

Algorithms and computer code for *ab initio* path integral molecular dynamics simulations

Joshua N. More

St Hilda's College, University of Oxford



A thesis submitted for the degree of *Doctor of Philosophy*

Hilary term, 2015

Algorithms and computer code for *ab initio* path integral molecular dynamics simulations

Joshua N. More, St Hilda's College

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor
of Philosophy of the University of Oxford

Hilary Term, 2015

Abstract

This thesis presents i-PI, a new path integral molecular dynamics code designed to capture nuclear quantum effects in *ab initio* electronic structure calculations of condensed phase systems. This software has an implementation of estimators used to calculate a wide range of static and dynamical properties and of state-of-the-art techniques used to increase the computational efficiency of path integral simulations. i-PI has been designed in a highly modular fashion, to ensure that it is as simple as possible to develop and implement new algorithms to keep up with the research frontier, and so that users can take maximum advantage of the numerous electronic structure programs which are freely available without needing to rewrite large amounts of code.

Among the functionality of the i-PI code is a novel integrator for constant pressure dynamics, which is used to investigate the properties of liquid water at 750 K and 10 GPa, and efficient estimators for the calculation of single particle momentum distributions, which are used to study the properties of solid and liquid ammonia. These show respectively that i-PI can be used to make predictions about systems which are both difficult to study experimentally and highly non-classical in nature, and that it can illustrate the relative advantages and disadvantages of different theoretical methods and their ability to reproduce experimental data.

Acknowledgements

When historians later look back at our time, to try to pinpoint the exact point at which it all went wrong, I suspect that the moment which they will identify as the root cause was when the University of Oxford decided to bestow upon me the title of Doctor of Philosophy, with all the implied academic heritage and authority that entails. Little did they know that by doing this, as harmless as it might have seemed at the time (although given that "Dr More" could have come directly out of an Ian Fleming novel this is debatable), they had put the first step of my master plan in motion, and thus doomed the world. This is my supervillain origin story, and, like all good supervillains, I feel the need to gloat, safe in the knowledge that no-one ever reads a Doctoral thesis.

Since I'm currently in an uncharacteristically good mood, I will now deign to recognize those who, knowingly or unknowingly, contributed to my bringing about the end of civilization. Firstly, my project supervisor, David Manolopoulos, whose hands-off approach to supervising instilled in me the self-reliance I would need in later years when ruling as an all-powerful despot, and whose natural authority taught me much in the proper way of dealing with underlings.

Secondly, I would like to thank Michele Ceriotti, whose intelligence and ingenuity would have made him a fantastic ally in my schemes but for the unfortunate fact that he is also one of the nicest people I have ever met, for his help with the development of i-PI, teaching me much of the necessary background theory and for overseeing large parts of my work.

Next, my partner-in-crime and right hand man, David Wilkins, who has been the sounding board for my work, my writing and my evil schemes throughout my DPhil, and whose careful proofreading of this thesis has been instrumental in making it a document even a future ruler of the world can be proud of.

Finally, I'd like to acknowledge my friends and colleagues in the Wilson and Barford groups, who managed to keep me sane (or at least at an societally acceptable level of insanity) via our daily conversations over coffee, and to my family, for putting up with an evil genius for all these years without disowning me or reporting me to the authorities.

Contents

1	Introduction	1
2	Basic Theory	9
2.1	Path integral molecular dynamics	9
2.2	Generalized Langevin equations	15
3	Algorithmic details	23
3.1	Normal-mode transformation	23
3.2	Liouvillian formalism	26
3.3	Trotter splitting	27
3.4	Path integral Langevin equation dynamics	28
4	i-PI	33
4.1	Electronic structure	33
4.2	Empirical potentials	35
4.3	Density functional theory	38
4.4	Computational challenges	41
4.5	Program goals and language specification	43
4.6	Data communication	45
4.7	Program functionality and implementation	49

4.7.1	Ring polymer contraction	52
4.7.2	Reweighted sampling estimators	53
4.7.3	Dynamical property calculations	55
5	Constant pressure dynamics	59
5.1	Background	59
5.2	Theory	61
5.3	Implementation	65
5.4	Testing the implementation	70
5.5	An example application: high pressure water	73
5.5.1	Computational details	74
5.5.2	Results	76
6	Calculation of momentum distributions	83
6.1	Background	83
6.2	Experimental techniques	88
6.2.1	Deep inelastic neutron scattering	88
6.2.2	Impulse approximation	91
6.2.3	Momentum distribution calculation	92
6.2.4	Beyond the impulse approximation	93
6.2.5	Anisotropic Gaussian approximations	94
6.3	Theoretical techniques	97
6.3.1	Quasi-harmonic approximation	97
6.3.2	Open-path integral methods	99
6.3.3	Anisotropic Gaussian methods	103
6.4	Application to liquid and solid ammonia	111
6.4.1	Computational details	113

6.4.2 Results	116
7 Conclusions	133
Appendices	139
A Stochastic Liouvillian operators	139
A.1 General formalism	139
A.2 Drift and diffusion constants	142
B Constant pressure integration algorithm	145
B.1 Applying the Langevin thermostat	145
B.2 Propagating the forces	147
B.3 Propagating the centroid motion	149
B.4 Propagating the free ring polymer motion	150

Chapter 1

Introduction

Molecular dynamics (MD) has been used extensively as a sampling tool for investigating the properties of condensed phase materials [1, 2]. MD uses physical laws of motion to explicitly generate nuclear trajectories. Within any given theoretical model, this technique is in principle capable of calculating numerically exact statistical and dynamical properties for a particular system. This means that MD provides a baseline against which to test approximate physical models [3], and also that the validity of a particular form of the electronic potential can be gauged by comparison between MD and experiment [4–7].

Also, since the nuclear configurations are known and the dynamics that are generated produce physically realistic trajectories, MD can be used as a probe of microscopic static and dynamical properties that are not easily accessible to experiment [8–10]. Furthermore, MD can be used to study systems at high pressure or temperature which are difficult for experiments to observe [11–16]. MD and experiment are therefore complementary, in the sense that the strengths of one can often cover for the deficiencies of the other.

However, nuclear quantum effects (NQE) such as quantum tunnelling and zero point energy (ZPE) can often cause significant deviations from classical behaviour

when studying systems at low temperature or with light nuclei, especially hydrogen, and these effects are neglected in most MD simulations. It is therefore useful to have a theory that goes beyond Newton's laws of motion as a basis for generating trajectories, so that a broader range of phenomena can be studied using MD techniques.

The state-of-the-art method for including NQE in static property calculations is path integral molecular dynamics (PIMD) [17,18], which simulates quantum statistics by using classical dynamics in an extended phase space. Much of the original MD framework is maintained in this method, so many of the algorithms used to speed up or broaden the applicability of MD simulations can be carried over to PIMD. There has also been much research in recent years devoted to reducing the size of the extended phase space needed to converge NQE [19–21] and extending the use of the path integral formalism to calculate approximate dynamical properties [22–26]. As such, PIMD is being used for an increasingly large number of problems.

One way of further increasing the range of systems that can be treated accurately with PIMD is to calculate the interaction potential with *ab initio* electronic structure codes. These use approximate solutions to the electronic Schrödinger equation such as density functional theory (DFT) to calculate the potential energy surface on which nuclear trajectories are propagated.

DFT has some key advantages over empirical models. Firstly, the functionals commonly used in DFT calculations are highly transferable and so do not need to be re-parameterized for every new system studied. This means that they are particularly suited to modelling systems under extreme conditions where forcefields designed for use at room temperature and pressure tend to fail.

Secondly, while empirical potentials are usually fitted to reproduce experimental results, there is no guarantee that they will also be suitable to describe any properties not used in the fit. For example, most empirical models are unable to provide

any insight into bond dissociation or charge redistribution processes, even under the conditions at which they were parameterized. This is not a problem faced by DFT, as the electron density (on which the above properties depend) is generated during the calculation of the interaction potential.

However, while the *ab initio* path integral molecular dynamics (AI-PIMD) formalism was established about 20 years ago [27–29], the number of practical applications seen in the literature has been limited due to the extreme cost of the electronic structure calculation. The earliest examples of AI-PIMD required a computational tour-de-force, as can be seen in the work of Marx and Parrinello [30–32], and Tuckerman and Klein [33–36].

More recently, there has begun to be a more widespread adoption of these methods. In particular, Michaelides [37, 38], Morrone and Car [39, 40] have worked on the *ab initio* simulation of NQE in hydrogen bond networks, and Michaelides has also studied the properties of low temperature hydrogen [41–43]. However, the time required for these calculations is still a huge barrier to the widespread adoption of AI-PIMD.

This is particularly frustrating since modelling NQE explicitly is of paramount importance for *ab initio* molecular dynamics (AI-MD): while empirical potentials can be fitted to include NQE in a mean field sense, no such fitting is used for techniques such as DFT. It is therefore desirable to have methods which can reduce the computational cost of the path integral machinery, so that in situations where AI-MD will be systematically wrong AI-PIMD can be used instead.

Among the methods that have been developed to reduce the size of the PIMD extended phase space, the most generally applicable in the context of AI-PIMD simulations is the path integral generalized Langevin equation kinetic energy (PIGLET) algorithm [21]. This uses a non-equilibrium stochastic thermostat to bias the path

integral dynamics in such a way that correlations consistent with the quantum ensemble average are produced in the harmonic limit, even in classical MD. Away from the harmonic limit this technique increases the rate of convergence with respect to the size of the extended phase compared to PIMD alone, often by an order of magnitude, and so can make AI-PIMD simulations much more accessible.

However, computational expense is not the only obstacle that must be overcome before AI-PIMD can find widespread use. Another major issue is that DFT codes tend to be very complicated, so that a substantial implementation effort is required to combine these with path integral methods, let alone keep up with new developments in the field. To avoid this problem, the approach we have taken is to create a single library of PIMD algorithms that can be interfaced with any DFT code [44], with minimal implementation effort. This brings together the latest PIMD techniques and algorithms and makes them available for general use, with the aim of making AI-PIMD simulations routine.

The remainder of this thesis is organized as follows. Chapter 2 will derive a mathematical framework for the path integral formalism and define some of the concepts used throughout this work. There will also be a brief discussion of some of the new techniques that have been designed for PIMD simulations, to be developed further in later chapters. These discussions will motivate the rest of the work, as they show how recent progress in the field has made AI-PIMD simulations more feasible.

Chapter 3 will expand on this framework, giving some of the computational details which have been used in this work. These have been discussed previously in the literature, but are repeated here briefly as the notation will be needed in later chapters. In particular, the formalism that will be used to discretize the equations of motion will be discussed, and the well-documented [45] application to constant temperature dynamics shown.

Chapter 4 will introduce i-PI [44], original software developed to run PIMD simulations using *ab initio* electronic structure codes. The central design principle of this code is that it should reduce the implementation effort required to run AI-PIMD simulations as much as possible, firstly by reducing the amount of code that needs to be written by the user and secondly by being clear and easy to understand.

To achieve this aim i-PI uses a client-server model: the main dynamics loop is run on the server, which connects to one or many instances of a client electronic structure code via internet sockets. This method of data communication has the advantages of both keeping the codes completely separate (it is not even necessary to run them on the same computer) and allowing a primitive distributed programming paradigm to be realized, since as many connections can be made as necessary to fully exploit the trivial level of parallelism over the PIMD extended phase space. The implementation effort required to adapt the *ab initio* code to be run with i-PI is minimized by passing only the unit cell parameters and the atomic positions to the client, and returning only the forces, virial and potential.

The remainder of the thesis will discuss the functionality of the i-PI package which has either not been described before in the literature or has been explored in more detail by this work. Chapter 5 will discuss a new set of dynamical algorithms which allow the molecular dynamics formalism to be extended to sample constant pressure ensembles. These follow previous work [46–49] by including the cell volume as a dynamical variable in the equations of motion, the fluctuations of which mimic the action of an external barostat. This is particularly useful when working with electronic potentials that have not been parameterized to reproduce the experimental density curve, as in many cases experiment and theory can give very different equilibrium volumes. Since temperature and pressure are usually chosen as the independent experimental variables, constant pressure simulations can give results which

can be more directly compared with experiment than those done at constant volume.

In particular, constant pressure dynamics are preferred when accurate density measurements are difficult, for example when high temperatures and pressures are being considered. To illustrate this, constant pressure AI-PIMD will be used to study high-pressure water [44]. This will show that, even at high temperatures (750 K at 10 GPa), including NQE can be very important when investigating properties which depend on the tail of the bond length distribution. In particular, the frequency of transient fluctuations of the proton along the OH bond which change its nearest neighbour oxygen and a qualitative description of the bond dissociation rate will be compared in quantum and classical simulations to assess the magnitude of the effect of quantum delocalization.

Another property which is of particular interest when considering NQE is the single particle momentum distribution. While in the classical case this is given by a simple Maxwell-Boltzmann (MB) distribution, quantum mechanically the shape of the potential energy surface also has an effect. The momentum distribution therefore contains very useful information about the underlying physical potential, and can also be used as a sensitive probe of NQE. However, this has in the past been a difficult quantity to measure, for both experiments and simulations.

Chapter 6 will discuss recent developments in the field, describing both the latest experimental techniques and the various theoretical models that have been used in calculations. There will then be a comparison of the efficacy of some of these methods in a study of solid and liquid ammonia. It will be shown that good approximations to the single particle momentum distribution can be found by combining a simple ansatz for its functional form and the results of a PIMD simulation. This data will then be used to test several different theoretical methods, both in comparison with each other and with experiment. Finally, the results will be used to investigate how

the physics of ammonia are influenced by its hydrogen bond network, with specific reference to similar effects in the related and well-studied system of water.

To conclude, a summary of the aims and achievements of this work will be given in chapter 7, followed by two appendices used to derive important mathematical results used in earlier chapters.

Chapter 2

Basic Theory

2.1 Path integral molecular dynamics

Under the assumption that the Born-Oppenheimer approximation [50] holds, i.e. that the electronic and nuclear contribution to molecular motion can be separated, the properties of a system with N nuclei can be described by the Schrödinger equation

$$\hat{H} |\Psi\rangle = E |\Psi\rangle \quad (2.1.1)$$

where E is the energy of the nuclear eigenstate $|\Psi\rangle$, and \hat{H} is the quantum Hamiltonian of the system, which we will assume has the form

$$\hat{H} = \hat{U} + \hat{K} = \hat{U}(\hat{\mathbf{r}}_1, \hat{\mathbf{r}}_2, \dots, \hat{\mathbf{r}}_N) + \sum_{j=1}^N \left(\frac{|\hat{\mathbf{p}}_j|^2}{2m_j} \right) \quad (2.1.2)$$

where the subscript gives the nuclear label, m_j is the mass of the j -th nucleus, and $\hat{\mathbf{p}}_j$ and $\hat{\mathbf{r}}_j$ are the momentum and position operators of the j -th nucleus respectively. $\hat{U}(\hat{\mathbf{r}}_1, \hat{\mathbf{r}}_2, \dots, \hat{\mathbf{r}}_N)$ is the operator which gives the total electronic energy given that the nuclear quantum state is fixed to be $|\Psi\rangle$, and combines the contribution of the

electron motion with that of the Coulomb potential energy between the different charged particles in the system.

In the position representation of quantum mechanics the state $|\Psi\rangle$ is described by a probability amplitude, $\psi(\mathbf{r})$ [51]. This is a very complicated object, as it is a complex valued function of all the nuclear positions over all space, and so in practice it is not possible to measure it for any given state of a condensed phase system. Therefore we are not interested in the properties of a particular quantum state, but rather in the average of these properties over the ensemble of states which have the same macroscopic observables, as this average can be determined experimentally.

If we choose to describe the system using the number of nuclei (N), the volume (V) and the temperature (T), then the relevant set of states is the canonical ensemble. The sum over states for this ensemble, Q , known as the partition function, is defined as [52]

$$Q = \text{Tr} \left[e^{-\beta \hat{H}} \right] \quad (2.1.3)$$

where the trace is over any complete set of quantum states with N nuclei, volume V , and a temperature corresponding to $\beta = 1/k_B T$, where k_B is the Boltzmann constant.

Q is a central quantity in statistical mechanics, as it acts as a normalization factor for ensemble averages of a system in thermal equilibrium. In particular, the average of an observable \hat{O} is equal to [52]

$$\langle \hat{O} \rangle = \frac{1}{Q} \text{Tr} \left[e^{-\beta \hat{H}} \hat{O} \right] \quad (2.1.4)$$

where we use angled brackets to denote an average over an ensemble.

However, while this equation can be solved in theory, a brute force calculation suffers from highly unfavourable scaling with system size and so is not feasible for most condensed phase systems.

This problem can be made tractable by transforming this quantum partition function into a classical partition function in an extended phase space using the path integral formalism [53]. In discretised form this classical partition function is

$$Q = \lim_{n \rightarrow \infty} \left[\left(\frac{1}{2\pi\hbar} \right)^{3nN} \int d\mathbf{r} \int d\mathbf{p} \exp(-\beta_n H_n) \right] \quad (2.1.5)$$

where

$$H_n = \sum_{i=0}^{n-1} \left[U(\mathbf{r}_1^{(i)}, \mathbf{r}_2^{(i)}, \dots, \mathbf{r}_N^{(i)}) + \sum_{j=1}^N \left(\frac{|\mathbf{p}_j^{(i)}|^2}{2m_j} + \frac{m_j}{2\hbar^2\beta_n^2} |\mathbf{r}_j^{(i)} - \mathbf{r}_j^{(i+1)}|^2 \right) \right] \quad (2.1.6)$$

$\beta_n = \beta/n$ and $\mathbf{r}_j^{(n)} = \mathbf{r}_j^{(0)}$.

Equation 2.1.6 is mathematically equivalent to the Hamiltonian of a system of N classical n -particle polymer chains, with a potential term that includes harmonic bonds between neighbouring particles on the same chain and an interaction potential between particles with the same value of the index i .


Due to their appearance (see figure 2.1), these chains are known as ring polymers and the particles they are made up of are called beads. Also, since the particles specified by a particular value of the bead index i correspond to a classical analogue of the original system, each of these sets of particles is called a classical replica of the system, or just a replica.

For a position-dependent observable $O(\hat{\mathbf{r}})$ we can simply use the classical analogue of equation 2.1.4 to obtain the quantum average [53], given by

$$\langle O(\hat{\mathbf{r}}) \rangle = \frac{1}{Q} \lim_{n \rightarrow \infty} \left[\left(\frac{1}{2\pi\hbar} \right)^{3nN} \int d\mathbf{r} \int d\mathbf{p} \exp(-\beta_n H_n) O_n(\mathbf{r}) \right] \quad (2.1.7)$$

$$= \lim_{n \rightarrow \infty} \left(\langle O_n(\mathbf{r}) \rangle^{\text{PI}} \right) \quad (2.1.8)$$

Classical



Quantum

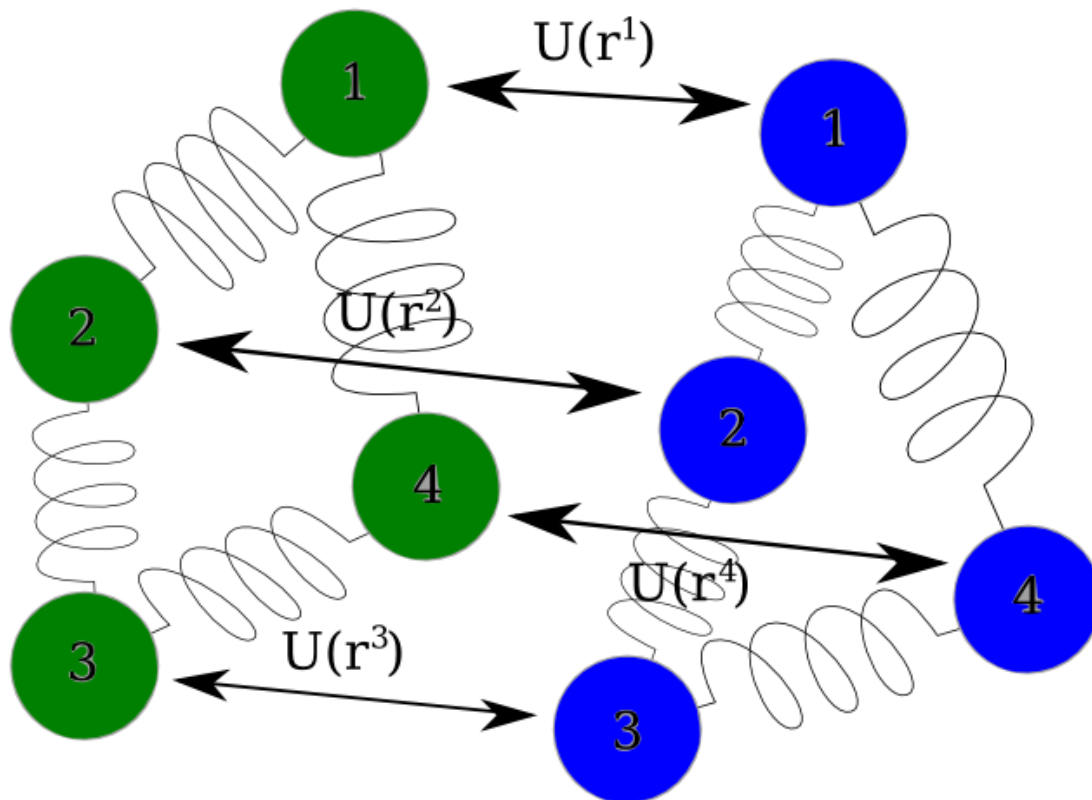


Figure 2.1: An illustration of how two interacting particles are described in classical and path integral simulations. Note in particular that in the quantum simulation the interaction potential is only calculated between beads with the same index.

where

$$O_n(\mathbf{r}) = \frac{1}{n} \sum_{i=0}^{n-1} O(\mathbf{r}^{(i)}) \quad (2.1.9)$$

and $O(\mathbf{r}^{(i)})$ is the classical value of the observable for a replica of the system with configuration $\mathbf{r}^{(i)}$. Note that the trace in equation 2.1.4 has been replaced by the classical equivalent: an integral over positions and momenta.

While equation 2.1.7 is formally equivalent to equation 2.1.4, it has an entirely classical structure rather than a quantum mechanical one, and so can be solved using well-established MD techniques [1, 2].

Under the assumption that the system is ergodic, i.e. that for any sufficiently long trajectory every state of the system is equally likely to be visited, then the sum over states in equation 2.1.7 can be replaced with an average over time. This means that if we take a trajectory through phase space generated using the physical laws of motion, then the average value of the property of interest over the length of the trajectory should be equal to the ensemble average.

The relevant laws of motion for classical mechanics are Hamilton's equations, given by

$$\dot{\mathbf{r}}_j^{(i)} = \frac{\partial H_n}{\partial \mathbf{p}_j^{(i)}} \quad (2.1.10)$$

$$\dot{\mathbf{p}}_j^{(i)} = -\frac{\partial H_n}{\partial \mathbf{r}_j^{(i)}} \quad (2.1.11)$$

While these coupled differential equations are not analytically soluble for systems of chemical interest, sufficiently accurate numerical solutions can be obtained by integrating these equations in discrete time steps (for a more detailed discussion of what constitutes a sufficiently accurate numerical solution, see ref. [1]). In this work we will use a symplectic Trotter splitting [54, 55] to derive the integration algorithm, as will be discussed more in detail in chapter 3.

