Paesani, Joel Bowman, and Michele Ceriotti, *Communication: On the consistency of approximate quantum dynamics simulation methods for vibrational spectra in the condensed phase*, J. Chem. Phys., **141**, 181101, 2014.

52. Michele Ceriotti, David E Manolopoulos, and Michele Parrinello, *Accelerating the convergence of path integral dynamics with a generalized Langevin equation.*, J. Chem. Phys., **134**, 84104, 2011.

53. Michele Ceriotti and David E Manolopoulos, *Efficient First-Principles Calculation of the Quantum Kinetic Energy and Momentum Distribution of Nuclei*, Phys. Rev. Lett., **109**, 100604, 2012.

54. C W Gardiner, *Handbook of Stochastic Methods*, Springer, Berlin, third edition, 2003.

55. P Langevin, *The theory of Brownian movement*, CR Acad. Sci, **146**, 530, 1908.

56. Robert Zwanzig, *Nonequilibrium statistical mechanics*, Oxford University Press, New York, 2001.

57. Robert Zwanzig, *Memory Effects in Irreversible Thermodynamics*, Phys. Rev., **124**, 983–992, 1961.

58. B J Berne and Dieter Forster, *Topics in Time-Dependent Statistical Mechanics*, Annu. Rev. Phys. Chem., **22**, 563–596, 1971.

59. L Kantorovich, *Generalized Langevin equation for solids. I. Rigorous derivation and main properties*, Phys. Rev. B, **78**, 94304, 2008.

60. L. Stella, C. D. Lorenz, and L. Kantorovich, *Generalized Langevin equation: An efficient approach to nonequilibrium molecular dynamics of open systems*, Phys. Rev. B, **89**, 134303, 2014.

61. A A Maradudin, T Michel, A R McGurn, and E R Méndez, *Enhanced backscattering of light from a random grating*, Ann. Phys. (N. Y)., **203**, 255–307, 1990.

62. F Marchesoni and P Grigolini, *On the extension of the Kramers theory of chemical relaxation to the case of nonwhite noise*, J. Chem. Phys., **78**, 6287, 1983.

63. Craig C. Martens, *Qualitative dynamics of generalized Langevin equations and the theory of chemical reaction rates*, J. Chem. Phys., **116**, 2516–2528, 2002.

64. J Luczka, *Non-Markovian stochastic processes: Colored noise*, Chaos, **15**, 26107, 2005.

65. Michele Ceriotti, Giovanni Bussi, and Michele Parrinello, *Colored-Noise Thermostats à la Carte*, J. Chem. Theory Comput., **6**, 1170–1180, 2010.

66. Joseph A Morrone, Thomas E Markland, Michele Ceriotti, and B J Berne, *Efficient multiple time scale molecular dynamics: Using colored noise thermostats to stabilize resonances.*, J. Chem. Phys., **134**, 14103, 2011.

67. Michele Ceriotti, Giovanni Bussi, and Michele Parrinello, *Nuclear Quantum Effects in Solids Using a Colored-Noise Thermostat*, Phys. Rev. Lett., **103**, 30603, 2009.

68. Scott Habershon and David E Manolopoulos, *Zero point energy leakage in condensed phase dynamics: An assessment of quantum simulation methods for liquid water*, J. Chem. Phys., **131**, 244518, 2009.

69. Hichem Dammak, Yann Chalopin, Marine Laroche, Marc Hayoun, and Jean-Jacques Greffet, *Quantum Thermal Bath for Molecular Dynamics Simulation*, Phys. Rev. Lett., **103**, 190601, 2009.

70. O. N. Bedoya-Martínez, Jean-Louis Barrat, and David Rodney, *Computation of the thermal conductivity using methods based on classical and quantum molecular dynamics*, Phys. Rev. B, **89**, 014303, 2014.

71. Michele Ceriotti, *A novel framework for enhanced molecular dynamics based on the generalized Langevin equation*, PhD thesis, ETH Zürich, 2010.

72. Michele Ceriotti, Giacomo Miceli, Antonino Pietropaolo, Daniele Colognesi, Angeloclaudio Nale, Michele Catti, Marco Bernasconi, and Michele Parrinello, *Nuclear quantum effects in ab initio dynamics: Theory and experiments for lithium imide*, Phys. Rev. B, **82**, 174306, 2010.

73. Ali A Hassanali, Jerome Jérôme Cuny, Michele Ceriotti, Chris J Pickard, and Michele Parrinello, *The Fuzzy Quantum Proton in the Hydrogen Chloride Hydrates*, J. Am. Chem. Soc., **134**, 8557–8569, 2012.

74. Scott Habershon, Thomas E Markland, and David E Manolopoulos, *Competing quan-*

tum effects in the dynamics of a flexible water model., J. Chem. Phys., **131**, 24501, 2009.