However, equations 2.1.10 and 2.1.11 alone are not sufficient to properly explore phase space, as it has been shown that the redistribution of energy between the normal modes of the ring polymer is very poor [56], leading to non-ergodic behaviour.

To get over this difficulty, an extra step is included in the equations of motion of the system which simulates an exchange of energy with an external heat bath. This allows constant (N, V, T) dynamics to be generated and ensures ergodicity by redistributing energy amongst the normal modes. Again, a more detailed discussion of the algorithmic steps used will be given in chapter 3.

For observables that depend on momentum, equation 2.1.7 cannot be used directly. However, estimators can often be derived for these properties using relations derived from thermodynamics. For example, one of the quantities that is likely to be of interest is the average kinetic energy. This is momentum dependent, but can be calculated within the PIMD framework by considering the following thermodynamic identity

$$\begin{aligned} \langle \hat{H} \rangle &= \langle \hat{U} \rangle + \langle \hat{K} \rangle = -\frac{\partial \ln(Q)}{\partial \beta} \\ \therefore \langle \hat{K} \rangle &= -\frac{1}{Q} \frac{\partial Q}{\partial \beta} - \langle \hat{U} \rangle \end{aligned} \quad (2.1.12)$$

Using 2.1.5 to give $(\partial Q / \partial \beta) / Q$ and 2.1.7 to calculate $\langle \hat{U} \rangle$, and doing an integration by parts to get an estimator with a lower variance [57, 58], gives

$$\langle \hat{K} \rangle = \lim_{n \rightarrow \infty} (\langle K_n \rangle^{\text{PI}}) \quad (2.1.13)$$

where

$$K_n = \frac{3N}{2\beta} - \frac{1}{2n} \sum_{i=0}^{n-1} \sum_{j=1}^N \left(\mathbf{r}_j^{(i)} - \mathbf{r}_j^c \right) \cdot \mathbf{f}_j^{(i)} \quad (2.1.14)$$

$\mathbf{f}_j^{(i)}$ is the force acting on each bead, given by

$$\mathbf{f}_j^{(i)} = -\frac{\partial U(\mathbf{r}_1^{(i)}, \mathbf{r}_2^{(i)}, \dots, \mathbf{r}_N^{(i)})}{\partial \mathbf{r}_j^{(i)}} \quad (2.1.15)$$

and \mathbf{r}_j^c is the centroid coordinate for atom j , given by

$$\mathbf{r}_j^c = \frac{1}{n} \sum_{i=0}^{n-1} \mathbf{r}_j^{(i)} \quad (2.1.16)$$

This method can also be used to calculate other properties that can be written as a derivative of a free energy such as the magnetic susceptibility [17], the stress tensor [48] and the heat capacity [57].

2.2 Generalized Langevin equations

The formalism given above is only exact in the $n \rightarrow \infty$ limit. Fortunately, while the motion of an infinite number of beads cannot be simulated, in practice only a finite number of beads are necessary to converge the results. By increasing the number of beads until the properties of interest no longer vary, the values that would be calculated from a full solution of the Schrödinger equation can be obtained to within statistical accuracy.

Still, the computational cost of a PIMD simulation scales linearly with n , and so if large numbers of replicas are needed then the time taken to run a simulation may become prohibitive, especially if *ab initio* potentials are used. In a system for which the highest vibrational frequency is ω_{\max} the number of replicas needed is of the order of $\beta\hbar\omega_{\max}$ [59]. This means that large numbers of replicas may be required at low temperatures and for systems with stiff vibrational modes.

To improve the rate of convergence with respect to n , several techniques have been

developed. One such method is to use a higher-order approximation to the quantum partition function in equation 2.1.3 which will converge more rapidly with respect to the number of beads [19, 60–62].

This introduces a new effective Hamiltonian which includes commutators of \hat{U} and \hat{K} , and thus the force. As such, applying these methods to PIMD has been difficult, since a direct application of equation 2.1.11 will result in terms involving the spatial derivative of the force (the Hessian matrix), which is a very difficult quantity to calculate for large systems (especially with *ab initio* potentials).

However, recent work has avoided this difficulty by noting that the difference between the new effective Hamiltonian and the standard PIMD Hamiltonian can be treated as a perturbation, so higher-order estimates can be obtained from standard PIMD simulations by simply re-weighting the configurations after each step according to this perturbation [63–65]. This has allowed this method to be scaled up from small Lennard-Jones clusters [64] to condensed phase systems of chemical interest such as liquid water [63, 65, 66].

Unfortunately, this re-weighting procedure is only effective within certain limits, as it can be shown [67] that the variance of the higher-order estimators increases exponentially with the variance in the perturbation. This makes this technique much less efficient than standard PIMD above a threshold system size, or below a threshold temperature.

An alternative method that has been more recently developed is to separate the short- and long-range potential contributions and calculate them using different numbers of beads, which is known as ring-polymer contraction (RPC) [20, 68, 69]. Since in many cases the computational cost of calculating the long-range part of empirical potentials is much greater than that of the short-range, especially if Coulomb terms are present in the interaction potential [70, 71], this can substantially reduce the total

simulation time.

RPC has found applications to empirical models of liquid water, [69, 72, 73], but is difficult to apply to *ab initio* simulations since these do not have a well-defined separation of length scales. While some attempts have been made to get around this difficulty [74–77], it still remains a difficult problem and as of yet RPC has not been applied to AI-PIMD.

Instead, the method used in this thesis to speed up AI-PIMD simulations is non-equilibrium stochastic dynamics [21, 59]. This approach combines the path integral method with a tailored generalized Langevin equation (GLE) thermostat [45, 78, 79], which uses coloured noise and estimates from the harmonic limit to try to enforce the correct quantum ensemble with fewer beads.

The Langevin equation is one of the simplest examples of a stochastic differential equation. It was originally designed as a model for Brownian motion [80], combining a model of random collisions with an external heat bath with a hydrodynamical equation for viscous drag, giving

$$\dot{\mathbf{p}} = \sqrt{\frac{2m_j\gamma}{\beta}}\boldsymbol{\xi} - \gamma\mathbf{p} \quad (2.2.1)$$

where $\boldsymbol{\xi}$ is a vector of uncorrelated normally-distributed numbers, each defined such that $\langle \xi(t) \rangle = 0$ and $\langle \xi(t)\xi(t') \rangle = \delta(t - t')$, and γ is a friction term that determines the strength of the random noise.

It was later realized that propagating the momentum by a combination of equations 2.2.1 and 2.1.11 would allow MD simulations to sample the canonical ensemble [81, 82]. However, because this algorithm lacks a well-defined conserved quantity and gives trajectories that are not repeatable, except by re-running with exactly the same random number generator, other approaches, such as Nosé-Hoover thermostats [83–85], have traditionally been more widely used for this purpose.

Recently this has begun to change, as new research has introduced stochastic thermostats that both couple to the total kinetic energy rather than the individual Cartesian components [86,87] (thus reducing the disruption of the dynamics) and have a well-defined conserved energy quantity [88]. This has sparked renewed interest in stochastic thermostating, with many recent papers on the subject.

One of the most important developments is the generalization of equation 2.2.1 to coloured noise, i.e. where the random numbers are no longer uncorrelated with each other in time. This generalization is non-Markovian, in the sense that the equations of motion during any given time step depend on the preceding trajectory as well as the current state of the system, and so can be difficult to implement. Fortunately, it can be shown [89] that this is equivalent to a purely Markovian formulation in an extended momentum space, as given by

$$\begin{pmatrix} \dot{\mathbf{p}} \\ \dot{\mathbf{s}} \end{pmatrix} = -\mathbf{A} \begin{pmatrix} \mathbf{p} \\ \mathbf{s} \end{pmatrix} + \mathbf{B} \boldsymbol{\xi} \quad (2.2.2)$$

where $\dot{\mathbf{s}}$ is a set of auxilliary momenta, \mathbf{A} is the drift matrix and \mathbf{B} is the diffusion matrix. \mathbf{A} describes the damping of the momenta due to the thermostat as γ does in the original Langevin equation, and \mathbf{B} governs the influence of the stochastic forces.

Note that while equation 2.2.1 has only one adjustable parameter, equation 2.2.2 has many. The added flexibility that this gives has found many applications [45, 78, 90, 91], but here we are interested in how it can speed up the rate of convergence of PIMD simulations with respect to the number of replicas.

The fitting procedure used to choose the parameters in equation 2.2.2 relies on two observations. Firstly, the original formulation of the Langevin equation [80] used the fluctuation-dissipation theorem to derive the diffusion constant, as this ensures that the algorithm samples the equilibrium canonical ensemble. In dynamics generated

using equation 2.2.2, this theorem puts a constraint [92] on the possible choices for the matrices \mathbf{A} and \mathbf{B} . However, this is not a necessary condition for well-defined dynamics. As such, the parameters in equation 2.2.2 can be chosen without this restriction, if so desired. Secondly, the only non-linear term in equations 2.2.2 and 2.1.11 is the force term, which becomes linear in the harmonic limit. Therefore, the statistical properties of a harmonic oscillator which obeys equation 2.2.2 can be calculated analytically [93].

For example, the mean squared fluctuations in position, c_{xx} , for an one-dimensional quantum harmonic oscillator with $U(\hat{x}) = m\omega^2\hat{x}^2/2$ are given by [21]

$$\langle \hat{x}^2 \rangle = c_{xx}(\omega) = \frac{\hbar}{2m\omega} \coth\left(\frac{\beta\hbar\omega}{2}\right) \quad (2.2.3)$$

To obtain the equivalent PIMD average, the correlations introduced by the harmonic spring terms in equation 2.1.6 have to be taken into account. This can be done by changing the coordinate system from the bead positions to the normal modes of the ring polymers, using the transformation matrix given in chapter 3. Summing over the n normal modes gives

$$\langle x_n^2 \rangle^{\text{PI}} = \frac{1}{n} \sum_{k=0}^{n-1} c_{xx}(\omega'_k) = \frac{\hbar}{2nm} \sum_{k=0}^{n-1} \frac{1}{\omega'_k} \coth\left(\frac{\beta\hbar\omega'_k}{2}\right) \quad (2.2.4)$$

where $\omega'_k = \sqrt{\omega^2 + \omega_k^2}$ and ω_k is the vibrational frequency of the k -th normal mode of the free ring polymer.

Bringing all of this together, the correct quantum position fluctuations for a harmonic oscillator can be obtained for any number of beads if the adjustable parameters in equation 2.2.2 are chosen such that the PIMD average in 2.2.4 equals the ensemble average in equation 2.2.3. This can be done using an iterative fitting procedure, as discussed in ref. [21].

There are a few practical concerns that should be addressed before this method is used in practice. Firstly, in order for the fluctuations of each normal mode to sum to give the correct total, a frequency-dependent temperature distribution must be sampled. Such a non-equilibrium system might be expected to show unphysical ZPE leakage between the normal modes of different frequency [94], which would lead to incorrect statistical averages. Fortunately, so long as the parameters in equation 2.2.2 are chosen to maximize the energy transfer between the system and the thermostat, this error can be kept at an acceptable level [21].

Secondly, the above analysis only applies in the harmonic limit, and could break down in anharmonic systems. However, for most problems of interest the convergence of averages with respect to the number of replicas is dominated by the highest frequency vibrational mode, which is usually well approximated by a harmonic oscillator. Therefore, while a single bead GLE calculation may not give the correct result as it would for a harmonic oscillator, fewer beads are usually required if a GLE thermostat is used than for PIMD alone.

Finally, it is known that this formulation reduces to an equilibrium, canonical-sampling thermostat in the infinite bead limit [21], and so will give the correct average if the number of beads used is already in the regime where PIMD gives converged results, regardless of the above concerns. This method must therefore be at least as good as PIMD for any system, and will be better as long as a harmonic model is not too inaccurate.

While the above discussion is sufficient to motivate the use of this method, a more general approach will be necessary for several applications. In particular, so far only the convergence of position fluctuations has been optimized, which will only be sufficient to converge position-dependent observables such as the potential energy.

However, the argument used in this chapter could equally be applied to any fluc-

tuation. The kinetic energy, for example, depends on the fluctuations of the centroid coordinate, as it can be shown [59] that, for the same harmonic oscillator as before

$$\langle K_n \rangle^{\text{PI}} = \langle U_n \rangle^{\text{PI}} + \frac{1}{2\beta} - \frac{1}{2}m\omega^2 \langle (x^c)^2 \rangle \quad (2.2.5)$$

As the kinetic and potential energy for a harmonic oscillator are equal due to the quantum virial theorem [95], this suggests that

$$\frac{1}{2}m\omega^2 \langle (x^c)^2 \rangle = \frac{1}{2\beta} \quad (2.2.6)$$

This shows us that the centroid variable must be distributed classically, while the other normal modes must be optimized for position correlations as before. It has been shown that it is sufficient to attach a Langevin equation thermostat to the centroid, and the same non-equilibrium GLE thermostat to the remaining modes. This method, which converges both the potential and kinetic energy more rapidly than PIMD, is known as PIGLET [59], and will be used to calculate single particle momentum distributions in chapter 6.

Chapter 3

Algorithmic details

3.1 Normal-mode transformation

In theory, simply generating dynamics based on the Hamiltonian given in equation 2.1.6 would be enough to sample the correct quantum ensemble. However, unlike standard MD simulations, the existence of the spring potential means there are several problems associated with the internal motion of the ring polymers if a naïve implementation of the dynamical algorithms is used.

Firstly, as already discussed in chapter 2, for weakly quantum systems there is such a large energy separation between the highest frequency ring polymer vibrational modes that there will be very little energy transfer between them, leading to non-ergodic behaviour [56]. Secondly, even if this problem is solved, these high frequency modes necessarily have a very short timescale of motion, and so a small time step must be used in the integration algorithm to prevent sampling errors.

To illustrate how these internal modes are dealt with in a practical PIMD simulation, let us first consider the Hamiltonian of a single free ring polymer, in the absence

of any interaction potential. This has a very simple form, given by

$$H_n = \sum_{i=0}^{n-1} \left(\frac{|\mathbf{p}^{(i)}|^2}{2m} + \frac{m}{2\hbar^2\beta_n^2} |\mathbf{r}^{(i)} - \mathbf{r}^{(i+1)}|^2 \right) \quad (3.1.1)$$

Using equations 2.1.10 and 2.1.11, we get

$$\dot{\mathbf{r}}^{(i)} = \frac{\mathbf{p}^{(i)}}{m} \quad (3.1.2)$$

$$\dot{\mathbf{p}}^{(i)} = \frac{m}{\hbar^2\beta_n^2} (\mathbf{r}^{(i+1)} + \mathbf{r}^{(i-1)} - 2\mathbf{r}^{(i)}) \quad (3.1.3)$$

The only difficulty that arises in solving these equations is that the momentum equations of motion each involve more than one bead. Therefore, the greatest simplification can be made by finding a unitary transformation to a basis in which the problem is diagonal. Since the Hamiltonian is harmonic, let us consider the normal modes of the ring polymer, $\tilde{\mathbf{r}}^{(k)}$, which are the solutions to

$$\mathbf{D}\tilde{\mathbf{r}}^{(k)} = m\omega_k^2\tilde{\mathbf{r}}^{(k)} \quad (3.1.4)$$

where ω_k are the normal mode frequencies and \mathbf{D} is given by

$$\mathbf{D}_{ij} = -\frac{\partial \dot{\mathbf{p}}^{(i)}}{\partial \mathbf{r}^{(j)}} \quad (3.1.5)$$

\mathbf{D} has two useful properties which make this problem easier to solve. Firstly, since the equilibrium geometry of the ring polymer has spherical symmetry (with all beads being on top of each other), this means that each of the Cartesian components are equivalent, essentially reducing equation 3.1.4 to 3 separate n -dimensional problems. This symmetry also means that we only expect 3 zero frequency normal modes.

Secondly, since the ring polymer only has nearest-neighbour harmonic bonds \mathbf{D}

is tri-diagonal. In fact, it is mathematically equivalent to the Hamiltonian of a cyclic Hückel polyene [96], the eigenstates of which are well known. Defining a unitary transformation matrix \mathbf{C} by $\tilde{\mathbf{r}}_j^{(k)} = \sum_{i=0}^{n-1} \mathbf{r}_j^{(i)} C_{ik}$, the solution is given by [45]

$$\sqrt{n} C_{ik} = \begin{cases} 1 & k = 0 \\ \sqrt{2} \cos\left(\frac{2\pi ik}{n}\right) & 1 \leq k \leq \frac{n}{2} - 1 \\ (-1)^i & k = \frac{n}{2} \\ \sqrt{2} \sin\left(\frac{2\pi ik}{n}\right) & \frac{n}{2} + 1 \leq k \leq n - 1 \end{cases} \quad (3.1.6)$$

and the normal mode frequencies are

$$\omega_k = \frac{2}{\hbar\beta_n} \sin\left(\frac{k\pi}{n}\right) \quad (3.1.7)$$

Note that as well as being defined by the above simple matrix multiplication, the normal modes can also be calculated by a fast-Fourier transform routine [97], which is an $\mathcal{O}(N \log(N))$ algorithm. As such, in most cases of interest this transformation takes a negligible time, and so we will use the bead and normal mode representations interchangeably throughout this thesis.

In this new basis, the equations of motion have a very simple form, given by

$$\dot{\tilde{\mathbf{r}}}^{(k)} = \frac{\tilde{\mathbf{p}}^{(k)}}{m} \quad (3.1.8)$$

$$\dot{\tilde{\mathbf{p}}}^{(k)} = -m\omega_k^2 \tilde{\mathbf{r}}^{(k)} \quad (3.1.9)$$

These equations can be integrated exactly (see ref. [45] or appendix B), using

$$\begin{pmatrix} \tilde{\mathbf{p}}^{(k)}(t + \Delta t) \\ \tilde{\mathbf{r}}^{(k)}(t + \Delta t) \end{pmatrix} = \begin{pmatrix} \cos(\omega_k \Delta t) & -m\omega_k \sin(\omega_k \Delta t) \\ \frac{1}{m\omega_k} \sin(\omega_k \Delta t) & \cos(\omega_k \Delta t) \end{pmatrix} \begin{pmatrix} \tilde{\mathbf{p}}^{(k)}(t) \\ \tilde{\mathbf{r}}^{(k)}(t) \end{pmatrix} \quad (3.1.10)$$

As this integration scheme is exact it does not require a small timestep to model the high frequency harmonic motion, and so is much preferred to working in the bead representation. Also, as we will see in section 3.4, this is also the optimal framework in which to apply the thermostat needed to ensure ergodicity.

3.2 Liouvillian formalism

When considering the creation and implementation of molecular dynamics algorithms, there are several important considerations. Firstly, it is important to be able to determine any conserved quantities and the stationary distribution, since these specify which ensemble is being sampled. Secondly, since all but the simplest equations of motion must be solved numerically, we must have a method of generating integration algorithms with desirable qualities.

So far during this thesis only Hamiltonian mechanics has been used as a framework for understanding MD integrators. However, this formalism is only valid when the phase space volume in equation 2.1.5 is conserved [98,99], which will not necessarily be the case for every algorithm of interest. In particular, the constant pressure algorithm that we derive in chapter 5 has several steps which are individually non-Hamiltonian. As such, a more general method is needed.

At the most fundamental level, MD algorithms are a system of equations which describe the time evolution of a phase space. For any equation of motion then, there is a corresponding operator equation describing its effect on a probability distribution, ρ , of the form

$$\mathcal{L} \rho(\mathbf{r}, \mathbf{p}, t) = -\frac{\partial \rho(\mathbf{r}, \mathbf{p}, t)}{\partial t} \quad (3.2.1)$$

where \mathcal{L} is known as the Liouvillian operator.

This is the formalism which will be used in the rest of this thesis, firstly because

it will work even for non-Hamiltonian dynamics, and secondly because it greatly simplifies the derivation of many important properties of MD algorithms. For example, the stationary distribution for equations of motion described by the Liouvillian \mathcal{L} is necessarily a solution to $\mathcal{L}\rho = 0$ (Note, however, that this is not a sufficient condition for ρ to be the correct stationary distribution, see ref. [98] for more details).

For the current discussion, the exact form of \mathcal{L} is not important, but will, for completeness, be included any time the explicit form of an equation of motion is given. For most degrees of freedom, x say, this operator will be given by

$$\mathcal{L}_x = \frac{\partial(\dot{x} \cdot)}{\partial x} \quad (3.2.2)$$

where the \cdot is used as a placeholder, such that

$$\mathcal{L}_x \rho = \frac{\partial(\dot{x} \rho)}{\partial x} \quad (3.2.3)$$

However, stochastic equations of motion are governed by Itô calculus and so are more complicated. A full discussion of this in the case of Markovian stochastic dynamics is given in the books [89] and [100], and a summary of the argument needed to derive results important to this thesis is given in appendix A.

3.3 Trotter splitting

Equation 3.2.1 has a formal solution of the form

$$\rho(\mathbf{r}, \mathbf{p}, t + \Delta t) = \exp(-\Delta t \mathcal{L}) \rho(\mathbf{r}, \mathbf{p}, t) \quad (3.3.1)$$

This formal solution cannot be exactly solved for most problems of interest. However, let us suppose that \mathcal{L} is the sum of \mathcal{N} Liouvillian operators, the equations of

motion of each of which can be propagated analytically. An approximation to the correct dynamics, known as a Trotter splitting [54, 55, 101], can then be obtained by using

$$\exp(-\Delta t \mathcal{L}) = \prod_{i=\mathcal{N}}^1 \exp\left(-\frac{1}{2}\Delta t \mathcal{L}_i\right) \prod_{j=1}^{\mathcal{N}} \exp\left(-\frac{1}{2}\Delta t \mathcal{L}_j\right) + \mathcal{O}(\Delta t^3) \quad (3.3.2)$$

where the two products are in the opposite order to each other to take into account the non-commutativity of the different Liouvillian operators.

Neglecting the term of $\mathcal{O}(\Delta t^3)$, the effect of the operator product on the right hand side of this equation can be calculated analytically. This is done by propagating for half a time step the equations of motion corresponding to the action of each individual Liouvillian, and then repeating this for another half a time step in reverse order.

This is a powerful method for generating new dynamical algorithms, not only because it can be applied easily to any set of equations of motion, but also because this type of splitting has many desirable properties, such as time reversibility.

3.4 Path integral Langevin equation dynamics

To take an illustrative example of this methodology, let us go back to the problem of non-ergodic behaviour in microcanonical PIMD simulations. The method of choice to solve this is to instead sample the canonical ensemble using a thermostating algorithm. These induce energy transfer between the different modes of the system via an external heat bath, and so can rapidly equilibrate all of the degrees of freedom.

Here we discuss a simple, computationally efficient method for canonical sampling which uses all the results of this chapter, the path integral Langevin equation (PILE) thermostat, as derived in ref. [45].

To build this thermostat, the first thing that is required is to choose the equations

of motion. In section 3.1 it was shown that the trajectory of a free ring polymer can be integrated analytically. The dynamics can therefore be split into three exactly soluble components; the motion of the free ring polymer, the action of the Langevin equation thermostat, and the application of the interatomic forces.

As per the rationale of section 3.1, the action of the forces is applied in the bead representation and the propagation of the free ring polymer equations of motion in the normal mode representation. The Langevin algorithm will also be integrated in the normal mode representation, for reasons which will be explained later.

Therefore, associating $\mathcal{L}_f^{(jk)}$ with the evolution of each normal mode of the free ring polymer, $\mathcal{L}_\gamma^{(jk)}$ with the interaction between each normal mode and the Langevin thermostat and $\mathcal{L}_U^{(ij)}$ with the action of the interatomic forces on each bead, the equations of motion corresponding to each operator are

$$\mathcal{L}_U^{(ij)} : \dot{\mathbf{p}}_j^{(i)} = \mathbf{f}_j^{(i)} \quad (3.4.1)$$

$$\mathcal{L}_f^{(jk)} : \begin{cases} \dot{\mathbf{p}}_j^{(k)} = -m_j \omega_k^2 \tilde{\mathbf{r}}_j^{(k)} \\ \dot{\tilde{\mathbf{r}}}_j^{(k)} = \frac{1}{m_j} \tilde{\mathbf{p}}_j^{(k)} \end{cases} \quad (3.4.2)$$

$$\mathcal{L}_\gamma^{(jk)} : \dot{\mathbf{p}}_j^{(k)} = \sqrt{\frac{2m_j \gamma_k}{\beta_n}} \boldsymbol{\xi}_j^{(k)} - \gamma_k \tilde{\mathbf{p}}_j^{(k)} \quad (3.4.3)$$

Using the above notation for the individual Liouvillian operators, the total Liouvillian operator is given by

$$\mathcal{L} = \sum_{j=1}^N \sum_{i=0}^{n-1} \mathcal{L}_U^{(ij)} + \sum_{j=1}^N \sum_{k=0}^{n-1} \mathcal{L}_f^{(jk)} + \sum_{j=1}^N \sum_{k=0}^{n-1} \mathcal{L}_\gamma^{(jk)} \quad (3.4.4)$$

$$= \mathcal{L}_U + \mathcal{L}_f + \mathcal{L}_\gamma \quad (3.4.5)$$

where [44]

$$\mathcal{L}_U^{(ij)} = \mathbf{f}_j^{(i)} \frac{\partial \cdot}{\partial \mathbf{p}_j^{(i)}} \quad (3.4.6)$$

$$\mathcal{L}_f^{(jk)} = \frac{\tilde{\mathbf{p}}_j^{(k)}}{m_j} \frac{\partial \cdot}{\partial \tilde{\mathbf{r}}_j^{(k)}} - m_j \omega_k^2 \tilde{\mathbf{r}}_j^{(k)} \frac{\partial \cdot}{\partial \tilde{\mathbf{p}}_j^{(k)}} \quad (3.4.7)$$

$$\mathcal{L}_\gamma^{(jk)} = -\gamma_k \frac{\partial \left(\tilde{\mathbf{p}}_j^{(k)} \cdot \right)}{\partial \tilde{\mathbf{p}}_j^{(k)}} - \left(\frac{m_j \gamma_k}{\beta_n} \right) \frac{\partial^2 \cdot}{\left(\partial \tilde{\mathbf{p}}_j^{(k)} \right)^2} \quad (3.4.8)$$

As an aside, note that this Liouvillian is consistent with canonical sampling as it gives the correct stationary distribution, i.e.

$$\mathcal{L} \exp(-\beta_n H_n) = 0 \quad (3.4.9)$$

Multiple Trotter splittings can be applied to this set of Liouvillian operators. We will follow ref. [45] and use

$$e^{-\Delta t \mathcal{L}} \approx e^{-(\Delta t/2) \mathcal{L}_\gamma} e^{-(\Delta t/2) \mathcal{L}_U} e^{-\Delta t \mathcal{L}_f} e^{-(\Delta t/2) \mathcal{L}_U} e^{-(\Delta t/2) \mathcal{L}_\gamma} \quad (3.4.10)$$

This particular splitting has the advantage that the canonical ensemble is stationary with respect to both \mathcal{L}_γ and $\mathcal{L}_U + \mathcal{L}_f$, and so the only sampling errors are in splitting \mathcal{L}_U and \mathcal{L}_f . This avoids the problem commonly associated with Langevin dynamics that the accuracy of the integration algorithm degrades as the strength of the friction increases. In fact, it can be shown [88] that this integrator will work even in the infinite friction limit.

Combining this Trotter splitting with the finite time solutions to the differential equations 3.4.1, 3.4.2 and 3.4.3 (see appendix B for their derivation) gives the

integration scheme

$$\begin{aligned}
\tilde{\mathbf{p}}_j^{(k)} &\leftarrow \exp\left(-\frac{\Delta t}{2}\gamma_k\right)\tilde{\mathbf{p}}_j^{(k)} + \boldsymbol{\xi}_j^{(k)}\sqrt{\frac{m_j}{\beta_n}(1-\exp(-\gamma_k\Delta t))} \\
\mathbf{p}_j^{(i)} &\leftarrow \mathbf{p}_j^{(i)} + \frac{\Delta t}{2}\mathbf{f}_j^{(i)} \\
\begin{pmatrix} \tilde{\mathbf{p}}_j^{(k)} \\ \tilde{\mathbf{r}}_j^{(k)} \end{pmatrix} &\leftarrow \begin{pmatrix} \cos(\omega_k\Delta t) & -m_j\omega_k\sin(\omega_k\Delta t) \\ \frac{1}{m_j\omega_k}\sin(\omega_k\Delta t) & \cos(\omega_k\Delta t) \end{pmatrix} \begin{pmatrix} \tilde{\mathbf{p}}_j^{(k)} \\ \tilde{\mathbf{r}}_j^{(k)} \end{pmatrix} \\
\mathbf{p}_j^{(i)} &\leftarrow \mathbf{p}_j^{(i)} + \frac{\Delta t}{2}\mathbf{f}_j^{(i)} \\
\tilde{\mathbf{p}}_j^{(k)} &\leftarrow \exp\left(-\frac{\Delta t}{2}\gamma_k\right)\tilde{\mathbf{p}}_j^{(k)} + \boldsymbol{\xi}_j^{(k)}\sqrt{\frac{m_j}{\beta_n}(1-\exp(-\gamma_k\Delta t))}
\end{aligned} \tag{3.4.11}$$

The last thing that must be done is to decide on the values of the thermostat parameters γ_k . Given that the rationale behind PILE is to rapidly repartition the energy between the ring polymer normal modes, a sensible choice of the parameters is the one which minimizes the autocorrelation time of the normal mode Hamiltonian. The most obvious way to define this Hamiltonian is the contribution of each normal mode to equation 3.1.1, given by

$$H_j^{(k)} = \frac{|\tilde{\mathbf{p}}_j^{(k)}|^2}{2m_j} + m_j\omega_k^2|\tilde{\mathbf{r}}_j^{(k)}|^2 \tag{3.4.12}$$

The quantity that we want to minimize is then

$$\tau_H = \frac{\int_0^\infty \langle (H_j^{(k)}(0) - \langle H_j^{(k)} \rangle) (H_j^{(k)}(t) - \langle H_j^{(k)} \rangle) \rangle dt}{\langle (H_j^{(k)})^2 \rangle - \langle H_j^{(k)} \rangle^2} \tag{3.4.13}$$

This can be calculated exactly for the free ring polymer attached to a Langevin thermostat [92, 100], with the optimal parameters [45] for the internal normal modes being given by $\gamma_k = 2\omega_k$. This only leaves the centroid ($k = 0$) thermostat unspecified,

the optimal choice of this parameter depending on the time scale of molecular motion.

While this is only formally correct in the free ring polymer limit, this is still likely to be close to optimal for realistic condensed phase systems. In fact, the large separation in energy scales which causes the non-ergodicity problems also means that there will be little deviation from the free ring polymer limit for the high frequency modes. Conversely, if there is a large deviation from the free ring polymer limit then there will also be a large coupling between the intermolecular potential and the internal modes of the ring polymer, so in either case the dynamics should be ergodic.

This then is a fully general method for running canonical PIMD simulations which avoids the problems associated with the high frequency ring polymer motion. What is more, the methodology used in deriving this integrator makes it very easy to change the algorithms used. For example, if we wished to use PIGLET rather than a white-noise Langevin equation thermostat then we would just replace equation 3.4.3 with equation 2.2.2 and repeat the above analysis.

This framework will be used in the creation of a novel integrator for constant pressure dynamics in chapter 5.

Chapter 4

i-PI

4.1 Electronic structure

So far, we have discussed the theoretical framework in which atomistic modelling of bulk systems can be used to determine their properties. In particular, chapter 2 discussed how fundamental quantities in the quantum statistical mechanics of distinguishable particles such as the partition function and an ensemble average can be reformulated in terms of a classical phase space, and chapter 3 showed how integration algorithms can be derived for any given ensemble.

However, up until now we have glossed over one of the most important parts of any PIMD simulation: the calculation of the forces $\mathbf{f}_j^{(i)}$. These were defined in equation 2.1.15 as position derivatives of the interaction potential, but up until now the form of this potential has been left unspecified.

The Hamiltonian of the system can be represented as the sum of the kinetic energy of each nucleus and electron and the various Coulomb interactions between

them, using

$$\begin{aligned} \hat{H} = & - \sum_{J=1}^{N_e} \frac{\hbar^2}{2m_e} \nabla_J^2 - \sum_{j=1}^N \frac{\hbar^2}{2m_j} \nabla_j^2 \\ & + \frac{e^2}{4\pi\epsilon_0} \left(\sum_{J<K}^{N_e} \frac{1}{|\hat{\mathbf{R}}_J - \hat{\mathbf{R}}_K|} - \sum_{j,K}^{N,N_e} \frac{Z_j}{|\hat{\mathbf{r}}_j - \hat{\mathbf{R}}_K|} + \sum_{j<k}^N \frac{Z_j Z_k}{|\hat{\mathbf{r}}_j - \hat{\mathbf{r}}_k|} \right) \end{aligned} \quad (4.1.1)$$

where N_e is the total number of electrons in the system, Z_j is the charge on the j -th nucleus and $\hat{\mathbf{R}}_J$ is the position of the J -th electron.

In theory, it is possible to find the total energy and the combined nucleus-electron wavefunction by finding the eigenvalues and eigenstates of this operator directly. However, we intuitively expect the motion of the electrons and the nuclei to be effectively decoupled from each other because their masses differ by several orders of magnitude. Therefore, it is generally both highly accurate and computationally efficient to assume that the total wavefunction can be separated into a product of a nuclear and an electronic wavefunction.

This is the theory behind the Born-Oppenheimer approximation [50], which essentially splits the problem into two: firstly, calculating the electronic structure of a system in which the nuclear positions are fixed, and secondly, studying the motion of the nuclei on the ground adiabatic potential energy surface created by their interaction with the electrons and each other.

The second of these is exactly given by equation 2.1.1, and has been the focus of this thesis so far. The first is equally simple; if the electronic wavefunction is given by $|\Phi\rangle$, then the ground adiabatic potential is calculated by solving

$$\left(\hat{H} - \hat{K}_{\text{nuc}} \right) |\Phi\rangle = U(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N) |\Phi\rangle \quad (4.1.2)$$

where \hat{K}_{nuc} is the nuclear kinetic energy, and on both sides of this equation each

of the nuclear positions \mathbf{r}_i are held constant. The $U(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N)$ calculated by this equation defines both the Hamiltonian originally given in equation 2.1.2 and the forces which are used to propagate the nuclear dynamics.

However, this is generally a much more challenging problem than the nuclear Schrödinger equation given in equation 2.1.1, since quantum effects are much more important for electrons than for nuclei. In particular, here we cannot ignore identical particle exchange or correlation, and so in almost all cases an exact calculation is far beyond the reach of even the most powerful computers. As such, further approximations usually have to be made.

4.2 Empirical potentials

In fact, the cheapest method for calculating the forces is not to solve the Schrödinger equation, but instead to represent $U(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N)$ using a physically motivated functional form and fit the results of PIMD simulations obtained with it either to accurate electronic structure calculations or experimental data. Empirical potentials like this often split the interaction potential into components that only depend on a small number of nuclear positions, such as bond lengths, bond angles and partial charge sites. This usually leaves a small number of parameters, such as bond strengths, to be adjusted to give the best fit to the chosen set of data.

For example, one of the earliest and most popular empirical potential models for liquid water is the SPC/E model of Berendsen *et al.* [102]. This treats each molecule as rigid, with fixed OH bond lengths and HOH bond angles, and so ignores intramolecular motion. The interaction potential is approximated using Coulomb attractions between partial charge sites on the H and O atoms, and a Lennard-Jones interaction [103] between O atoms on neighbouring molecules.

This only requires us to determine 5 parameters: the equilibrium bond length

and bond angle, the two Lennard-Jones parameters and the partial charge on each hydrogen (the oxygen partial charge being fixed by requiring the molecule to be neutral overall). The molecular geometry was fixed using a tetrahedral bond angle and a bond length of 1 Å. The remainder of the model was parameterized to reproduce the heat of vaporization and the pressure at 300 K and at the experimental density, and has been shown to give reasonable values of properties such as the diffusion constant.

However, the simplicity of SPC/E comes at the price of reduced applicability. Firstly, this model of water is only reliable for a small range of state points. In particular, it totally fails to reproduce the pressure-temperature phase diagram of ice [104], and so in general we can only trust results obtained with a SPC/E model if they pertain to liquid water.

Many of the qualitative features of this phase diagram can be obtained by moving the negative point charge away from the oxygen nucleus and closer to the hydrogens, as done for models such as TIP4P/2005 [105]. However, this comes at the cost of adding another parameter to the functional form of $U(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N)$, making fitting the model more difficult.

This problem of models becoming more complex as they become more detailed is then exacerbated when considering a second, even greater weakness of empirical potentials: that any simple functional form for the potential energy will be unable to describe certain kinds of phenomena, even in principle.

For example, since SPC/E is a rigid model it can never be used to study molecular vibrations. While the extension to flexible bonds is straightforward [106], this again comes at the cost of increasing the number of parameters required. Similarly, even this improved model only uses point charges, and so cannot be used for any calculation that involves significant charge redistribution, such as capturing the small

difference between the melting point and the temperature of maximum density [107]. This would require a model that involves explicit electronic polarizability, such as the TTM2.1-F model of Fanourgakis and Xantheas [108], which would again increase the computational expense and the difficulty of both fitting the parameters and implementing the dynamics.

In the case of liquid water, for which there has been decades of research into developing new empirical potential models, it is usually possible to study any given property without having to parameterize a model from scratch. Recent research into nuclear quantum effects in water, for example, has used empirical models fitted to MD simulations as a basis for ones designed specifically for PIMD [6, 109].

However, there are very large areas of condensed phase physics for which this will not be true, either because they involve a system at a state point which has not previously been studied, or because previous models were too simplistic to simulate the desired physics. On the other hand, developing a completely new empirical potential can be incredibly time-consuming: the space of possible functional forms grows exponentially with the number of parameters in $U(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N)$, the computational cost of the simulations needed to test each individual choice of parameterization can be prohibitively high, especially if explicit electronic polarization is included, and the properties used to inform the fit can often depend in a highly non-linear way on the parameters.

This suggests that for problems whose correct description requires a large number of parameters or a complicated model, the cost of developing an empirical potential will outweigh the effort saved by not solving equation 4.1.2 directly. In this case, a compromise needs to be made between a full solution of the Schrödinger equation, which is too expensive, and the use of an empirical potential, which would be too difficult to develop.

4.3 Density functional theory

In applications to MD, the most commonly used approximation to equation 4.1.2 is density functional theory (DFT). Instead of using the electronic wavefunction directly, in this method the system is described by its electron density, ρ_e , which is defined as

$$\rho_e(\hat{\mathbf{R}}) = N_e \int d\hat{\mathbf{R}}_2 \dots \int d\hat{\mathbf{R}}_{N_e} \left| \Phi(\hat{\mathbf{R}}, \hat{\mathbf{R}}_2, \dots, \hat{\mathbf{R}}_{N_e}) \right|^2 \quad (4.3.1)$$

Despite being a much lower-dimensional object than the full electronic wavefunction, it can be shown that it contains all the information necessary to determine the properties of the system [110, 111]. The simplicity of this formulation gives it the important advantage that formally it scales linearly with the number of electrons in the system [112, 113]. While most implementations of DFT do not achieve this in practice, they are still generally fast, making this the only *ab initio* electronic structure technique which can currently be used in large-scale MD simulations (the exception of ref. [114] notwithstanding).

For reasons of accuracy most DFT codes do not work with the electron density directly, and instead represent it with a set of one-electron wavefunctions [115, 116] using

$$\rho_e(\hat{\mathbf{R}}) = \sum_{J=1}^{N_e} \left| \phi_J(\hat{\mathbf{R}}) \right|^2 \quad (4.3.2)$$

Once this substitution is made, the problem of finding $U(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N)$ then reduces to solving a set of secular equations for which these one-electron wavefunctions (known as the Kohn-Sham orbitals) are the eigenstates, given by [117]

$$\left(-\frac{1}{2}\nabla^2 + u_N[\rho_e] + u_C[\rho_e] + u_{xc}[\rho_e] \right) \phi_J(\hat{\mathbf{R}}) = \epsilon_J \phi_J(\hat{\mathbf{R}}) \quad (4.3.3)$$

where ϵ_J is the Kohn-Sham orbital energy, $u_N[\rho_e]$ and $u_C[\rho_e]$ are the fields caused by

electron-nuclear attraction and the classical Coulomb electron repulsion respectively, and $u_{xc}[\rho_e]$ contains both the effect of identical-particle exchange, and that of electron correlation on both the kinetic energy and the Coulomb interaction.

The functions ϕ_J are found using a self-consistent procedure: firstly an initial guess is made for ϕ_J , secondly this is used to form ρ_e and thirdly a new set of ϕ_J is derived by solving the set of equations 4.3.3. Steps two and three are then repeated until the wavefunctions from one step to another agree to within a pre-set error bound. $U(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N)$ can then be calculated from the converged results using

$$U(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N) = \sum_{J=1}^{N_e} \left\langle \phi_J \left| -\frac{1}{2} \nabla^2 \right| \phi_J \right\rangle + E_{xc}[\rho_e] + \int d\hat{\mathbf{R}} \rho_e(\hat{\mathbf{R}}) u_N[\rho_e] + \frac{1}{2} \int d\hat{\mathbf{R}} \rho_e(\hat{\mathbf{R}}) u_C[\rho_e] \quad (4.3.4)$$

where E_{xc} is related to u_{xc} by

$$u_{xc}[\rho_e] = \frac{\partial E_{xc}[\rho_e]}{\partial \rho_e} \quad (4.3.5)$$

Note that this energy is an entirely local function of ρ_e , and so can take electron correlation into account approximately without having to converge the cusp that appears in $|\Phi\rangle$ when two electrons occupy the same space, unlike wavefunction-based methods. This has the consequence that the Kohn-Sham orbitals tend to converge very quickly with respect to basis set size, further reducing the computational expense of DFT.

While in theory this is an exact method, in practice the correct form of $E_{xc}[\rho_e]$ is not known and so it must be approximated, either from certain known analytical bounds [118–120] or from fitting to more accurate electronic structure calculations [121–123]. Fortunately, this is generally by far the smallest energy term, and many

common functional forms have been developed which can determine the physical properties of systems of chemical interest to high levels of accuracy. For example, it has been shown [124] that this method can be used to calculate very challenging properties such as hydrogen bond strength in small water clusters to within 5-10 %.

A vast literature exists on this subject, describing both the implementation of efficient algorithms and the development of improved exchange-correlation functionals which can be applied to accurately calculate a greater range of properties. As this work is not the focus of this thesis, we refer any interested reader to review articles such as refs [115, 125–129] for a more detailed discussion.

One more important discussion point is how to calculate the forces. Theoretically, in cases where the basis set is constant with respect to the nuclear positions, we can calculate the forces with only a knowledge of the converged Kohn-Sham orbitals for the current nuclear geometry using the Hellmann-Feynman theorem [130]. This says that the forces can be derived from the already converged Kohn-Sham orbitals and the known analytical derivatives of the fields in equation 4.3.3.

However, in most cases this method is numerically unstable, since it is very sensitive to the level of convergence of the ground-state density. Furthermore, in many cases the natural basis set to use are hydrogenic orbitals centred at the nuclear positions, which therefore depend explicitly on the nuclear positions.

Fortunately, it can be shown that the “Pulay forces” generated from the change of the basis set with respect to nuclear positions can be taken into account in DFT analytically without having to do more self-consistent loops, by taking an appropriate definition of the derivative of the density matrix with respect to the nuclear coordinates [131, 132]. Accurate forces can therefore be calculated without a large increase in the computational time, making it possible to run MD simulations using DFT.

4.4 Computational challenges

At this point, we have all the ingredients that are needed to run *ab initio* path integral molecular dynamics (AI-PIMD) simulations. However, the real test of whether this is a practical method for studying condensed phase materials is whether or not these techniques can be brought together to be used for actual problems.

Density functional theory is a relatively mature field, and yet one that is still an area of active research [128]. The breadth of the subject means that writing code which can apply even a small subset of the techniques developed for DFT calculations is a huge challenge to implement (see for example the size of the author list of ref. [133]). Trying to write a DFT code from scratch is therefore not generally a practical proposition. Fortunately, there exist a wide range of programs which are either freely available or for which access can be bought. The problem then becomes how to integrate these with an implementation of path integral molecular dynamics.

The first possible strategy is to directly modify the DFT code so that it can run PIMD simulations. However, this is also not a perfect solution, for several reasons. Firstly, the size and complexity of most DFT libraries means it can be very difficult to make any major additions to their functionality, especially if they involve changing the way that the code is organized.

In particular, adapting a code designed specifically for electronic structure calculations so that it can do MD simulations, or one which can run MD to do PIMD, would involve rewriting its main loop. Even if we neglect all the extra functions needed to run the PIMD algorithms, this would require reworking many of the input and output routines and a major change to the structure of the code.

Also, since most electronic structure codes are designed for parallel execution, a large number of arrays of numerical data must be kept up-to-date across all processors. During a MD simulation, where the nuclear positions, and perhaps even the

simulation box, will be repeatedly updated, this can be non-trivial.

Even those codes which do have a working MD/PIMD implementation tend not to use the latest algorithms, and so can be both severely limited in functionality and much less efficient than ones specifically designed for PIMD. Since AI-PIMD simulations are much slower than those with empirical potentials even with the current state-of-the-art techniques, this can make studying some properties too computationally expensive to be feasible.

Finally, as mentioned earlier, there are many different electronic structure codes in use, each with their own strengths and weaknesses. Therefore, to get the best results, different DFT codes should be used for different problems, and as such the effort of implementing PIMD routines may have to be repeated. Even if one research group only needs one DFT library, others may not, and so directly modifying DFT codes for use with PIMD is not a practical method for making AI-PIMD simulations available to the wider community.