75. Hsiao S Mei, Mark E Tuckerman, Diane E Sagnella, and Michael L Klein, *Quantum Nuclear ab Initio Molecular Dynamics Study of Water Wires*, J. Phys. Chem. B, **102**, 10446–10458, 1998.

76. P L Geissler, C Dellago, D Chandler, J Hutter, and M Parrinello, *Autoionization in liquid water*, Science, **291**, 2121–2124, 2001.

77. Mark E Tuckerman, Dominik Marx, and Michele Parrinello, *The nature and transport mechanism of hydrated hydroxide ions in aqueous solution.*, Nature, **417**, 925–9, 2002.

78. G. Reiter, C. Burnham, D. Homouz, P. Platzman, J. Mayers, T. Abdul-Redah, A. Moravsky, J. Li, C.-K. Loong, and A. Kolesnikov, *Anomalous Behavior of Proton Zero Point Motion in Water Confined in Carbon Nanotubes*, Phys. Rev. Lett., **97**, 247801, 2006.

79. Giovanni Romanelli, Michele Ceriotti, David E Manolopoulos, Claudia Pantalei, Roberto Senesi, and Carla Andreani, *Direct Measurement of Competing Quantum Effects on the Kinetic Energy of Heavy Water upon Melting*, J. Phys. Chem. Lett., **4**, 3251–3256, 2013.

80. 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, *Isotope Fractionation and Atmospheric Oxygen: Implications for Phanerozoic O2 Evolution*, Science, **287**, 1630–1633, 2000.

81. T E Markland and B J Berne, *Unraveling quantum mechanical effects in water using isotopic fractionation*, Proc. Natl. Acad. Sci. USA, **109**, 7988–7991, 2012.

82. Yuki Nagata, Ruben E Pool, Ellen H G Backus, and Mischa Bonn, *Nuclear Quantum Effects Affect Bond Orientation of Water at the Water-Vapor Interface*, Phys. Rev. Lett., **109**, 226101, 2012.

83. Jian Liu, Richard S. Andino, Christina M. Miller, Xin Chen, David M. Wilkins, Michele Ceriotti, and David E. Manolopoulos, *A Surface-Specific Isotope Effect in Mixtures of Light and Heavy Water*, J. Phys. Chem. C, **117**, 2944–2951, 2013.

84. Lu Wang, Michele Ceriotti, and Thomas E. Markland, *Quantum fluctuations and isotope effects in ab initio descriptions of water*, J. Chem. Phys., **141**, 104502, 2014.

85. X.-Z. Li, B Walker, and A Michaelides, *Quantum nature of the hydrogen bond*, Proc. Natl. Acad. Sci. USA, **108**, 6369–6373, 2011.

86. Ross H McKenzie, Christiaan Bekker, Bijyalaxmi Athokpam, and Sai G Ramesh, *Effect of quantum nuclear motion on hydrogen bonding.*, J. Chem. Phys., **140**, 174508, 2014.

87. Michele Ceriotti, Jérôme Cuny, Michele Parrinello, and David E Manolopoulos, *Nuclear quantum effects and hydrogen bond fluctuations in water.*, Proc. Natl. Acad. Sci. USA, **110**, 15591–6, 2013.

88. Michele Ceriotti, Joshua More, and David E. Manolopoulos, *i-PI: A Python interface for ab initio path integral molecular dynamics simulations*, Comput. Phys. Commun., **185**, 1019–1026, 2014.

89. Ali Hassanali, Meher K Prakash, Hagai Eshet, and Michele Parrinello, *On the recombination of hydronium and hydroxide ions in water.*, Proc. Natl. Acad. Sci. USA, **108**, 20410–5, 2011.

90. Glenn J Martyna, Adam Hughes, and Mark E Tuckerman, *Molecular dynamics algorithms for path integrals at constant pressure*, J. Chem. Phys., **110**, 3275, 1999.
91. "CP2K", http://www.cp2k.org.
92. Joost VandeVondele, Matthias Krack, Fawzi Mohamed, Michele Parrinello, Thomas Chassaing, and Jürg Hutter, *Quickstep: Fast and accurate density functional calculations using a mixed Gaussian and plane waves approach*, Comput. Phys. Commun., **167**, 103–128, 2005.
93. Steve Plimpton, *Fast Parallel Algorithms for Short-Range Molecular Dynamics*, J. Comput. Phys., **117**, 1–19, 1995.
94. Volker Blum, Ralf Gehrke, Felix Hanke, Paula Havu, Ville Havu, Xinguo Ren, Karsten Reuter, and Matthias Scheffler, *Ab initio molecular simulations with numeric atom-centered orbitals*, Comput. Phys. Commun., **180**, 2175–2196, 2009.