However, if we consider that we only need the DFT code to calculate the forces and potential energy, a better strategy is to write the code which propagates the PIMD algorithms separately to the electronic structure routines. This avoids the need to make major modifications to the electronic structure libraries, or to have a different PIMD implementation for every different DFT code. The only remaining difficulty is then how to design the communication interface between the two programs.

The simplest method to use would be a scripting approach, where at each time step the PIMD program calls an external DFT code using the current nuclear configuration as input, and searches the output files for the potential and forces once it is finished. However, this would not be very efficient, because a large proportion of the time spent in the DFT code would be on initialization, which would have to be repeated every time it was called. Also, many DFT routines contain efficient extrapolation

techniques (see for example ref. [134]) which can rapidly converge the ground state wavefunctions from those of the previous steps; these cannot be used if the program is rerun every time.

To avoid these problems, our approach is to have the PIMD routines running alongside one or more instances of the DFT program, and to dynamically send the required information between them. Since the DFT code does not have to be restarted at any point, it need only be initialized once and can easily use any wavefunction extrapolation routines available. Our implementation of this, i-PI [44], is discussed in the next section.

4.5 Program goals and language specification

Before introducing i-PI, let us first remind ourselves of the reasoning behind its development and the characteristics we would therefore like it to have, so that we can motivate our choice of implementation. As discussed in the introduction, one of the primary aims of this project is to develop computer software to help make AI-PIMD simulations more widespread. As such, there are several important attributes that i-PI is required to have.

Firstly, we need i-PI to be easy to use, even by people not familiar with the path integral formalism. This is because, while the use of MD techniques is fairly widespread, there are relatively few researchers who are studying NQE specifically. However, as computers become ever more powerful many more researchers are beginning to consider how quantum effects may influence the phenomena they are interested in. Similarly, people who have previously included NQE with more approximate methods have begun to run PIMD simulations instead as it has become feasible to complete them in a reasonable time frame. As such, the structure and implementation of the code should be as transparent as possible, so as to accommodate these

potential users.

Secondly, as mentioned earlier, many groups favour one particular choice from the many different available DFT programs in their research, and are unlikely to be willing to learn to use a new code just to run simulations with i-PI. As such, it must be possible to adapt any set of DFT routines to work with i-PI with as small an implementation effort as possible.

Thirdly, as path integral methods are still a topic of active research, i-PI should be as simple as possible to modify and extend, so that it can keep up with the latest developments in the field. This is again much simpler if the code is clear and concise, and also suggests that using a modular programming style, where any particular algorithm or technique can be easily replaced by a different one, would be highly advantageous.

Finally, we note that one of the things that is *not* important to our particular problem is computational efficiency. This is because the electronic structure calculation will take up the lion's share of the running time, even if the most efficient modern DFT algorithms are used. As such, in any case where there is a trade-off between clarity and speed, we will always opt for the method which gives the most user-friendly code.

In order to achieve these goals we have chosen to write i-PI using the Python programming language, for multiple reasons. Firstly, it is very easy to learn. The language has an intuitive syntax (particularly when dealing with iterators over lists), has an excellent interpreter which allows code to be tested without having to compile it, and has extensive online documentation. This should reduce the barrier to anyone wishing to develop the code and give it new functionality.

Secondly, Python files are generally easier to read than those of similar languages. As well as having an easy to understand syntax, the language specifications enforce

a layout which has large amounts of whitespace and indentation, preventing some of the poor habits which can lead to obfuscated code. Readability is therefore one of the fundamental advantages of the Python programming language, which will make it easier for both users and developers to understand how i-PI works.

These things make it a very good choice for an open-source code with contributions from multiple collaborators from different groups. However, it also has features that make it well suited to the problem at hand. Like many modern programming languages, it is object-oriented, which means that as well as the standard data types, such as floating point numbers and character strings, the user can define their own classes containing related data and functions. In this way, we can group together all information related to one theoretical concept in a simple and transparent manner. This is the perfect programming paradigm to use if we wish to add new functionality to keep up with the research frontier, since replacing an existing algorithm simply involves creating a new class of object containing the appropriate code and using this in place of the original class.

On top of this, one of the major advantages of Python is that its standard library is very large, so there is a lot of functionality that has already been coded, tested and is regularly updated by members of the open-source community. As i-PI must have a clear input file syntax and be able to communicate with an *ab initio* electronic structure code in some way, it is particularly useful that Python has standard libraries to use for inter-process communication and reading multiple input file formats.

4.6 Data communication

With this decided, the remaining challenge is how to implement the communication between i-PI and the DFT code. Our approach is to use a client-server model in which i-PI acts as the server, dealing exclusively with evolving the nuclear degrees

of freedom according to the equations of motion. The evaluation of the forces, the potential and the virial is delegated to one or more instances of the client code (see figure 4.1).

The communication is kept to a minimum, not so much for its impact on performance as for the fact that exchanging more information would require a more substantial implementation effort on the client side. All the details of the force evaluation, such as the parameters of the electronic structure calculation, are left to the input of the external code: i-PI only sends the coordinates of the nuclei and the dimensions of the simulation box, \mathbf{h} , and expects the atomic species to be stored internally by the client code in the same order they are in the i-PI input.

The client code computes and returns the forces, electronic potential energy and virial tensor, \mathbf{W} , possibly supplemented by a string that contains any further information on the electronic structure (e.g. the dipole moment, the partial charges of the atoms, etc.) that might be required for post-processing; i-PI simply outputs this string verbatim.




Figure 4.1: Schematic overview of the client-server model underlying i-PI. The communication is kept to a minimum, and so are the modifications that need to be made to adapt an existing electronic structure code to act as the client. The client is not restarted between successive force evaluations, so that the overhead associated with initialisation is avoided.

By making connections to multiple client codes, up to the number of beads n , this allows the trivial level of parallelism over the different replicas of the system to be fully exploited, since each different electronic structure calculation can be done simultaneously by a different client code. If n instances of the DFT code are used then this can give nearly perfect scaling with the number of beads, apart from the fact that each time step is limited by the slowest force calculation.

Instances of the *ab initio* code can register themselves dynamically, and i-PI maintains a list of active clients, with force evaluations being distributed among those that are currently available. This means that, if desired, one can run a n -bead calculation using only $m < n$ clients, as the server takes care of sending multiple replicas to each client per MD step. To avoid having clients idling for a substantial amount of time, m should be a divisor of n . The main advantage of this approach, compared to one that rigidly assigns one instance of the client to each bead, is that if each client is run as an independent job on a computational cluster, i-PI can start performing PIMD as soon as a single job has started, and can carry on advancing the simulation even if one of the clients becomes unresponsive.

On the other hand, there is an advantage to running simulations with $m = n$, which is particularly important for *ab initio* calculations. i-PI will always try to send the coordinates for one path integral replica to the client that computed it at the previous step: this reduces the change in the particle positions between force evaluations, so that the charge density/wavefunction from the previous step is a better starting guess and self-consistency can be achieved faster. Also, receiving coordinates that represent a continuous trajectory makes it possible to use any extrapolation strategies that might be available in the client code.

Obviously, most electronic-structure client codes provide a further level of parallelization, based on OpenMP and/or MPI. This is fully compatible with i-PI, as it

does not matter how the client does the calculation since only the forces, potential and virial are sent to the server, and the communication is typically performed by the master process of the client.

The client-server communication is implemented using sockets. A socket is a data transfer device that is designed for internet communication, so it supports both multiple client connections to the same server and two-way communication. This makes sockets ideal for use in i-PI, where each calculation may require multiple instances of the client code.

A socket interface can actually function in two different modes. UNIX-domain sockets are a mechanism for local, inter-process communication. They are fast, and best suited for when i-PI is run with empirical potentials, when the latency of the communication with the client becomes a significant overhead for the calculation. UNIX-domain sockets create a special file in the local file system that serves as a rendezvous point between server and clients, and are uniquely identified by the name of the file itself.

Unfortunately, UNIX sockets do not allow i-PI and the clients to be run on different computers, which greatly limits their utility when one needs to run massively parallel calculations. In these cases, as is typical when performing *ab initio* simulations, the force calculation becomes the bottleneck, so there is no need for fast communication with the server and so internet sockets, which are specifically designed for communication over a network, can be used instead.

This has the advantage that it allows the i-PI server to be held on a network where it can be run indefinitely, while the client codes can be run efficiently on a cluster that has been optimized for fast computation. In this way, the likelihood of the main loop crashing, which would involve the loss of data, is much reduced, without compromising on run time. As such, almost all AI-PIMD simulations should

be run using internet sockets.

Finally, let us consider how to implement the socket interface in each new DFT code. i-PI is distributed with a C interface which contains the functions required for socket initialization and data transfer. Most of the programming languages which have been used to create electronic structure programs, including C++, C and Fortran, can run these routines without requiring an interface, and so the only problem remaining is how to adapt the main loop such that the dynamics are done by i-PI.

In essence, this is done by finding the place in the code which represents the self-consistent loop, either in the main code or, if there are already working MD/PIMD routines, during the dynamical algorithms, and reading in the nuclear positions from i-PI just before it starts and returning the potential and forces just after it has finished.

Other than this, the only required changes are usually a slight modification of the initialization routines, so that the client socket is opened correctly, and the addition of a new input file option, so that the information used to identify the correct server can be read in. This information usually consists of a host name and socket number, possibly supplemented by a string asking for either a UNIX or an internet socket.

So far, this procedure has been performed on the electronic structure codes FHI-aims [135], CPMD [136], Quantum Espresso [133], CP2K [137], and the empirical potential code LAMMPS [138]. Patch files which can be applied to the standard distributions of the last three of these are supplied with i-PI.

4.7 Program functionality and implementation

Before discussing the capabilities of i-PI, first let us consider in general terms how the code has been designed. i-PI has a hierarchical structure, where the different steps of the calculation correspond to different levels of the code organization. To realize this,

each abstract concept used in the algorithms has its own dedicated Python class.

For example, let us consider the algorithms associated with the Trotter splitting of the NVT ensemble Liouvillian described in section 3.4. There is one step associated with heat transfer to the thermostat, one for the influence of the electronic interaction potential, and one propagating the internal ring polymer motion.

This would be represented in the code by an “ensemble” class, which then delegated propagation of the dynamics to a “thermostat” class, a “forces” class and a “normal modes” class. The choice of “ensemble” class determines both the stationary ensemble and how the equations of motion are discretized, while the other classes determine the particular choice of integration algorithms. For example, if the user wished to use a PIGLET thermostat instead of a PILE thermostat then in the code this would correspond to choosing a different “thermostat” class.

This is what underlies the flexibility of the i-PI code, as each item of functionality is designed to be easily replaceable. If an improved algorithm for any stage of the calculation is developed, then in many cases i-PI can be updated to include it by simply writing one new module with its implementation and making a few amendments to the relevant input routines.

With this in mind, let us look in more detail at what has currently been implemented. As per its stated aims, i-PI incorporates many of the most recent developments in path integral and coloured-noise MD, so as to most effectively make AI-PIMD simulations more widely practicable. In summary, the list of functionality available in the current version of i-PI is:

- MD/PIMD in the NVE , NVT and NPT ensembles, with the high-frequency internal vibrations of the ring polymer propagated in the normal-mode representation [45];
- ring polymer contraction [20, 68], used to speed up calculations with empirical

potentials;

- efficient stochastic velocity rescaling [86] and PILE thermostats [45];
- various GLE thermostats, including the optimal sampling [78, 79], quantum [90], and δ [91] thermostats, the parameters for which can be downloaded from an on-line repository [139];
- mixed path integral–generalized Langevin equation techniques for accelerated convergence, including both PI+GLE [21] and its more recent and effective version, PIGLET [59];
- all the standard estimators for structural properties, the quantum kinetic energy, pressure, etc.;
- more sophisticated estimators such as the scaled-coordinate heat capacity estimator [64], estimators to obtain isotope fractionation free energies by reweighting a simulation of the most abundant isotope [140], and a displaced-path estimator for the particle momentum distribution [141];
- the infrastructure that is needed to perform ring polymer molecular dynamics (RPMD) [24, 25], including the latest thermostatted version T-RPMD [26], and partially-adiabatic centroid molecular dynamics (PA-CMD) [22, 23] approximate quantum dynamics calculations.

To elaborate on this, let us discuss the more important functionality in more detail. Most of the capabilities of the code broadly fall into two categories: techniques to speed up the convergence with respect to the number of replicas, and estimators or methods which allow the user to calculate properties of interest.

4.7.1 Ring polymer contraction

Currently, two techniques have been implemented to reduce the number of beads required for PIMD simulations: ring polymer contraction and generalized Langevin equations. The latter has already been discussed at length in section 2.2, however the former has only been mentioned in passing and so we will now describe this in more detail.

Ring polymer contraction has been developed from the same principle used for multiple time-step methods, i.e. that the long-range part of the potential, which tends to require a large amount of time to compute as there is a greater number of interactions to consider than for short-range potentials, also tends to be slowly varying over a wide length and time scale. As such, the long-range part of the potential can be calculated with fewer beads than the short-range part, since a smaller portion of the distribution needs to be sampled to achieve the same accuracy.

This can be achieved by using the normal mode formalism given in section 3.1; mapping a ring polymer with n beads to one with n' beads is done by selecting the n' lowest frequency normal modes of the n bead ring polymer, and then multiplying them by a normalization factor so that the centroid is unchanged by the transformation. The long-range forces are then calculated on the reduced ring polymer, whereas the short-range forces are calculated using the full number of beads.

The only remaining step is to rescale the resultant force vector and virial tensor so that they correspond to the original number of beads. Let us suppose that the original ring polymer is $\tilde{\mathbf{r}}$ and the reduced ring polymer is $\tilde{\mathbf{r}}'$, and that they can be interconverted using

$$\tilde{\mathbf{r}}' = \sqrt{\frac{n'}{n}} \mathbf{\Gamma} \tilde{\mathbf{r}} \quad (4.7.1)$$

where $\mathbf{\Gamma}$ is a $n' \times n$ matrix, where the left-most $n' \times n'$ portion is an identity matrix and all other elements are zero. In essence this is a truncation matrix, which selects

the appropriate normal modes to use in the contracted ring polymer.

With this definition, the force vector and each component of the virial tensor can be rescaled from n' beads to n beads using

$$\tilde{\mathbf{f}} = \sqrt{\frac{n}{n'}} \mathbf{\Gamma}^T \tilde{\mathbf{f}}' \quad (4.7.2)$$

Having more than one force field requires there to be two different socket interfaces, one dealing with the calculation of the long-range component of the potential and one the short-range. This gives an example of the power of the object-oriented approach, since in the code this is implemented by taking a three-tier hierarchy for the force calculation: the lowest giving the socket interface for each bead, the middle calculating each component of the force, and the highest consolidating all of the data. Two different force components can then be included by simply using two different instances of the middle class.

4.7.2 Reweighted sampling estimators

A large number of estimators are included in the standard version of i-PI, most of which are fairly standard and should require no further explanation. However, some of the more specialized estimators are more complicated, and so the implementation of the necessary algorithms required some careful attention.

One of the more common advanced theoretical techniques is reweighted sampling, which takes the results of a simulation at one state point and uses them to investigate either a system with adjusted physical parameters or a different ensemble. These estimators, instead of taking averages of the sort given in equation 2.1.7, are formed by

$$\langle \hat{O} \rangle' = \lim_{n \rightarrow \infty} \left(\frac{\langle O_n(\mathbf{r}) \exp(-\beta_n \Delta U) \rangle^{\text{PI}}}{\langle \exp(-\beta_n \Delta U) \rangle^{\text{PI}}} \right) \quad (4.7.3)$$

where the prime in the first average indicates that it is an average over a different ensemble to the one in which the dynamics is propagated, and ΔU is the energy difference between these two ensembles.

This formulation is used to calculate isotope fractionation ratios [140] and the displaced-path proton momentum distribution estimator [141], the latter of which will be discussed later in chapter 6. We will therefore leave a further discussion about the benefits and weaknesses of the reweighting method until then.

For the current discussion, the most important implementation detail is that, for the displaced-path momentum distribution estimator and one version of the isotope fractionation ratio, the reweighted Hamiltonian involves a different ring polymer than that used for the dynamics. The potential difference given in equation 4.7.3 will then require the potential of a different configuration to be calculated.

Again, the machinery used to calculate the forces make this simple. The property calculation routines define another instance of the top-level “forces” class, with the same forcefields and socket interface as the one used for the dynamics, but a different set of ring polymer positions. This is then used to determine the potential energy of this new configuration so that the estimator can be calculated without disturbing the dynamics.

Similarly, the specific heat capacity estimator of Yamamoto [64] is defined in terms of a derivative of the potential energy with respect to the temperature, and is calculated using a finite-difference method. This also requires the forces to be evaluated at configurations other than those generated by the dynamics, and so is dealt with in the code in the same way.

4.7.3 Dynamical property calculations

The last main class of properties which can be calculated by i-PI, included here for completeness despite not being the main focus of this thesis, are dynamical properties such as correlation functions. These describe the relation between the value of two observables at different times, and take the form

$$S_{AB}(t) = \frac{1}{Q} \text{Tr} \left[e^{-\beta \hat{H}} \hat{A} e^{it\hat{H}/\hbar} \hat{B} e^{-it\hat{H}/\hbar} \right] \quad (4.7.4)$$

where $S_{AB}(t)$ is the correlation function of the operators \hat{A} and \hat{B} .

These are more difficult to calculate than the static properties considered so far, since phase information and quantum coherence can have a large impact on the results, especially at long times. In the context of PIMD, this means that while the quantum statistics is calculated correctly, the fact that the dynamics are entirely classical and have no direct relation to the correct quantum dynamics embodied in the time propagator operator $e^{-it\hat{H}/\hbar}$ mean that this formalism is incapable of calculating quantities such as $S_{AB}(t)$ exactly. Therefore, in order to make progress with a path integral approach it will have to be approximated.

On the other hand, while quantum dynamical effects are very important in highly anharmonic non-dissipative systems [24], for the condensed phase systems we are interested in much of this information is lost, due to the multi-dimensional nature of the problem. Therefore, this motivates the use of a classical approximation to the dynamics, as long as the correct quantum statistics are maintained.

As it turns out, classical dynamics does not give the best approximation to the correlation function given in equation 4.7.4, but instead to the Kubo-transformed

version, $\tilde{S}_{AB}(t)$, given by

$$\tilde{S}_{AB}(t) = \frac{1}{\beta Q} \int_0^\beta d\lambda \text{Tr} \left[e^{-(\beta-\lambda)\hat{H}} \hat{A} e^{-\lambda\hat{H}} e^{it\hat{H}/\hbar} \hat{B} e^{-it\hat{H}/\hbar} \right] \quad (4.7.5)$$

This is because $\tilde{S}_{AB}(t)$ has many of the same time symmetries as classical correlation functions, both being real and even functions of t in the case where $\hat{A} = \hat{B}$, for example. As such, this is the form we will compare to when using a classical approximation to the dynamics. Fortunately, this Kubo-transformed correlation function can be related to the one given in equation 4.7.4 via a simple Fourier transform [24], and so it does not matter which of the two is used.

Many different semi-classical approximations to equation 4.7.5 exist, but the two that are implemented in i-PI are partially-adiabatic centroid molecular dynamics (PA-CMD) [142, 143] and ring polymer molecular dynamics (RPMD) [24].

PA-CMD uses an adjusted free ring polymer Hamiltonian to propagate the dynamics, given by

$$H_n^0 = \sum_{k=0}^{n-1} \sum_{j=1}^N \left(\frac{|\tilde{\mathbf{p}}_j^{(k)}|^2}{2m_j^{(k)}} + \frac{1}{2} m_j \omega_k^2 |\tilde{\mathbf{r}}_j^{(k)}|^2 \right) \quad (4.7.6)$$

where the fictitious masses $m_j^{(k)}$ have been chosen such that the ring polymer vibrational frequencies have been shifted out of the range of the motions of physical interest in the system.

The PA-CMD approximation to the Kubo-transformed correlation function, at least in the most common implementations, is then calculated using

$$\tilde{S}_{AB}(t) \approx \lim_{n \rightarrow \infty} \left(\langle A(\mathbf{r}^c, 0) B(\mathbf{r}^c, t) \rangle^{\text{PA-CMD}} \right) \quad (4.7.7)$$

where the superscript PA-CMD implies the use of equation 4.7.6 when propagating the dynamics.

RPMD is simpler as it uses the PIMD dynamics directly. In this scheme then, the Kubo-transformed correlation function is approximated using

$$\tilde{S}_{AB}(t) \approx \lim_{n \rightarrow \infty} \left(\langle A_n(\mathbf{r}, 0) B_n(\mathbf{r}, t) \rangle^{\text{PI}} \right) \quad (4.7.8)$$

In the original formulation, this average was done in the microcanonical ensemble, and thus had problems associated with the non-ergodicity of *NVE* PIMD dynamics [56]. T-RPMD instead uses canonical sampling of the internal ring polymer modes and microcanonical sampling of the centroid (as could be achieved by setting the variable γ_0 , as described in section 3.4, to zero), without affecting the validity of any of the analytical results derived for RPMD, and so is generally to be preferred.

Several comparisons of RPMD and PA-CMD have been discussed elsewhere [26, 142, 143], so we will only briefly summarize the main points. The original RPMD method performs very well when calculating reaction rates [144, 145], and has both been shown to be related to quantum instanton approaches to rate theory [146] and to rigorously reduce to quantum transition-state theory in the limit as $t \rightarrow 0_+$ [147–149]. This is related to its excellent short time behaviour, which is better than any of the related methods with different choices of the bead masses, as shown in ref. [150].

On the other hand, RPMD tends to show unphysical resonance behaviour when applied to calculate vibrational spectra [142, 143] or incoherent structure factors [151], due to coupling of the nuclear motion to the internal vibrations of the ring polymer. PA-CMD does not show this problem since averages are calculated on the centroid and so the ring polymer internal motion can be integrated out, so this method tends to perform better in this case. It also fares better than RPMD in strongly anharmonic problems [144], since the centroid potential of mean force on which the nuclear motion is propagated tends to be more harmonic than the underlying interaction potential, especially at low temperatures [152]. On the other hand, rescaling the

mass parameters causes its own unphysical artefacts such as the curvature problem [153, 154], which is due to the softening of the centroid potential of mean force by rotational-vibrational coupling.

T-RPMD, which damps the internal vibrational motion of the ring polymers, and so should have a much reduced resonance effect without having to introduce the curvature problem by rescaling the ring polymer normal masses, may be expected to outperform both these methods. Early results [26] appear to support this claim, although more research into the usefulness of T-RPMD for cases such as rate theory, where RPMD performs well, have yet to be completed.

In terms of implementation details, we first note that RPMD and T-RPMD require no additional functionality over that which is already needed for PIMD. PA-CMD was implemented by creating a custom class to deal with the adjusted masses of the ring polymer normal modes, and by making sure that all relevant estimators can be evaluated at the centroid coordinate as well as at the positions of the individual beads.

The only remaining piece of functionality that has not yet been discussed is the extension of the PIMD algorithms to generate constant pressure ensembles in a similar way to those used in section 3.4 to maintain constant temperature. Since our implementation of the NPT dynamics is novel, it will be discussed further in the next chapter.

Chapter 5

Constant pressure dynamics

5.1 Background

The majority of PIMD simulations are done using the NVT ensemble, since this forms a good compromise between computational efficiency and implementational complexity. However, most experiments are carried out under conditions of constant temperature and constant pressure, for which the NPT ensemble is appropriate. This suggests that constant pressure dynamics is the natural choice if we wish to compare simulations with experiments.

This is doubly true when considering that we do not propagate the dynamics using an exact solution to the Schrödinger equation, and instead simplify the calculation using either empirical potentials or approximate methods such as DFT. This means that the calculated density for a given temperature and pressure may be inconsistent with that found by experiment. Therefore, even when the experimental density is known we cannot assume that an NVT simulation will sample the correct ensemble [155].

When the experimental density is not known accurately, it becomes even more

difficult to do constant volume dynamics. It is possible to calculate optimized zero-temperature lattice constants; for example, performing an energy minimization at different volumes and fitting the results to a simple equation of state such as the Birch-Murnaghan equation [156], can be used to find the equilibrium volume [157]. However, in most systems of interest the change in the density due to thermal expansion will cause large errors in ensemble averages at finite temperature if these optimized lattice parameters are used.

This is particularly important when studying systems under extremes of pressure and/or temperature, where experimental studies are difficult to carry out. For example, there has been much research recently on the dynamics and properties of the interiors of the gas giant planets, where there are water/ammonia mixtures under pressures and temperatures of up to 600 GPa and 7000 K respectively [158]. These conditions are very difficult to produce experimentally, the only generally practical technique being shock pressure experiments [159], which are limited in that they can only produce particular state points [160]. On the other hand these are relatively simple to study with *NPT* AI-PIMD simulations [11, 12], which can be used to fill in the areas of the phase diagram that are inaccessible to experiments.

On top of these advantages, constant pressure simulations often make it much easier to study certain properties. Firstly, many important dynamical processes are related to pressure changes over time, and these can only be studied computationally if the pressure is used as an external variable. For example, the formation of glasses is characterized by changes in temperature and pressure over very short time scales, leading to a large degree of disorder since the structure is unable to relax to the lowest energy state. The study of these metastable states can be achieved computationally via *NPT* AI-MD simulations (see for example ref. [161]).

Similarly, constant pressure ensembles are also advantageous from a mathematical

perspective. Since they include the volume as a dynamical variable, they give us access to fluctuations that would not otherwise be available. While calculations in different ensembles are equivalent in the $N \rightarrow \infty$ limit, for many properties it is much more natural to represent them by considering changes in the volume of the system. For example, elastic constants, which describe the stress-strain relations of a solid, can be understood much more intuitively when described in terms of fluctuations in the cell parameters used to describe the system [162].

Finally, to study any structural rearrangement we must allow the cell volume to fluctuate, both because this makes it easier for molecules to diffuse past one another, and because the final structure may have a significantly different equilibrium geometry. Therefore, any simulation of phase transitions or explosive chemistry must necessarily be done at constant pressure.

5.2 Theory

Having discussed in detail the rationale for using NPT simulations, let us now see how they are implemented in practice. To begin with, let us consider the difference between the microcanonical, or NVE , ensemble and the isobaric-isoenthalpic ensemble, NPH . The latter has a similar role for calculations under constant pressure as the former does for constant volume, that of simulating a system attached to a barostat but otherwise isolated.

As the name suggests, the isobaric-isoenthalpic ensemble is characterized by the constants of motion N , P and the enthalpy, H , given by $(E + PV)$. Note that the enthalpy is not to be confused with the Hamiltonian. The stationary distribution which must be generated by any NPH dynamics is

$$\rho_{NPH} \propto \delta(-\beta(H - H_0)) \quad (5.2.1)$$

where H_0 is the initial enthalpy.

To realize this ensemble, fluctuations in the energy of the particles must be compensated for by compression or expansion of the system. However, Hamiltonian mechanics, as illustrated by equations 2.1.10 and 2.1.11, cannot be straightforwardly extended to model volume fluctuations, since it is not immediately obvious how fluctuations in the volume should influence the particle dynamics.

The major theoretical breakthrough which made constant-pressure dynamics feasible was due to Andersen [46]. He made an analogy between a system of interacting particles, as discussed so far in this thesis, and a scaled system with unit volume. The transformation between these systems was given by

$$\mathbf{r}'_j = \frac{\mathbf{r}_j}{V^{\frac{1}{3}}} \quad (5.2.2)$$

$$\mathbf{p}'_j = \mathbf{p}_j V^{\frac{1}{3}} \quad (5.2.3)$$

In the original system, assuming for now that we are working with classical mechanics, the enthalpy is

$$H = \sum_{j=1}^N \frac{|\mathbf{p}_j|^2}{2m_j} + U(\mathbf{r}_1, \dots, \mathbf{r}_N) + PV \quad (5.2.4)$$

It is then trivial to use the transformations in equations 5.2.2 and 5.2.3 to give a new function in terms of the scaled coordinates. Since the volume of the scaled system is constant, Hamiltonian mechanics can be applied to its dynamics. To do this, we first define the conjugate momentum of the volume as Π , and the mass associated with its motion as μ . The Hamiltonian we will use to simulate its motion is then the enthalpy given above, with an additional kinetic energy term corresponding to the

volume fluctuations, and is given by

$$H' = \sum_{j=1}^N \frac{|\mathbf{p}'_j|^2}{2m_j V^{2/3}} + U(\mathbf{r}'_1 V^{1/3}, \dots, \mathbf{r}'_N V^{1/3}) + PV + \frac{\Pi^2}{2\mu} \quad (5.2.5)$$

The equations of motion required to generate NVE dynamics for this scaled system can then be found trivially by applying equations 2.1.10 and 2.1.11 to H' . In terms of the original variables, these are

$$\dot{V} = \frac{\Pi}{\mu} \quad (5.2.6)$$

$$\dot{\Pi} = \sum_{j=1}^N \left[\frac{|\mathbf{p}_j|^2}{3m_j V} - \frac{\mathbf{r}_j}{3V} \frac{\partial U(\mathbf{r}_1, \dots, \mathbf{r}_N)}{\partial V} \right] - P \quad (5.2.7)$$

$$\dot{\mathbf{r}}_j = \frac{\mathbf{p}_j}{m_j} + \frac{\mathbf{r}_j \eta}{\mu} \quad (5.2.8)$$

$$\dot{\mathbf{p}}_j = \mathbf{f}_j - \frac{\mathbf{p}_j \eta}{\mu} \quad (5.2.9)$$

where η is proportional to the log-derivative of the volume, and is given by

$$\eta = \frac{\mu}{3} \frac{d \ln(V)}{dt} \quad (5.2.10)$$

By construction, these dynamics conserve H' , which is closely related to the target conserved quantity H . Also, the choice of scaling factors means that there is an exact correspondence between any given state point with a particular value of V in the scaled system and one in the original system. In fact, it can be shown that, once the dummy variable Π has been integrated out, these dynamics sample the NPH ensemble in the original system, except for negligible errors [46].

This leaves two further complications. Firstly, we must find a way to go from the NPH ensemble to the NPT ensemble, in an analogous way to the approach discussed in section 3.4 used to sample the NVT ensemble. Secondly, we need to extend this

formalism to work with path integrals.

To achieve the first of these, we use the approach of ref. [49], and attach a stochastic thermostat both to the nuclear degrees of freedom and the volume degree of freedom. This work also uses a slightly different set of equations of motion, where η is used as the conjugate momentum variable rather than Π , and it is these equations which will form the basis for our implementation of constant pressure dynamics.

When applying this to PIMD simulations we must note two things. Firstly, similar to how the effective temperature of the classical PIMD dynamics is nT , the effective pressure applied by the barostat must be nP for it to give the correct ensemble averages.

Secondly, it is much more efficient to apply the barostat only to the centroid coordinate, and propagate the internal motion of the ring polymers in the same way as for NVT dynamics. This is because, as for the kinetic energy, the preferred choice of pressure estimator uses the centroid virial form, since its variance does not scale with the number of replicas used and as it has a short auto-correlation time. However, this is derived based on the assumption that only the centroid motion is affected by volume fluctuations [48], and so it can only be used if we make the approximation that the internal motion of the ring polymers is not connected to the external barostat.

Fortunately, this turns out to be a good approximation as long as the radius of gyration of the ring polymers is much smaller than the simulation box. Since this length scale is in almost all cases smaller than that which would give rise to large finite-size effects, in practice this assumption leads to negligible errors in simulations of bulk-phase systems [48], and so this is the form of the equations of motion we will use.

The form of the pressure estimator we will use is given by

$$P_{\text{int}} = -\frac{1}{n} \sum_{i=0}^{n-1} \frac{dU(\mathbf{r}_1^{(i)}, \dots, \mathbf{r}_N^{(i)})}{dV} + \frac{1}{3nV} \left[\sum_{j=1}^N \frac{|\tilde{\mathbf{p}}_j^{(0)}|^2}{m_j} + \sum_{i=0}^{n-1} \sum_{j=1}^N (\mathbf{r}_j^{(i)} - \mathbf{r}_j^c) \cdot \frac{\partial U(\mathbf{r}_1^{(i)}, \dots, \mathbf{r}_N^{(i)})}{\partial \mathbf{r}_i^{(j)}} \right] \quad (5.2.11)$$

Note that the kinetic energy term in this estimator could equally well be substituted by its mean, $\langle \sum_j |\tilde{\mathbf{p}}_j^{(0)}|^2 / m_j \rangle = 3N/\beta_n$. However, we use the instantaneous value as this is needed for the dynamics to have a well-defined conserved quantity. Note also that we use dU/dV to signify the *total* derivative of the potential energy with respect to volume, which typically contains a term obtained from the virial.

We shall not enter into the details of how the virial should be computed, since the evaluation of the potential energy component of the pressure is delegated to the client program. Regardless of the implementation, one should make sure that the client returns the total derivative, including terms that depend explicitly on the volume such as tail corrections.

This now gives us all the ingredients needed to run constant-pressure *NPT* simulations. For clarity, we will first describe the equations of motion, and then the corresponding Liouvillian operators. We will then show that the correct ensemble is being sampled, and finally we will discuss how these equations can be integrated to give an algorithm which we can implement.

5.3 Implementation

Our choice of the equations of motion bring together the *NPT* dynamics discussed in section 3.4, the classical barostat discussed in ref. [49], and previous work on

constant-pressure PIMD [48]. The resultant equations can be summarized using

$$\dot{\tilde{\mathbf{p}}}_j^{(k)} = \sqrt{\frac{2m_j\gamma_k}{\beta_n}} \boldsymbol{\xi}_j^{(k)} - \gamma_k \tilde{\mathbf{p}}_j^{(k)} \quad (5.3.1a)$$

$$+ \tilde{\mathbf{f}}_j^{(k)} \quad (5.3.1b)$$

$$- \tilde{\mathbf{p}}_j^{(k)} \delta_{k0} \eta / \mu \quad (5.3.1c)$$

$$\left. \begin{aligned} & - m_j \omega_k^2 \tilde{\mathbf{r}}_j^{(k)} \\ \dot{\tilde{\mathbf{r}}}_j^{(k)} = \tilde{\mathbf{p}}_j^{(k)} / m_j & \end{aligned} \right\} \quad (5.3.1d)$$

$$+ \tilde{\mathbf{r}}_j^{(k)} \delta_{k0} \eta / \mu \quad (5.3.1e)$$

$$\dot{V} = 3V \eta / \mu \quad (5.3.1f)$$

$$\dot{\eta} = \sqrt{\frac{2\mu\gamma_\eta}{\beta_n}} \xi_\eta - \gamma_\eta \eta \quad (5.3.1g)$$

$$+ 3 \left[nV(P_{\text{int}} - P_{\text{ext}}) + \frac{1}{\beta_n} \right] \quad (5.3.1h)$$

where for clarity the external pressure defining the ensemble is shown as P_{ext} rather than P .

We also note that the equations of motion given in [49] have been slightly modified by including just one factor of $1/\beta_n$ in equation 5.3.1h rather than two, so that we obtain the same NPT stationary distribution as was derived in ref. [48] (see equation 5.3.4).

Before we continue, we should make sure that these equations of motion are consistent with NPT dynamics. Firstly, consider the quantity E_{cons} , which we define to be

$$E_{\text{cons}} = H_n + nP_{\text{ext}}V + \frac{\eta^2}{2\mu} - \frac{1}{\beta_n} \ln V + \Delta H, \quad (5.3.2)$$

where ΔH is the heat transferred between the system and the barostat during steps in equations 5.3.1a and 5.3.1g, as explained in refs. [45, 88].

The rate at which this quantity changes can be calculated with the chain rule.

Since E_{cons} does not explicitly depend on time, this is given by

$$\dot{E}_{\text{cons}} = \sum_x \dot{x} \frac{\partial E_{\text{cons}}}{\partial x} \quad (5.3.3)$$

where the sum is taken over all degrees of freedom, including V and η .

Noting that the equations of motion give us \dot{x} , and that \dot{E}_{cons} is trivially zero over the thermostating steps by construction due to the term ΔH , we find that E_{cons} is a constant of the motion. This is important, since the logic of refs. [88] and [49] suggests that this gives us an objective measure of the sampling errors, which will help us when deciding the integration time step. In general, if the drift in E_{cons} is not much smaller than the energy scales of interest, then the time step should be reduced.

Next, we should make sure that the correct stationary distribution is being sampled. While more than one NPT ensemble has been described in the literature [163], we will take it to be [1, 2]

$$\rho(\mathbf{p}, \mathbf{r}, \eta, V) \propto \exp\left(-\beta_n \left[H_n + nP_{\text{ext}}V + \frac{\eta^2}{2\mu} \right]\right) \quad (5.3.4)$$

where we again note that the effective pressure is nP_{ext} , and that the distribution includes a contribution from the dummy variable η , which can nevertheless be simply integrated out to give the desired isobaric-isothermal stationary distribution.

We can verify that this distribution is stationary with respect to the equations of motion given in equations 5.3.1a - 5.3.1g using the Liouvillian formalism of section 3.2. To do this, we simply note that $\mathcal{L}\rho(\mathbf{p}, \mathbf{r}, \eta, V) = 0$, where the total Liouvillian,

\mathcal{L} , is given by

$$\begin{aligned}
\mathcal{L} &= \sum_{k=0}^{n-1} \sum_{j=1}^N \mathcal{L}_\gamma^{(jk)} + \mathcal{L}_{\gamma_\eta} + \sum_{k=0}^{n-1} \sum_{j=1}^N \mathcal{L}_U^{(jk)} + \mathcal{L}_\eta \\
&+ \left(\sum_{j=1}^N \mathcal{L}_0^{(j0)} + \sum_{j=1}^N \mathcal{L}_p^{(j)} + \sum_{j=1}^N \mathcal{L}_r^{(j)} + \mathcal{L}_V \right) + \sum_{k=1}^{n-1} \sum_{j=1}^N \mathcal{L}_0^{(jk)} \\
&= \mathcal{L}_\gamma + \mathcal{L}_{\gamma_\eta} + \mathcal{L}_U + \mathcal{L}_\eta + \mathcal{L}_0^0 + \mathcal{L}_0'.
\end{aligned} \tag{5.3.5}$$

and where the Liouville operators corresponding to each of the equations of motion 5.3.1a - 5.3.1g are

$$\mathcal{L}_\gamma^{(jk)} = -\gamma_k \frac{\partial(\tilde{\mathbf{p}}_j^{(k)} \cdot)}{\partial \tilde{\mathbf{p}}_j^{(k)}} - \frac{m_j \gamma_k}{\beta_n} \frac{\partial^2 \cdot}{(\partial \tilde{\mathbf{p}}_j^{(k)})^2} \tag{5.3.6a}$$

$$\mathcal{L}_U^{(jk)} = \tilde{\mathbf{f}}_j^{(k)} \frac{\partial \cdot}{\partial \tilde{\mathbf{p}}_j^{(k)}} \tag{5.3.6b}$$

$$\mathcal{L}_p^{(j)} = -\frac{\eta}{\mu} \frac{\partial(\tilde{\mathbf{p}}_j^{(0)} \cdot)}{\partial \tilde{\mathbf{p}}_j^{(0)}} \tag{5.3.6c}$$

$$\mathcal{L}_0^{(jk)} = \frac{\tilde{\mathbf{p}}_j^{(k)}}{m_j} \frac{\partial \cdot}{\partial \tilde{\mathbf{r}}_j^{(k)}} - m_j \omega_k^2 \tilde{\mathbf{r}}_j^{(k)} \frac{\partial \cdot}{\partial \tilde{\mathbf{p}}_j^{(k)}} \tag{5.3.6d}$$

$$\mathcal{L}_r^{(j)} = \frac{\eta}{\mu} \frac{\partial(\tilde{\mathbf{r}}_j^{(0)} \cdot)}{\partial \tilde{\mathbf{r}}_j^{(0)}} \tag{5.3.6e}$$

$$\mathcal{L}_V = 3 \frac{\eta}{\mu} \frac{\partial(V \cdot)}{\partial V} \tag{5.3.6f}$$

$$\mathcal{L}_{\gamma_\eta} = -\gamma_\eta \frac{\partial(\eta \cdot)}{\partial \eta} - \frac{\mu \gamma_\eta}{\beta_n} \frac{\partial^2 \cdot}{\partial \eta^2} \tag{5.3.6g}$$

$$\mathcal{L}_\eta = 3 \left[nV (P_{\text{int}} - P_{\text{ext}}) + \frac{1}{\beta_n} \right] \frac{\partial \cdot}{\partial \eta} \tag{5.3.6h}$$

Having shown that these equations of motion give the correct ensemble, we now must make a choice of how to propagate the dynamics. The integration scheme we have implemented in i-PI follows closely that of ref. [49] and is based on the following

symmetric Trotter splitting of the propagator

$$e^{-\mathcal{L}\Delta t} \approx e^{-(\mathcal{L}_\gamma + \mathcal{L}_{\gamma\eta})\Delta t/2} e^{-(\mathcal{L}_U + \mathcal{L}_\eta)\Delta t/2} e^{-(\mathcal{L}_0^0 + \mathcal{L}_0^1)\Delta t} e^{-(\mathcal{L}_U + \mathcal{L}_\eta)\Delta t/2} e^{-(\mathcal{L}_\gamma + \mathcal{L}_{\gamma\eta})\Delta t/2} \quad (5.3.7)$$

which leaves us with four sets of differential equations, the solutions of which are derived in appendix B for reference.

Firstly, the thermostat is applied to the ring polymer and cell momenta for half a time step according to equations 5.3.1a and 5.3.1g, using

$$\begin{aligned} \eta &\leftarrow e^{-\gamma\eta\Delta t/2}\eta + \sqrt{\frac{\mu}{\beta_n}(1 - e^{-\gamma\eta\Delta t})}\xi_\eta \\ \tilde{\mathbf{p}}_j^{(k)} &\leftarrow e^{-\gamma_k\Delta t/2}\tilde{\mathbf{p}}_j^{(k)} + \sqrt{\frac{m_j}{\beta_n}(1 - e^{-\gamma_k\Delta t})}\boldsymbol{\xi}_j^{(k)} \end{aligned} \quad (5.3.8)$$

Next, the momenta are evolved under the action of the pressure term and the inter-atomic forces for half a time step according to equations 5.3.1b and 5.3.1h

$$\begin{aligned} \eta &\leftarrow \eta + 3\frac{\Delta t}{2} \left[nV(P_{\text{int}} - P_{\text{ext}}) + \frac{1}{\beta_n} \right] \\ &\quad + \left(\frac{\Delta t}{2} \right)^2 \sum_j \frac{1}{m_j} \tilde{\mathbf{f}}_j^{(0)} \cdot \tilde{\mathbf{p}}_j^{(0)} + \left(\frac{\Delta t}{2} \right)^3 \sum_j \frac{1}{3m_j} \tilde{\mathbf{f}}_j^{(0)} \cdot \tilde{\mathbf{f}}_j^{(0)} \\ \tilde{\mathbf{p}}_j^{(k)} &\leftarrow \tilde{\mathbf{p}}_j^{(k)} + \frac{\Delta t}{2} \tilde{\mathbf{f}}_j^{(k)} \end{aligned} \quad (5.3.9)$$

Note that these two steps should be done in this order, so that the value of $\tilde{\mathbf{p}}_j^{(0)}$ used to update η is correct.

This brings us to the central term in equation 5.3.7, which comprises two independent (and commuting) Liouvillians. One of these, \mathcal{L}_0^0 , evolves the centroid using

an algorithm identical to the stochastic classical barostat of ref. [49],

$$\begin{aligned}\tilde{\mathbf{p}}_j^{(0)} &\leftarrow \tilde{\mathbf{p}}_j^{(0)} e^{-\Delta t \eta / \mu} \\ \tilde{\mathbf{r}}_j^{(0)} &\leftarrow \tilde{\mathbf{r}}_j^{(0)} e^{\Delta t \eta / \mu} + \frac{\sinh(\Delta t \eta / \mu)}{\eta / \mu} \frac{\tilde{\mathbf{p}}_j^{(0)}}{m_j} \\ V &\leftarrow V e^{3\Delta t \eta / \mu},\end{aligned}\tag{5.3.10}$$

while the other, \mathcal{L}'_0 , evolves the internal modes using an algorithm identical to the free ring polymer *NVE* integrator shown in section 3.4

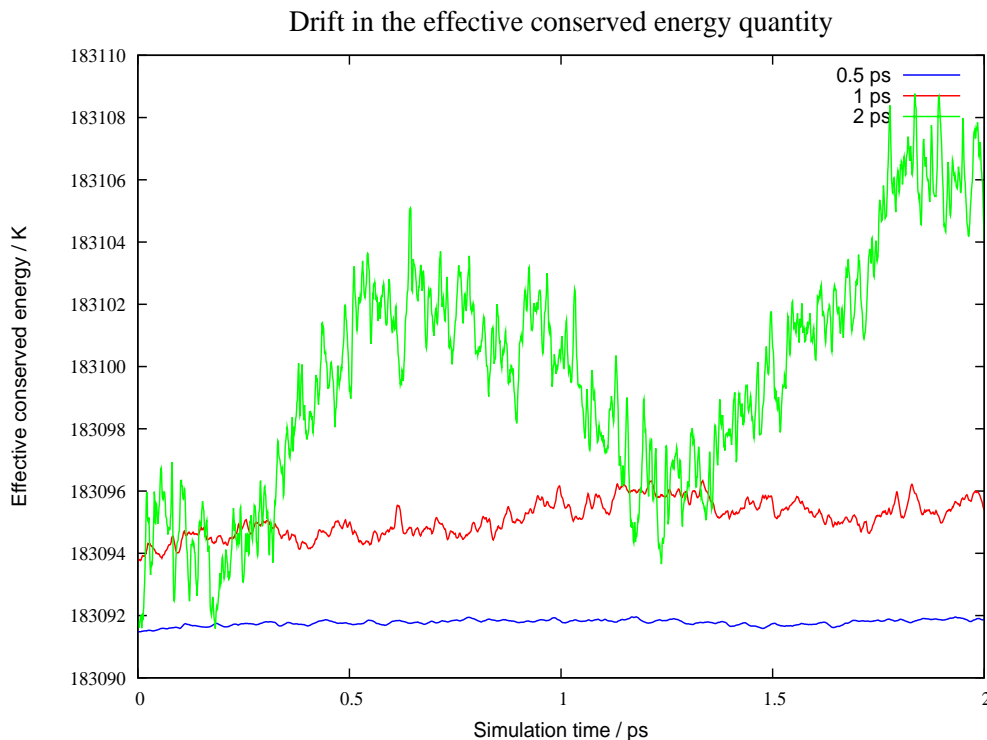
$$\begin{pmatrix} \tilde{\mathbf{p}}_j^{(k)} \\ \tilde{\mathbf{r}}_j^{(k)} \end{pmatrix} \leftarrow \begin{pmatrix} \cos(\omega_k \Delta t) & -m_j \omega_k \sin(\omega_k \Delta t) \\ \frac{1}{m_j \omega_k} \sin(\omega_k \Delta t) & \cos(\omega_k \Delta t) \end{pmatrix} \begin{pmatrix} \tilde{\mathbf{p}}_j^{(k)} \\ \tilde{\mathbf{r}}_j^{(k)} \end{pmatrix}\tag{5.3.11}$$

At this point, one continues with the second part of the integrator, executing again the steps in equation 5.3.9, and finally applying the thermostats as in equation 5.3.8. Note that the modular nature of this integration scheme means it is very easy to implement a different thermostating approach for the ions or the cell, simply by changing equation 5.3.8 to the appropriate propagator.

5.4 Testing the implementation

To test that the code was working correctly, we compared the results obtained by i-PI against previous work. Fortunately, there already exists an implementation of constant pressure PIMD in the literature, as reported in ref. [48]. This paper, among other things, looked at the simple system of liquid para-hydrogen, for which an accurate isotropic interaction potential is well known [164]. This system can therefore be considered a pseudo-atomic liquid for which nuclear quantum effects are non-negligible, and so is often used as a first test for path integral methods.

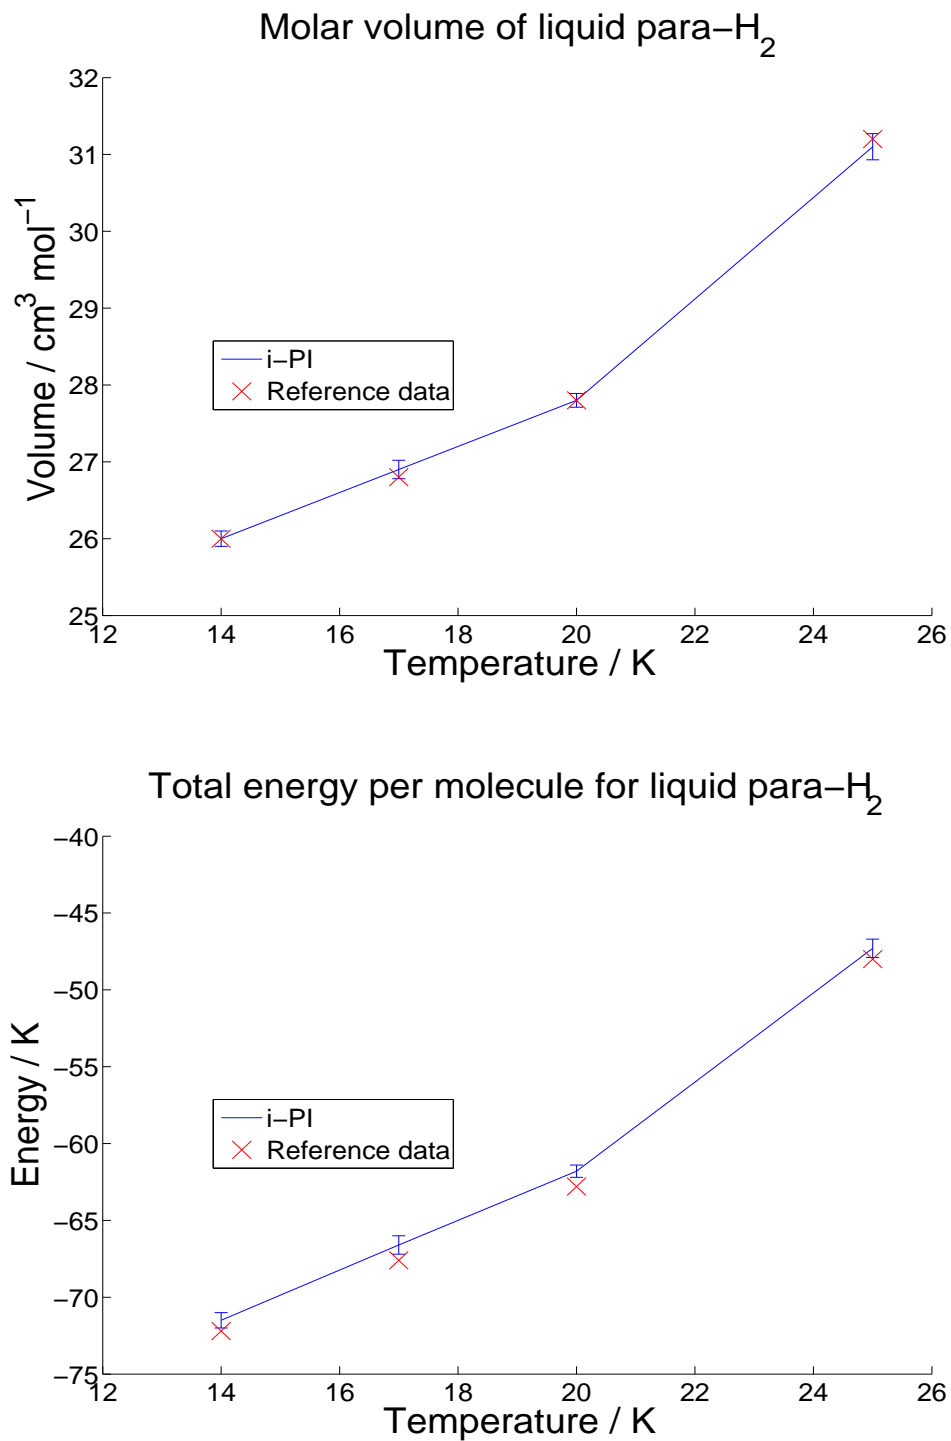
Figure 5.1: This plot shows how the drift in the conserved energy quantity can be systematically reduced by decreasing the time step.



The simulation parameters were chosen so that our results could be compared directly to those of ref. [48]. All calculations were run using 172 hydrogen molecules and with a time step of 1 fs, unless otherwise stated. Atmospheric pressure was used, and the temperatures considered were 14 K, 17 K, 20 K and 25 K. These required 32, 24, 20 and 16 ring polymer beads respectively.

The system was initially thermalized by running a short *NVT* run of 10 ps, for which we used a stochastic velocity rescaling thermostat with a time constant of 10 fs. The final configuration from this trajectory was used as a starting point for a 20 ps *NPT* simulation, for which the time constant of both the barostat and the

Figure 5.2: This plot shows how i-PI can reproduce the results of comparable constant pressure PIMD implementations.



Langevin thermostat attached to the cell momentum was 100 fs.

As an additional test, the 25 K simulation was run for 2 ps with three different time steps from the same initial configuration to see if the drift in the effective conserved quantity, as defined in ref. [88], could be controlled. As we can see from figure 5.1, reducing the time step by a factor of 4 drastically reduced the observed drift, giving us confidence that the dynamics do indeed have the correct conserved quantity and are sampling the required ensemble.

Similarly, looking at figure 5.2, we see that the values obtained by i-PI for important properties such as the total energy and the average volume agree well with those reported in the literature. While the total energy may on average be slightly high, this can probably be accounted for by slight differences in how the cut-off radius for the interaction potential is treated.

Given that there are no apparent errors in sampling the NPT ensemble, and that i-PI gives averages that agree with previous results, we believe that our implementation of constant pressure dynamics is robust. With this established, we now turn our attention to the more theoretically interesting system of high-pressure water.

5.5 An example application: high pressure water

To show how this novel NPT integrator works in practice, we have used it to perform an illustrative *ab initio* PIGLET [59] simulation of supercritical water at 750 K and 10 GPa. This simulation is both technically challenging and physically interesting. We shall use it to reveal the importance of NQEs under conditions similar to those explored in the pioneering work of ref. [165], in which calculations were performed at constant volume and without quantum effects in the nuclear motion.

5.5.1 Computational details

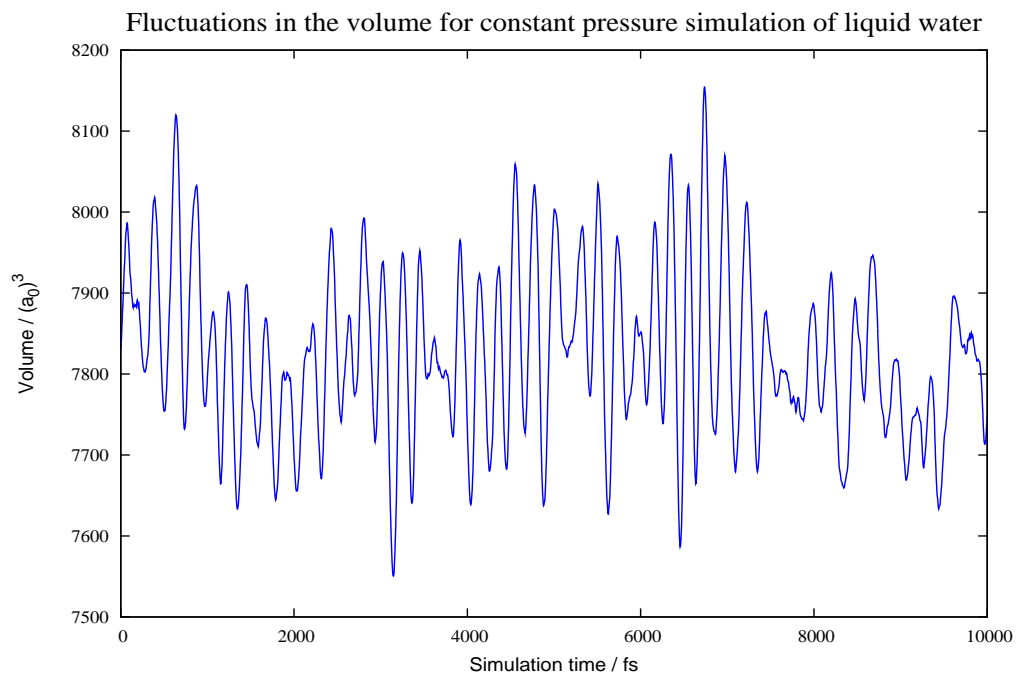
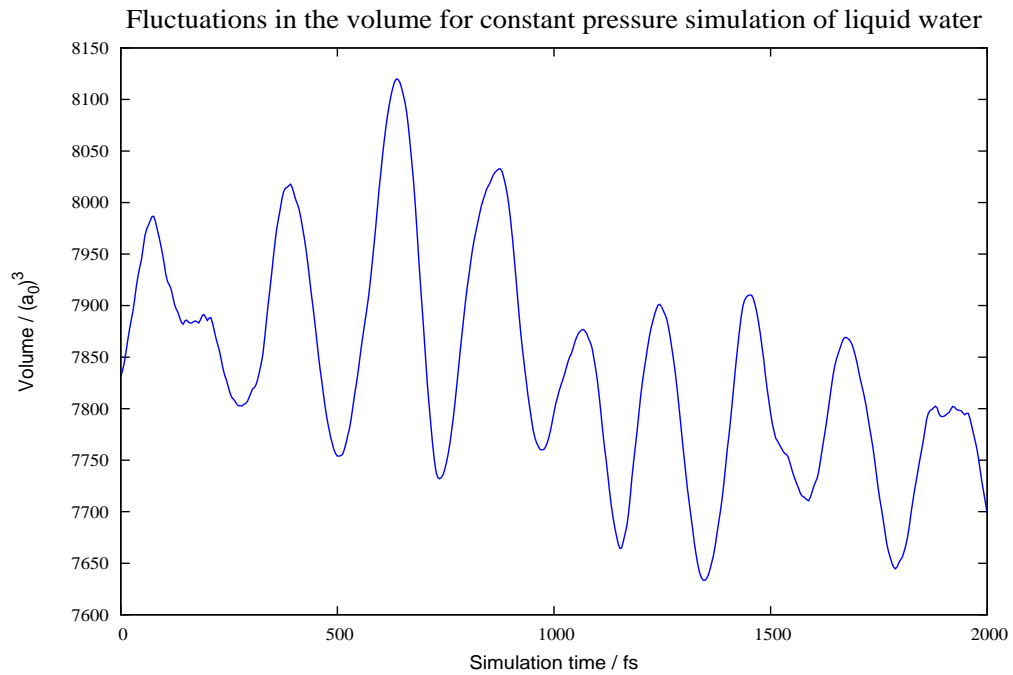
All the present calculations used a simulation box containing 64 water molecules, initially equilibrated at the desired state point using an empirical forcefield [6]. Starting from this configuration, we performed 25 ps of *NVT* dynamics with classical nuclei and *ab initio* forces. We used i-PI [44], with a patched version of CP2K [137] as the electronic structure client. We used the BLYP exchange-correlation functional [121, 122] and GTH pseudopotentials [166]. Wave functions were expanded in the Gaussian DZVP basis set, while the electronic density was represented using an auxiliary plane wave basis, with a kinetic energy cutoff of 300 Ry.

In our constant-pressure simulations we used a TZV2P basis set, a cutoff of 800 Ry, and a plane wave grid consistent with a cubic reference cell with a side of 10.5 Å (the higher cutoff was necessary to converge the stress tensor used to define the pressure virial). We included D3 empirical Van der Waals corrections [167], which are essential to obtain a reasonable description of water at constant pressure [168, 169].

After the preliminary *NVT* equilibration, we performed a reference 40 ps *NPT* trajectory with classical nuclei, with the first 5 ps discarded for equilibration. The integration time step was 0.5 fs. We used a stochastic velocity rescaling thermostat [86] for the ions, with a time constant of 20 fs, and an optimal-sampling GLE thermostat [79] for the fictitious cell momentum, η , for which we chose a fictitious mass consistent with $\tau_\alpha = 500$ fs.

We then performed 40 ps of simulation with quantum ions, modelling the quantum nature of nuclei using a 4-bead path integral supplemented with PIGLET coloured noise. The fictitious cell momentum was thermostatted using an optimal-sampling, classical GLE thermostat. Input files for both i-PI and CP2K are provided among the examples attached to the source code of i-PI.

Figure 5.3: Plot to show how the simulation box volume varies with time on short and long time scales. For clarity, only 10 ps of trajectory from the middle of the trajectory are shown.






Figure 5.4: Histogram of the density (upper panel) and of the instantaneous internal pressure estimator in equation 5.2.11 (lower panel), comparing NPT simulations with classical and quantum nuclei.

5.5.2 Results

Firstly, we looked at the fluctuations in the simulation volume, as shown in figure 5.3. We can clearly see that there is a time scale of fluctuations of the order of the barostat time constant, as we would expect. This gives another reason to believe that our implementation of the constant pressure dynamics is reasonable.

Figure 5.4 reports the histograms of density and internal pressure obtained during the simulations. The average density in the classical simulation is 1.644 ± 0.001 g/cm^3 , which is qualitatively consistent with the simulation of ref. [165], which observed an average pressure of 14.5 GPa in the NVT ensemble at a density of 1.72 g/cm^3 . The density appears to be slightly higher for simulations with quantum nuclei, 1.654 ± 0.003 g/cm^3 , but one should note that the correlation time for ρ tends to be underestimated in relatively short *ab initio* runs, so the quantum effect is barely significant given that our computed error bars are likely to be optimistic. Note that the estimator for the internal pressure is correctly peaked around the value fixed by the definition of the ensemble.




Figure 5.5: Radial distribution functions from *NPT* simulations of water at a temperature of 750 K and a pressure of 10 GPa. Simulations with classical and quantum nuclei are compared.

The small role that NQEs play in determining the equilibrium density is not surprising: even at room temperature, the isotope effect on the number density of water is extremely small. However, one should not think that at these high temperatures the quantum nature of light nuclei can be neglected. It is clear from the radial distribution functions reported in figure 5.5 that, while nuclear quantum effects do not change the long-range structure, they do have a very sizeable impact on the short-range part of $g(r)$.

There is very little effect on the oxygen-oxygen $g_{OO}(r)$ (the difference is barely significant considering the statistical error bars), but $g_{HH}(r)$ and $g_{OH}(r)$ are noticeably over-structured for $r < 2 \text{ \AA}$ in the classical simulation. Notice in particular how the quantum $g_{OH}(r)$ departs significantly from zero between the intra-molecular and inter-molecular regions. This is a sign that the combined effect of pressure and NQEs leads to significant delocalization of protons along the hydrogen bonds.

To investigate this delocalization further, we have adopted the simple protocol used in ref. [165] to identify charged species. Each proton in the simulation is assigned




Figure 5.6: These bar charts depict the probability of finding n_{+-} charged pairs in a snapshot of our simulation box. The upper panel is for a classical simulation, and the lower panel for a quantum simulation. The dots joined with lines represent the values for a binomial distribution fitted to the computed probabilities – each snapshot contains 64 water molecules, so can accommodate a maximum of 32 ion pairs. The best-fit probability of a pair being ionized is found to be $p_{+-} = 0.019$ for the quantum simulation and $p_{+-} = 0.0014$ for the classical simulation.

to the closest oxygen atom. One can then distinguish between neutral, positively and negatively charged species based on the number of protons assigned to each oxygen. We will refer to “+” species as those oxygen atoms to which three protons have been assigned, and “–” species as those with just one.

Given the artificial nature of this procedure, we do not intend to imply that these species correspond to water, hydronium and hydroxide. However, they do provide a simple way to characterise how much each frame departs from the conventional picture of a molecular fluid composed of neutral molecules. In all of our simulations, we only detected a single instance of a “doubly ionised” oxygen, so in the discussion that follows we shall assume that the charged O atoms always form in singly charged pairs.

Figure 5.6 shows very clearly just how important NQE are in determining the fluctuations of protons along hydrogen bonds. Most frames in the classical simulation

only contain neutral species; the concentration of $+$ or $-$ is $\approx 0.07\%$, even smaller than that observed in ref. [165], where the density was higher.

In the quantum simulation, the concentration is much larger, about 0.97% , and there is a fairly large probability of having more than one pair of ions in any given frame. Interestingly, in both cases the probability of finding n_{+-} ionised pairs follows closely a binomial distribution, which is a sign that the ionisation events are only weakly correlated with each other.




Figure 5.7: Radial distribution functions for positively and negatively charged defects in *NPT* simulations of water at 10 GPa and 750 K. Results from quantum and classical simulations are compared, and the lower panel shows the long-range part of $g_{+-}(r)$ on an enlarged scale. Given the low concentration of defects, the radial distribution functions are very noisy, and the normalisation is problematic, so the scale of the ordinate should be considered as arbitrary.

It is interesting to investigate the spatial distribution of these ion pairs: in figure 5.7 we show the radial distribution functions $g_{+-}(r)$. The distributions are very noisy, particularly in the classical case, where we have only a handful of snapshots containing more than a single ion pair.

In the quantum simulation, we could collect better statistics, because the concentration of ion pairs is larger. In this case, it is clear that the radial distribution function is almost flat for $r > 3 \text{ \AA}$ (see lower panel of figure 5.7), which is consistent

with the binomial distribution of n_{+-} in figure 5.6. Furthermore, the upper panel of figure 5.7 clearly shows that, in the vast majority of cases, what is detected as a pair of ions is simply an excursion of a proton along a compressed hydrogen bond. This is consistent with the observation of short-lived charged species in classical simulations of water at considerably higher pressure [170].

As was shown in a recent study of water under ambient conditions [171], nuclear quantum effects enhance the delocalization of the proton along the hydrogen bond. This suggests that perturbations that modulate the average O–O distance, such as pressure, might trigger auto-ionisation more easily than in the classical case.

In fact, not all of the ion pairs in the present simulations are close together and short-lived: if we define an isolated ion as a positive species that has no negative counterpart within 3 Å of it, we find a concentration of $6 \times 10^{-3}\%$ of these isolated ions in our classical simulations, and of 0.24% in our quantum simulations. This is not far from the value of 0.5 % measured in shock-compressed water at 13 GPa [159,172]. However, our definition of an isolated ion is somewhat arbitrary, and the concentrations we quote will be strongly system-size dependent.

Nevertheless, the ratio between these concentrations and the overall fraction of ionised species indicates that quantum effects have an even more dramatic impact on genuine auto-ionisation events than they do on local hydrogen bond fluctuations. This is perhaps not too surprising, given the significant isotope effect on the pH of water under ambient conditions.

The presence of a significant fraction of ionised species increases the mobility of protons, which can hop from molecule to molecule in a Grotthuss-like fashion. Even though PIMD (and particularly the heavily-thermostatted PIGLET method) does not allow us to make quantitative statements about how quantum effects enhance proton mobility, one can clearly see that by the end of our quantum simulation a




Figure 5.8: A snapshot from the final stage of our 40 ps quantum simulation of water at 750 K and 10 GPa. Protons that are bound to a different oxygen than the one they were bound to in the initial configuration of the trajectory are highlighted in blue. Note that many exchanges have occurred. In our classical simulation, no exchanges were observed.

significant fraction of the protons has been exchanged between water molecules (see figure 5.8). On the contrary, in the classical case all the water molecules maintain their chemical identity throughout our simulation.

Using an approximate quantum dynamics technique such as RPMD [25] or CMD [22, 23] to study the dynamics in high pressure water would be an interesting future application of i-PI.

Chapter 6

Calculation of momentum distributions

6.1 Background

One area of interest for which NQE must be included to get even a qualitative understanding of the underlying physics is in the calculation of nuclear momentum distributions. The momentum distribution of a classical particle is completely determined by its mass and the temperature of the system; however, this is not true quantum mechanically since, due to the non-commutativity of the position and momentum operators, the momentum distribution of a particle is affected by the underlying potential energy surface.

As such, any deviation from a Maxwell-Boltzmann distribution is both purely quantum mechanical in nature and contains useful information about the physics of the system. In particular, the position-momentum uncertainty principle suggests that a strongly localized particle, i.e. one in a steep potential well, will necessarily have a broad momentum distribution. This has been used, for example, to determine

whether there are single or double-well hydrogen bonds in the proton conductors $\text{Rb}_3\text{H}(\text{SO}_4)_2$ [173] and $\text{H}_3\text{OSbTeO}_6$ [174].

In the context of theoretical investigations of condensed phase materials, the main reason that we are interested in calculating nuclear momentum distributions is that the deviation from classical behaviour in the motion of nuclei underlies a wide variety of kinetic isotope effects (KIE) of current interest in chemical physics.

In particular, studying the different contributions to the quantum kinetic energy of water have helped explain some of its puzzling properties. Water is a well-studied system in physics and chemistry, both because of its ubiquity in nature and use as a common solvent, and because the properties are influenced in a non-trivial way by hydrogen bonding [175–177] and NQE [6]. Indeed, many of these properties are still not well understood.

For example, despite the fact that we might expect that the NQE in water will be large, the ZPE in the OH bond being much greater than the thermal energy even at room temperature, the difference in several properties such as the melting point [178–180] and diffusion constant [181, 182] between light and heavy water is actually very small.

More intriguingly, only some empirical potential models of water are able to reproduce this effect. When the dynamics of liquid water were studied [6], it was found that an anharmonic OH bond was necessary to get agreement with experimental data in the ratio between the diffusion constants of H_2O and D_2O , and that models without this feature showed a much larger nuclear quantum effect.

This suggests that the contribution of NQE to the dynamics of liquid water is in fact significant, but that only minor deviations from classical behaviour are observed in many properties because of competition between the influence of inter- and intra-molecular motion. Indeed, knowing that the diffusion of liquid water is

heavily influenced by the strength of the hydrogen bond network, this gives a simple hypothesis for the origin for this effect.

Regardless of how the interaction potential is calculated, all quantum simulations are expected to show a disruption of the hydrogen bond network due to the delocalization of the hydrogens, which will lead to an increase in the rate of diffusion. However, if potentials that include the effect of anharmonicity are used we also expect that the bonds will lengthen due to ZPE, increasing the molecular dipole moment. This will strengthen the hydrogen bond network, countering the effect of quantum delocalization.

To test this theory, the intra- and intermolecular motions must be studied separately. This can be done by comparing the components of the single particle quantum kinetic energy tensor, as these correspond approximately to the contributions of the OH bond stretch, the bending motion and the molecular libration. These then split cleanly into one intermolecular component, the librational motion, and two intramolecular components.

Kinetic energies are particularly interesting to study in the context of competing quantum effects (CQE) since they can be directly related to free energy differences. For example, the dependence of the Helmholtz free energy, A , on mass can be found [140, 183] starting from

$$\frac{\partial A}{\partial m} = -\frac{1}{\beta Q} \frac{\partial Q}{\partial m} = -\frac{\langle \hat{K}(m) \rangle}{m} \quad (6.1.1)$$

Integrating this equation with respect to mass gives [140]

$$A(m') - A(m) = \Delta_{m \rightarrow m'} A = -\int_m^{m'} \frac{\langle \hat{K}(\mu) \rangle}{\mu} d\mu \quad (6.1.2)$$

Assuming that the system is approximately harmonic, such that the kinetic energy is proportional to $m^{-1/2}$, this can be simplified to

$$\Delta_{m \rightarrow m'} A \approx 2 \langle \hat{K}(m) \rangle \left(\sqrt{\frac{m}{m'}} - 1 \right) \quad (6.1.3)$$

$$\approx 2 \langle \hat{K}(m') \rangle \left(1 - \sqrt{\frac{m'}{m}} \right) \quad (6.1.4)$$

Finally, comparing the above difference between two different chemical environments, X and Y say, gives

$$\Delta_{m \rightarrow m'}^X A - \Delta_{m \rightarrow m'}^Y A = \Delta \Delta A \quad (6.1.5)$$

$$\approx 2 \left(\langle \hat{K}(m) \rangle_X - \langle \hat{K}(m) \rangle_Y \right) \left(\sqrt{\frac{m}{m'}} - 1 \right) \quad (6.1.6)$$

$$\approx 2 \left(\langle \hat{K}(m') \rangle_X - \langle \hat{K}(m') \rangle_Y \right) \left(1 - \sqrt{\frac{m'}{m}} \right) \quad (6.1.7)$$

As an example, it has been shown experimentally that with 50:50 mixture of hydrogen and deuterium there is an excess of OH bonds at the surface of water [184–187], and that heavy water prefers to be in the liquid phase rather than the vapour phase at equilibrium [188–191]. The equilibrium constant, K_{eq} , for exchanging an atom of mass m in environment X with one of mass m' in environment Y is given by [72, 73]

$$K_{\text{eq}} = \exp(-\beta \Delta \Delta A) \quad (6.1.8)$$

Substituting X = surface/vapour, Y = liquid, $m = m_{\text{D}}$ and $m' = m_{\text{H}}$ in equations 6.1.6 and 6.1.8, and comparing to the experimental isotope fractionation ratios, suggests that the kinetic energy of the protons in the liquid phase must be higher than either those at the surface or those in the vapour.

However, it is also seen experimentally that the magnitude of this free-energy

change is relatively small compared to the ZPE. This suggests that a detailed study of the kinetic energy tensor and its components may show an experimentally detectable CQE in the behaviour of liquid water.

This was done in refs [72] and [73], which studied the liquid-vapour phase separation of light and heavy water and isotope fractionation of DOH/HOH mixtures at the liquid-air interface respectively. In both cases, it was found that, while the bulk liquid water phase had the highest total kinetic energy, the increase due to the intermolecular motion was partially cancelled by an decrease in the intramolecular kinetic energy components, as expected from the physical argument given at the beginning of this chapter.

More evidence of CQE was found in a study of the melting of heavy water [192]. This related the change in free energy given in equation 6.1.6 to another experimental observable, the difference in melting temperature between light and heavy water.

For a given function, F say, let $\Delta_{\text{fus}}F = F_{\text{solid}} - F_{\text{liquid}}$. If we take $T_{\text{melt}}(m)$ to be the melting point for the system with the isotope of mass m , then we know that $\Delta_{\text{fus}}A(m, T_{\text{melt}}(m)) = 0$. Given that under constant pressure conditions the Gibbs free energy, $\Delta_{\text{fus}}G$, could be used instead of $\Delta_{\text{fus}}A$, it then follows from the definition of G that

$$\Delta_{\text{fus}}H(m) = T_{\text{melt}}(m)\Delta_{\text{fus}}S(m) \quad (6.1.9)$$

Therefore, assuming that $\Delta_{\text{fus}}S(m)$ is approximately constant with respect to temperature, this gives the following free energy change with temperature

$$\Delta_{\text{fus}}G(m, T) = \Delta_{\text{fus}}S(m) (T_{\text{melt}}(m) - T) \quad (6.1.10)$$

Finally, after noting that equation 6.1.5 can be rearranged to give $\Delta\Delta A = \Delta\Delta G = \Delta_{\text{fus}}G(m') - \Delta_{\text{fus}}G(m)$, equations 6.1.6 and 6.1.10 can be combined at the melting

point of each species to give

$$\Delta_{\text{fus}} \langle \hat{K}(m_{\text{D}}, T_{\text{melt}}(m_{\text{H}})) \rangle = \frac{\Delta_{\text{fus}} S(m_{\text{D}})}{2 \left(\sqrt{\frac{m_{\text{D}}}{m_{\text{H}}}} - 1 \right)} \left(T_{\text{melt}}(m_{\text{H}}) - T_{\text{melt}}(m_{\text{D}}) \right) \quad (6.1.11)$$

$$\Delta_{\text{fus}} \langle \hat{K}(m_{\text{D}}, T_{\text{melt}}(m_{\text{D}})) \rangle = \frac{\Delta_{\text{fus}} S(m_{\text{H}})}{2 \left(\sqrt{\frac{m_{\text{D}}}{m_{\text{H}}}} - 1 \right)} \left(T_{\text{melt}}(m_{\text{H}}) - T_{\text{melt}}(m_{\text{D}}) \right) \quad (6.1.12)$$

This shows from simple thermodynamic arguments that CQE in the kinetic energy change of freezing should account for the fact that there is very little difference in the melting point of light and heavy water. Indeed the value of $\Delta_{\text{fus}} \langle \hat{K}(m_{\text{D}}, T_{\text{melt}}(m_{\text{D}})) \rangle$ is small enough that the simulations in ref. [192] found it to vanish to within statistical accuracy, again because of a cancellation of contributions from the inter- and intra-molecular motion.

6.2 Experimental techniques

6.2.1 Deep inelastic neutron scattering

While the last two sections establish that nuclear momentum distributions can give useful insight into the physics of systems of chemical interest, it still remains to be shown how they can be measured.

Experimentally, the only method currently capable of determining the single-particle momentum distribution of nuclei is deep inelastic neutron scattering (DINS). While the theory behind this technique has long been known [193], there were several technological hurdles which had to be overcome before it could be used in practice.

The first requirement of any DINS experiment is a beam of high-energy neutrons. However, this can be problematic, since the energy regime needed to generate Compton scattering of neutrons is very high (usually around 10-100 eV), and it is only with

the increasing availability of spallation neutron sources that reaching it has become possible. These take protons from a particle accelerator and collide them with a sample of high atomic mass, displacing large numbers of high-energy neutrons which can then be directed into beams and used to probe the sample.

Secondly, in order to be able to interpret the recorded spectrum the energy of either the incident or the scattered neutrons must be fixed at a particular known value. The two main methods of energy selection are to use a chopper or a nuclear resonance absorption foil [194].

A chopper is a rapidly rotating disc with a small aperture, which allows only a very narrow beam of neutrons to pass to the sample at a time. By placing this a certain distance from the source, only neutrons within a specific range of velocities will be able to pass through and reach the sample.

A resonance foil is made of a material that absorbs neutrons of a particular momentum (currently the Vesuvio instrument at the Rutherford Appleton laboratory uses either gold or uranium). The recorded spectrum then is the difference between a measurement with the foil in place and one without.

The resonance foil technique is generally the preferred method, for two reasons. Firstly, this method records a difference spectrum, and so the background white noise caused by multiple scattering events and detector cross-talk is automatically removed. Secondly, the maximum energy that can be produced by a chopper is limited by its rotational speed, whereas in theory a resonance foil can select neutrons with any momentum, with a resolution that only depends on the width of the absorption peak. Since the approximations used to interpret DINS measurements get more accurate as the incident neutron momentum increases, this is a distinct advantage.

Generally, to get the highest resolution two resonance foils of the same width are used [195]. One is attached directly to the detector, so that whenever it absorbs a

neutron the γ -ray cascade it causes is immediately picked up. The other is cycled in and out of the path of the neutron beam, so that the difference spectrum can be calculated.

This is known as the foil cycling (FC) technique, and has several advantages over using a single resonance foil. Firstly, with two foils the shape of the recorded spectrum is dependent on the square of the absorption coefficient of the material used, rather than being directly proportional to it. This drastically reduces the line width of the peak, and so improves energy resolution.

There is also a greater difference in count rates between when the foil is in and when it is out for the FC technique. Since the recorded spectrum is the difference between these two this makes this method more robust, ameliorating the effect of random noise and drift in the recorded spectrum over time.

Finally, this also has the consequence that the overall count rates are reduced by being absorbed through two foils. This means that greater incident neutron intensities can be used with a detector which has a given dead time, as the rate at which particles must be counted is lower.

Bringing these techniques together, the experimental set-up consists of a sample holder connected to the neutron beam, surrounded by detectors at different scattering angles, and with energy selectors positioned appropriately in the path of the neutron beam.

Since either the incident or the scattered neutron energy has been fixed, and the distance travelled is known from the experimental set-up, the remaining energy which is unknown can be determined from the time of flight. The momentum transfer can also be calculated from the energy change and the angle of deflection by classical mechanical arguments. Finally, this is combined with a knowledge of the absorption characteristics of the detector and the count rates to calculate the scattering cross-

section [194].

6.2.2 Impulse approximation

While it is in general difficult to interpret neutron scattering profiles [196], for sufficiently high neutron energies the analysis can be simplified considerably using the so-called impulse approximation (IA) [197].

This approximation comes in two parts. Firstly, it assumes that the collision takes place on a time scale which is small enough that only the neutron and the struck particle need be considered. Scattering that obeys this criterion is said to be incoherent.

Secondly, it is assumed that no energy is transferred to neighbouring particles during the collision. This gives rise to an energy conservation condition in the calculation of the structure factor and means that the struck particle recoils as a free particle, and so its wavefunction is a plane wave.

The momentum transfer, \mathbf{q} , is usually given in \AA^{-1} since then the conservation of momentum condition in the IA implies that

$$\mathbf{p}_n - \mathbf{p}'_n = \hbar\mathbf{q} \quad (6.2.1)$$

where \mathbf{p}_n and \mathbf{p}'_n are the incident and scattered neutron momentum respectively.

The IA is expected to be valid if the collision time scale is much shorter than the period of the highest frequency vibration of the struck particle, since this time scale characterizes the interaction of the particle with its environment, the neglect of which is the main approximation which has been made. If we take the magnitude of the momentum transfer to be q , then this frequency is denoted by $\omega(q)$, and the IA

is satisfied when [194]

$$\frac{m\omega(q)}{q\Delta p} \ll 1 \quad (6.2.2)$$

where Δp is the width of the momentum distribution.

In practice, this usually requires wavenumbers of 20 \AA^{-1} or higher, but these energies are easily achievable using spallation neutron sources such as ISIS at the Rutherford Appleton Laboratory (UK) [195].

6.2.3 Momentum distribution calculation

Using the IA gives the nuclear structure factor, $S(\mathbf{q}, \omega)$, as [194, 197]

$$S(\mathbf{q}, \omega) = \frac{m}{\hbar q} J(y, \hat{\mathbf{q}}) \quad (6.2.3)$$

where $J(y, \hat{\mathbf{q}})$, the neutron Compton profile, is given by

$$J(y, \hat{\mathbf{q}}) = \hbar \int n(\mathbf{p}') \delta(\hbar y - \mathbf{p}' \cdot \hat{\mathbf{q}}) d\mathbf{p}' \quad (6.2.4)$$

where $\hat{\mathbf{q}} = \mathbf{q}/q$, $n(\mathbf{p})$ is the momentum distribution, $y = m(\hbar\omega - \hbar\omega_r)/\hbar^2 q$, $\hbar\omega$ is the energy of the incident neutron and $\hbar\omega_r = \hbar^2 q^2/2m$.

In liquids, or polycrystalline solid state materials, what is actually measured is the spherical average of the neutron Compton profile, which can be calculated using

$$J(y) = 2\pi\hbar \int_{|\hbar y|}^{\infty} p n(p) dp \quad (6.2.5)$$

This is a central quantity in the analysis of the data, and is formally the Radon transform of the momentum distribution. Despite being an average over momentum space, this contains the same amount of information as the momentum distribution itself. The inversion required to retrieve the momentum distribution is well known

[198], and given by

$$n(p) = -\frac{1}{2\pi\hbar^3} \frac{d}{dy} J(y)|_{\hbar y=p} \quad (6.2.6)$$

Theoretically then, the process used to calculate the momentum distribution from the scattering data is therefore as follows:

1. The scattering cross-section is calculated from the detected intensity of scattered neutrons as a function of their energy and a knowledge of the experimental set-up.
2. From the scattering cross-section we get the structure factor, and hence the neutron Compton profile.
3. The momentum distribution is calculated from the inverse Radon transform of the neutron Compton profile.

This however ignores some subtleties, which we will now discuss.

6.2.4 Beyond the impulse approximation

Firstly, it was assumed that the experimental set-up is known precisely, such that the distance between the incident beam and the detector, the angle that the beam is scattered by and the energy of the incident/reflected beam (whichever is being passed through the chopper or the absorption foil) are all known exactly.

However in practice there will be an uncertainty in each of these measurements, which will lead to an observed broadening of the spectrum. Therefore if this is not taken into account the calculated momentum distribution will contain a systematic statistical error.

Fortunately, these statistical uncertainties generally follow Gaussian statistics, apart from the energy which has a Voigt profile [199]. The final spectrum can therefore

be fitted to a convolution of the wanted spectrum and a known Voigt function, which we shall call $R(y, q)$, so the effect of these uncertainties can be minimized.

Secondly, while the impulse approximation is exact in the infinite incident momentum limit, with any finite value of q there will be some deviation from this idealized picture. This is usually treated using a final-state correction term, $\Delta J(y, q)$, derived from the leading order terms in the Taylor series with respect to $1/q$ and \mathbf{r} of $J(y, q)$ [200, 201].

These are due to the non-zero interaction of the struck particle with its neighbours during the collision, and so involve collective properties such as the potential derivatives with respect to the position of each particle, i.e. the forces and Hessian matrix. Fortunately, usually only one or two terms of this series are needed [194], and these can be calculated automatically from the experimental data.

Taking all this together then, instead of applying the inverse Radon transform directly to the experimental spectrum, $F(y, q)$, first final-state effects and experimental uncertainties are taken into account by fitting $J(y)$ to the following convolution

$$F(y, q) = [J(y) + \Delta J(y, q)] \otimes R(y, q) \quad (6.2.7)$$

Once this has been done, the calculated $J(y)$ will then be the best estimate for the neutron Compton profile, which is then used in equation 6.2.6 to calculate the momentum distribution.

6.2.5 Anisotropic Gaussian approximations

A general momentum distribution needs a basis which spans three-dimensional space to describe it completely. For example, the momentum distribution is often expanded in terms of Laguerre polynomials and spherical harmonics for single crystal samples

[202], since the coefficients in this series can be calculated from the form of the neutron Compton profile given in equation 6.2.4.

However, it is not always convenient to use this expansion. Firstly, many terms are often required in the summation, especially for strongly anisotropic distributions [194]. This means that a large number of coefficients must be determined to fit the data to this functional form accurately. Secondly, it is often difficult to extract this general distribution from PIMD simulations, as it must be calculated by either more expensive [39] or less reliable [141] techniques than required for properties such as the kinetic energy.

This is disappointing, since the kinetic energy is clearly closely related to the momentum distribution. In particular, the tensor form of the average kinetic energy, $\langle \hat{\mathbf{K}} \rangle$, can be obtained from the momentum covariance matrix by

$$\langle \hat{K}_{\alpha\beta} \rangle = \frac{\langle \hat{p}_\alpha \hat{p}_\beta \rangle}{2m} \quad (6.2.8)$$

where throughout this thesis x_α and $X_{\alpha\beta}$ are defined in the usual manner as components of the vector \mathbf{x} and the matrix \mathbf{X} respectively.

In general, there is no way to reconstruct the momentum distribution from this tensor. It is however possible in the special case of an anisotropic Gaussian distribution, defined as

$$n(\mathbf{p}) = \frac{1}{\sqrt{8\pi^3\sigma_x\sigma_y\sigma_z}} \exp\left(-\frac{p_x^2}{2\sigma_x^2} - \frac{p_y^2}{2\sigma_y^2} - \frac{p_z^2}{2\sigma_z^2}\right) \quad (6.2.9)$$

where each σ_α is the variance of the momentum distribution in the direction of the principle component p_α . This can also be spherically averaged to compare with

experiments of isotropic systems using

$$4\pi p^2 n(p) = \int \delta(p - |\mathbf{p}|) n(\mathbf{p}) d^3\mathbf{p} \quad (6.2.10)$$

Under the assumption that the momentum distribution can be written in this form, each p_α is an eigenvector of $\langle \hat{\mathbf{K}} \rangle$, with an associated eigenvalue $\sigma_\alpha^2/2m$.

This approximation is widely used, for three reasons. Firstly, it is very simple to fit this functional form to the experimental data since it only has three parameters. This prevents overfitting of the data; in almost all cases distinguishing any more complicated functional form from an anisotropic Gaussian is beyond the current capabilities of technology.

Secondly, these parameters have a simple physical interpretation: the amplitude associated with the principal components of the atomic motion. In the case of a molecular solid or liquid, these components are often closely related to the molecular vibrations. This is particularly helpful when studying CQE, as this gives an easy way to separate the contributions of the inter- and intramolecular motion.

Finally, as the matrix in equation 6.2.8 is simple to calculate during standard PIMD simulations, unlike the momentum distribution itself, this is the framework in which theoretical and experimental results can be compared most easily. How this is done in practice will be discussed in more detail in section 6.3.

While generally this is only an approximate form of the nuclear momentum distribution, it is exact in the harmonic limit. More importantly, it is accurate whenever the potential of mean force is effectively harmonic, as has been shown to be true in even the highly quantum case of low temperature ice [203]. This suggests that this approximation will be valid for many problems of interest, although it must be noted that this may not always be the case.

6.3 Theoretical techniques

6.3.1 Quasi-harmonic approximation

The simplest approximation that can include NQE in the calculation of momentum distributions is a quasi-harmonic analysis. This uses the normal mode frequencies of a zero-temperature optimized system to approximate the finite temperature kinetic energy distribution, from which we can reconstruct the momentum distribution by assuming it can be fitted to an anisotropic Gaussian form.

The first step of a quasi-harmonic calculation is to find the optimized geometry of the system, i.e. the energy minimum of the potential energy surface. There are several different algorithms which can be used for this purpose of varying degrees of sophistication, from simple conjugate gradient routines [204] to quasi-Newton methods such as the BFGS algorithm [205–209].

Regardless of how the geometry is optimized it is important that a sufficiently high degree of accuracy is achieved, since the second step of a harmonic analysis is determining the Hessian matrix of equation 3.1.5. This requires taking finite difference calculations of the forces acting on the particles (the Hessian matrix being the position derivative of the forces), and so is very sensitive to the exact point taken as the potential minimum. Usually, the nuclear positions should be converged to within a small fraction of the finite difference parameter.

Once the Hessian matrix, \mathbf{D} , has been constructed the next step is to solve the multi-dimensional analogue of equation 3.1.4. These normal modes, calculated from the Cartesian basis vectors $\hat{\mathbf{r}}_\alpha$ by $\hat{\mathbf{x}}_\nu = \sum_{\alpha=1}^{3N} \hat{\mathbf{r}}_\alpha C_{\alpha\nu}$, are then the solutions to

$$\mathbf{M}^{-1}\mathbf{D}\mathbf{M}^{-1}\hat{\mathbf{x}}_\nu = \omega_\nu^2 \hat{\mathbf{x}}_\nu \quad (6.3.1)$$

where $M_{\alpha\beta} = \sqrt{m_\alpha} \delta_{\alpha\beta}$, ω_ν are the normal mode frequencies and m_α is the mass

associated with the degree of freedom α .

For any crystal there should be six normal modes with zero frequency due to the symmetry of the Hamiltonian, corresponding to the translational and rotational degrees of freedom. However, the imposition of periodic boundary conditions lead to the breaking of spherical symmetry, so the rotational normal mode frequencies are not rigorously zero, though they should have a small magnitude.

Approximations used in the finite difference calculation will introduce some error into the values of the normal mode frequencies, and so in practice both the translations and rotations will have a non-zero frequency. This can therefore be used to test whether all the stages in the calculation have been sufficiently well converged. For the current problem, the six zero-frequency modes should behave classically at the temperature of interest, and have a frequency that is much smaller than any other normal mode.

Once a satisfactory level of convergence has been reached, the normal modes can then be used to calculate the momentum distribution. For a harmonic oscillator of frequency ω_ν the kinetic energy a particular quantum state n is set by the quantum virial theorem [95] to be

$$\langle n | \hat{K}_\nu | n \rangle = \frac{1}{2} \langle n | \hat{H}_\nu | n \rangle = \frac{(2n+1)\hbar\omega_\nu}{4} \quad (6.3.2)$$

The finite temperature canonical kinetic energy average for a particular normal mode ν can then be calculated using equation 2.1.4 to give

$$\begin{aligned} \langle \hat{K}_\nu \rangle &= \frac{\sum_{n=0}^{\infty} \langle n | \exp(-\beta \hat{H}_\nu) \hat{K}_\nu | n \rangle}{\sum_{n=0}^{\infty} \langle n | \exp(-\beta \hat{H}_\nu) | n \rangle} \\ &= \frac{\hbar\omega_\nu}{4 \tanh(\frac{\beta\hbar\omega_\nu}{2})} \end{aligned} \quad (6.3.3)$$

The six zero-frequency modes correspond to motions of the whole system, and so are treated classically in this analysis, i.e. each give a contribution of $kT/2$ to the total kinetic energy as required by the equipartition theorem.

Finally, we use this information to construct the kinetic energy tensor of equation 6.2.8. To do this, we note that each Cartesian component can be decomposed into normal modes by inverting the original transformation, using $\hat{\mathbf{r}}_\alpha = \sum_{\nu=1}^{3N} \hat{\mathbf{x}}_\nu C_{\alpha\nu}$. This then gives [210]

$$\begin{aligned} \langle \hat{K}_{\alpha\beta} \rangle &= \sum_{\nu=1}^{3N} C_{\alpha\nu} C_{\beta\nu} \langle \hat{K}_\nu \rangle \\ &= \sum_{\nu=1}^{3N} C_{\alpha\nu} C_{\beta\nu} \frac{\hbar\omega_\nu}{4 \tanh\left(\frac{\beta\hbar\omega_\nu}{2}\right)} \end{aligned} \quad (6.3.4)$$

where the translational and rotational degrees of freedom are set to $\omega_\nu = 0$, so that $\langle \hat{K}_\nu \rangle = kT/2$.

This then gives a $3N \times 3N$ matrix giving the kinetic energy tensor for all the degrees of freedom in the system. The 3×3 matrix corresponding to each individual nucleus is then taken separately, as we are only interested in the limit of incoherent scattering in which the off-diagonal terms involving correlations between atoms are ignored.

The anisotropic Gaussian approximation to the momentum distribution for any given nucleus can then be calculated by identifying the eigenvectors of the resultant 3×3 matrix with p_α and its eigenvalues with $\sigma_\alpha^2/2m$ in equation 6.2.9.

6.3.2 Open-path integral methods

While the quasi-harmonic approximation is a simple and fast method for calculating an approximation to the momentum distribution of a low-temperature solid, it does not take into account the affect of anharmonicity in the potential and so is unsuitable

for use in liquid phase systems. Similarly, this approximation is unlikely to describe the behaviour of molecular systems very well, since the intermolecular bonding is often strongly anharmonic even at low temperatures.

To move beyond a purely harmonic analysis, we use path-integral molecular dynamics to do calculations at finite temperature. There are several methods which can be used for this purpose, of varying sophistication, some of which will now be discussed.

Open-path path integral molecular dynamics

At the highest level of theory there are open-path integral molecular dynamics simulations [40, 211], which can be used to calculate the momentum distribution directly. This differs from a standard closed-path integral calculation in that we no longer use cyclical boundary conditions in the imaginary time path integral. In terms of the simulation implementation, one of the ring polymers is “opened”, such that the polymer is no longer a ring but a chain.

This stems from the definition of the momentum distribution as a Fourier transform [211] over the end-to-end distribution, $n(\mathbf{x})$, given by

$$n(\mathbf{p}) = \frac{1}{(2\pi\hbar)^3} \int d\mathbf{x} \exp\left(\frac{i\mathbf{p} \cdot \mathbf{x}}{\hbar}\right) n(\mathbf{x}) \quad (6.3.5)$$

where $n(\mathbf{x}) = \frac{1}{Z} \int d\mathbf{r} d\mathbf{r}' \delta(\mathbf{r} - \mathbf{r}' - \mathbf{x}) \rho_n(\mathbf{r}, \mathbf{r}')$, $Z = \int d\mathbf{r} \rho_n(\mathbf{r}, \mathbf{r})$ and $\rho_n(\mathbf{r}, \mathbf{r}')$ is the density matrix.

In path integral form the single-particle density matrix for particle 1 can be calculated using

$$\rho_n(\mathbf{r}, \mathbf{r}') = \int d\mathbf{r}_2 \dots d\mathbf{r}_n \exp(-\beta U_n) \quad (6.3.6)$$

where

$$U_n = \frac{1}{2} \left[U(\mathbf{r}, \mathbf{r}_2^{(0)}, \dots, \mathbf{r}_N^{(0)}) + U(\mathbf{r}', \mathbf{r}_2^{(n)}, \dots, \mathbf{r}_N^{(n)}) \right] + \sum_{i=1}^{n-1} U(\mathbf{r}_1^{(i)}, \mathbf{r}_2^{(i)}, \dots, \mathbf{r}_N^{(i)}) + \sum_{j=1}^N \sum_{i=0}^{n-1} \left(\frac{n^2 m_j}{2\hbar^2 \beta^2} |\mathbf{r}_j^{(i)} - \mathbf{r}_j^{(i+1)}|^2 \right) \quad (6.3.7)$$

Under this formalism, the end-to-end distribution is then simply

$$n(\mathbf{x}) = \lim_{n \rightarrow \infty} (\langle \delta(\mathbf{r} - \mathbf{r}' - \mathbf{x}) \rangle^{\text{OPI}}) \quad (6.3.8)$$

where the superscript OPI means that the ensemble average is taken over the open-path integral Boltzmann density.

While this method is in principle exact (under the same approximations used throughout this thesis for closed PIMD simulations), it does have a few significant drawbacks, which have prevented its widespread use.

Firstly, the nature of the technique is that it is only exact in the case where only one path is opened at any given time. This means that it is not possible to average over all the equivalent nuclei (which we will refer to as horizontal statistical averaging), and so long simulation times are often necessary to get converged results. While it has been found that equivalent nuclei on different molecules can be opened at the same time in some systems without significant loss of accuracy [211], this is by no means guaranteed to always be the case.

Secondly, since the size of the opening between the two ends of the ring polymer is given by random walk statistics, the probability of a small ring opening of magnitude x occurring is proportional to x^2 [141]. This means that the $x \rightarrow 0$ limit of the end-to-end distribution tends to have even poorer statistics, and so cannot be converged in any reasonable time frame.

Displaced path method

To try and address the problems with the open-path method, a closed-path formulation of the same estimator was developed [141]. This uses a change of variables to take equation 6.3.8 and transform it into an average over a closed ring polymer, which results in

$$\begin{aligned}
 n(\mathbf{x}) &= \lim_{n \rightarrow \infty} \left(\langle \delta(\mathbf{r} - \mathbf{r}' - \mathbf{x}) \rangle^{\text{OPI}} \right) \\
 &= \lim_{n \rightarrow \infty} \left(\left\langle \exp \left(\frac{\beta}{n} \sum_{i=0}^{n-1} U \left(\mathbf{r}_1^{(i)} + y_i \mathbf{x}, \mathbf{r}_2^{(i)}, \dots, \mathbf{r}_N^{(i)} \right) \right. \right. \\
 &\quad \left. \left. - \frac{\beta}{n} \sum_{i=0}^{n-1} U \left(\mathbf{r}_1^{(i)}, \mathbf{r}_2^{(i)}, \dots, \mathbf{r}_N^{(i)} \right) - \frac{m_1 \mathbf{x} \cdot \mathbf{x}}{2\beta \hbar^2} \right) \right\rangle^{\text{PI}} \right)
 \end{aligned} \tag{6.3.9}$$

where $y_i = \frac{1}{2} - \frac{i}{n} - \frac{1}{2n}$

This is mathematically equivalent to the open-path formulation, and solves the two problems which were raised in the previous section; namely that the open-path average has poor horizontal statistics and is very difficult to converge in the limit as $x \rightarrow 0$.

In the displaced path formulation, the amount, \mathbf{x} , by which the path is opened is chosen by the user, so there is no sampling issue for small values of x . Similarly, as the above estimator uses the closed-path Hamiltonian, full use of horizontal statistics can be achieved with no loss of accuracy, although each calculation does require one more potential evaluation to be done for each path that is opened.

This does, however, raise its own statistical problems. As the estimator in equation 6.3.9 is an exponentiated quantity, there are configurations that are very unlikely to be visited as they have a very small Boltzmann weight in the closed-path phase space, which nevertheless have a large weight in the average due to the exponential of the potential energy difference. This can give rise to very poor statistical con-

vergence, as the correct average is only obtained once exponentially unlikely regions of phase space are explored. In fact it can be shown [67] that the variance of the estimator increases exponentially below a certain critical temperature; this method is therefore unfeasible for highly quantum problems.

6.3.3 Anisotropic Gaussian methods

The previous two methods both have statistical drawbacks that make them difficult to use in many problems of interest. However, as discussed in section 6.2.5, momentum distributions can be calculated with only a knowledge of the kinetic energy tensor if just one additional approximation is made, i.e. that the momentum distribution can be represented as an anisotropic Gaussian distribution.

In a PIMD simulation, the centroid virial estimator for the kinetic energy tensor of the j -th particle, a generalization of equation 2.1.14, is given by

$$\langle \hat{K}_{\alpha\beta}^j \rangle = \lim_{n \rightarrow \infty} \left(\langle K_{\alpha\beta,n}^j \rangle^{\text{PI}} \right) \quad (6.3.10)$$

where

$$K_{\alpha\beta,n}^j = \frac{3\delta_{\alpha\beta}}{2\beta} + \frac{1}{4n} \sum_{i=0}^{n-1} \left[\left(r_{j\alpha}^{(i)} - r_{j\alpha}^c \right) \frac{\partial U(\mathbf{r}^{(i)})}{\partial r_{j\beta}^{(i)}} + \left(r_{j\beta}^{(i)} - r_{j\beta}^c \right) \frac{\partial U(\mathbf{r}^{(i)})}{\partial r_{j\alpha}^{(i)}} \right] \quad (6.3.11)$$

In theory, the anisotropic Gaussian approximation to the momentum distribution of the j -th particle can be constructed by calculating the eigenvectors and eigenvalues of $\langle K_{\alpha\beta,n}^j \rangle^{\text{PI}}$. However, there is one major complication that must be addressed before this method can be applied to liquid systems, where it is expected to offer the most consistent improvement over harmonic approximations.

Instantaneous values of any estimator are not necessarily meaningful in a PIMD simulation, and so they must be averaged over many configurations. Unfortunately,

in a liquid system, where the molecular frame of reference changes over the course of a simulation due to rotational diffusion, these averages will be isotropic if done naïvely. To get around this issue several techniques have been developed, and these will now be discussed.

Mean square displacement method

The main issue that causes time averages to become isotropic in liquids is that the nuclei do not have a fixed reference frame. However, in some cases there is an internal reference frame which can be used instead. In particular, molecular systems have a consistent geometry which can be used to relate configurations at different times.

This is done in three steps. Firstly, for each molecule a rotation matrix is found which in some sense minimizes the difference between its configuration and that of a reference molecule. Secondly, this rotation matrix is applied to the kinetic energy tensor of each nucleus so that every time step can be treated on a level footing. Finally, the rotated tensors are averaged over time and over equivalent nuclei, and the eigenvectors and eigenvalues of the resultant matrix used to give the parameters needed to calculate the momentum distribution. This is known as the mean square displacement (MSD) method.

The first step, finding the optimum rotation matrix, is usually done by least squares fitting. Let \mathbf{r}_j be the instantaneous position and \mathbf{R}_j the reference position of atom j , where $j = (1, \dots, M)$ and M is the number of atoms per molecule. Both the instantaneous and reference molecules are then translated such that their centroid is at the origin, so that only rotations need to be considered. Once this is done, the best choice of matrix for this method, which we will call \mathbf{U} , is then the one which

minimizes the fitness function

$$E = \sum_{j=1}^M |\mathbf{U}\mathbf{r}_j - \mathbf{R}_j|^2 \quad (6.3.12)$$

under the constraint that it is unitary, i.e.

$$\mathbf{U}^T \mathbf{U} = \mathbf{I} \quad (6.3.13)$$

There are several ways of solving this set of equations, but in this thesis the Kabsch algorithm [212, 213] was used. This enforces the unitary constraint using Lagrange multipliers, and then finds the optimal rotation matrix by solving the resulting set of equations analytically.

The solution can be written in several ways: the original paper used the eigenvectors of the covariance matrix, $C_{\alpha\beta} = \sum_{j=1}^M R_{j\alpha} R_{j\beta}$, as a basis for \mathbf{U} [212], but a more numerically stable method is to perform a singular value decomposition (SVD) [214] on this covariance matrix, giving

$$\mathbf{C} = \mathbf{L}\mathbf{\Sigma}\mathbf{R}^T \quad (6.3.14)$$

where \mathbf{L} and \mathbf{R} are unitary and $\mathbf{\Sigma}$ is a diagonal matrix. From this it can be shown that \mathbf{U} is given by

$$\mathbf{U} = \mathbf{L}\mathbf{R}^T \quad (6.3.15)$$

Now that the optimum rotation matrix has been found, it can be used to transform the kinetic energy tensor for each atom in the molecule to the reference frame. If the instantaneous value of the tensor defined in 6.3.11 is \mathbf{k}_n^j , then the change of basis is applied using

$$\mathbf{K}_n^j = \mathbf{U}\mathbf{k}_n^j\mathbf{U}^T \quad (6.3.16)$$

Finally, \mathbf{K}_n^j is averaged over equivalent nuclei and over time. The eigenvalues and eigenvectors of the resulting tensor are then used to calculate the momentum distribution, using the procedure outlined in section 6.2.5.

One last complication that needs to be addressed in the particular case of ammonia is that this molecule has two inversion isomers, differing by a permutation of hydrogen nuclear labels. While one isomer cannot be mapped onto the other by a rigid rotation, during a simulation it is possible for the two to interconvert since the nitrogen electron lone pair can tunnel through the low energy barrier which separates them.


The above method is capable of dealing with this problem, since an improper rotation (a combination of a rotation and a reflection) is still a unitary transformation. However, most implementations of the Kabsch algorithm will only allow \mathbf{U} to be a proper rotation by default, using

$$\mathbf{U} = \mathbf{L}\mathbf{A}\mathbf{R} \tag{6.3.17}$$

where \mathbf{A} is either the identity matrix, or has the element corresponding to the smallest eigenvalue of \mathbf{C} replaced by -1 if \mathbf{U} would otherwise not be positive definite. It can be shown that this definition of \mathbf{U} gives the minimum possible value of E under the dual constraints of unitarity and molecular rigidity [213].

However, in the case of ammonia this procedure is not necessary, since the order in which the hydrogens are labelled does not mean anything physically. In fact, if the instantaneous geometry and the reference geometry happen to correspond to different inversion isomers, as in figure 6.1, then an improper rotation will be required to give the best fit. As such, either \mathbf{A} should not be included in the algorithm, or the results of both \mathbf{r} and $-\mathbf{r}$, the latter of which corresponds to an inversion of the instantaneous positions, should be compared and only the \mathbf{U} matrix which gives the lowest value of E used.

Figure 6.1: A diagram to show the two inversion isomers of the ammonia molecule. Assuming that the nitrogen atom is above the plane of the paper for both, we see that there is no rigid rotation that can interconvert them. Therefore, an improper rotation is required to get a correct MSD fit in the case where the instantaneous geometry corresponds to a different isomer than the reference molecule.



As long as distortions of the molecular geometry from the reference geometry are small this method will give a reliable estimate of the anisotropic kinetic energy tensor even for liquid phase systems. On top of this, the way the average is constructed means that the orientation of the eigenvectors can be directly related to the molecular axes, and so they can easily be given a physically useful interpretation.

While these are useful properties of the MSD technique, they also suggest that it will likely fail whenever the structure cannot be described by a fixed molecular geometry. In this case a more general method is required, one that does not require an internal reference frame.

Transient anisotropic Gaussian method

The key observation needed to calculate anisotropic averages without requiring a reference geometry is that the averaged kinetic energy tensor only becomes isotropic on timescales of the order of the rotational diffusion time. Therefore if a time-averaging window is chosen that is small enough to leave the environment around

each nucleus unchanged, but still large enough to give sufficiently good statistics, then we can get an average of the kinetic energy tensor that retains the anisotropy without needing to map the system on to a reference structure.

For each time step, a moving average is taken over a time window $2\Delta t$, giving

$$\bar{K}_{\alpha\beta,n}^j(t) = \frac{1}{\Delta t} \int_{t-\Delta t}^{t+\Delta t} K_{\alpha\beta,n}^j(t') \left(1 - \frac{|t-t'|}{\Delta t}\right) dt' \quad (6.3.18)$$

The eigenvalues of this matrix, which we will label θ_α^j , are then averaged over each time step, t , to give the parameters to be used in equation 6.2.9, i.e. $\langle \theta_\alpha^j \rangle = (\sigma_\alpha^j)^2/2m_j$. This is known as the transient anisotropic Gaussian (TAG) method [59].


While there is no *a priori* method of determining the optimal window parameter, Δt , it is usually found that there are three distinct regimes of behaviour. For small window parameters, it is found that the eigenvalues $\langle \theta_\alpha^j \rangle$ “repel” one another, i.e. the largest is found to be larger than for the converged result and vice-versa.

The reason for this can be seen from figure 6.2. In this case, even though the red distribution has a larger average than the blue distribution, for the given sample the instantaneous values are the other way around. Since it is not known *a priori* which distribution each eigenvalue corresponds to, the eigenvalues can only be ordered by size, and so this means that these points would be assigned incorrectly. The calculated values of the averages of these distributions will then differ from the correct ones, with the smaller value being too small and vice-versa.

This will be seen whenever the tails of two distributions overlap significantly, and so is particularly problematic whenever the eigenvalues are close together, or when their distributions are very broad. This means that for small Δt , where the variance in the distribution of each of the eigenvalues $\langle \theta_\alpha^j \rangle$ is large, their averages will be systematically wrong.

At the other extreme, in liquid systems $\bar{\mathbf{K}}_n^j$ will be spherically averaged in the

Figure 6.2: A plot to show the origin of the ordering problem in TAG and TNP when the distributions of the eigenvalues of $\bar{\mathbf{K}}_n^j$ show noticeable overlap. In this example, the red point will be assigned to the blue distribution and vice-versa, leading to a “repulsion” of the observed mean for each distribution.



limit as $\Delta t \rightarrow \infty$ due to rotational diffusion. As such, an isotropic tensor will be obtained, and so each of the eigenvalues $\langle \theta_\alpha^j \rangle$ will be identical.

Between these two limits there is usually a third regime where there is a plateau in each of the eigenvalues $\langle \theta_\alpha^j \rangle$ with respect to Δt . These plateau values are the ones which are used in equation 6.2.9 to give the TAG approximation to the momentum distribution.

Here we note two things. Firstly, the trade-off that has to be made to remove the requirement of a well-defined internal reference frame is that the TAG method cannot be used to determine physically meaningful eigenvectors, unlike the MSD method. Secondly, if there is no plateau in the eigenvalues $\langle \theta_\alpha^j \rangle$ between the timescale of diffusion and the time needed to sufficiently converge $\bar{\mathbf{K}}_n^j$ the TAG method will not give well-defined results.

Transient $n(p)$ method

To try to remove this last limitation of the TAG method, we need to average the data in such a way that which of the eigenvalues θ_α^j correspond to which of the three distributions is not important, so that smaller values of Δt can be used. One possible way this could be done is to calculate the spherically averaged momentum distribution, as defined in equation 6.2.10, at each time step and average this instead of the eigenvalues themselves. The shape of this distribution is independent of the ordering of the eigenvalues, and the spherical averaging can be inverted to return the values of σ_α^j and θ_α^j , if required. In analogy with the transient anisotropic Gaussian method we will call this the transient momentum distribution (TNP) method.


The first step is the same as for the TAG method: calculating the eigenvalues θ_α^j from equation 6.3.18. Next, these eigenvalues are used to calculate the instantaneous spherically averaged momentum distribution estimator, which we will label $\eta^j(p)$, using the procedure outlined in section 6.2.5. Finally, the best approximation to $n(p)$ is found from $\langle \eta^j(p) \rangle$ by fitting the parameters in equation 6.2.9 to minimize the quantity

$$E = \int dp |\langle \eta^j(p) \rangle - n(p)|^2 \quad (6.3.19)$$

The main problem with the TNP method is that, as will be demonstrated later, the shape of the $n(p)$ distribution is fairly insensitive to the individual components of the anisotropic momentum distribution, and is mostly determined by the total kinetic energy. It is therefore possible that the final step in the calculation, fitting the data to a spherically averaged $n(p)$, will be very sensitive to any statistical noise in the data, making the results unreliable. Whether or not this method is an improvement over the TAG method will be tested in the next section.

6.4 Application to liquid and solid ammonia

Figure 6.3: A representation of the unit cell of solid ammonia I, which is in the $P2_13$ space group. This was created using the VESTA crystal structure software [215].



To test these different theoretical and experimental techniques, we chose to look at the system of solid and liquid ammonia. We expect NQE to be important for ammonia for the same reasons that they are in water; it is a condensed phase molecular system which has both bonds with a large ZPE and a highly structured hydrogen bond network. It is therefore a good testing ground for the theoretical techniques given last section.

Ammonia is also an important system to study in its own right for several reasons.

It is commonly used as a solvent for chemical reactions; in particular alkali metals will dissolve in ammonia to give metal cations and solvated electrons. This solution is very interesting from a research perspective, both because it acts as a powerful and selective reducing agent [216–221], and since it has an unusual electronic structure and shows a range of conductive and magnetic behaviour [222].

Furthermore, ammonia is present in high concentrations in the atmosphere of the gas giant planets [11]. It is difficult to study this system experimentally, for obvious reasons, and so simulating how the bonding and properties of high-pressure ammonia are influenced by NQE is an important problem in astrophysics.

Finally, DINS measurements of ammonia were being made at the Rutherford Appleton Laboratory, Oxford, at the same time as our calculations by Andrew Seel with Carla Andreani (see figure 6.14). As such it was an excellent system to use to compare different theoretical approaches for calculating momentum distributions with experiment.

Both solid and liquid ammonia are studied here, for two reasons. Firstly, this will provide the most stringent tests of the theoretical methods, as they will have to work both in the limit where there are large quantum effects and when the average kinetic energy tensor is isotropic.

Secondly, by observing how the momentum distribution changes upon melting, this will provide indirect evidence for whether or not CQE are important to the properties of ammonia. In particular, the logic of section 6.1 suggests that the magnitude of the isotope effect on the melting point of ammonia will be reduced if there is a significant redistribution of energy between intra- and intermolecular motion upon melting.

Solid ammonia I, the equilibrium phase at atmospheric pressure for all temperatures considered and the only phase of solid ammonia studied in this thesis, is a

cubic crystal in the $P2_13$ space group [223–226]. Figure 6.3 shows the unit cell, which contains four ammonia molecules, and figure 6.4 shows the view down the principal C_3 axis and the environment around a representative ammonia molecule, each of the molecules in ammonia I being equivalent due to symmetry.

Importantly, we see that each nitrogen lone pair must form bonds with three different hydrogen atoms due to geometric constraints. This means that each hydrogen bond will be considerably weaker than the water analogue, and so we expect the hydrogen bonding network to exert a smaller influence on the dynamics of the system.


6.4.1 Computational details

The electronic structure calculations were done using the Quantum Espresso DFT package [133], which uses plane wave basis sets to calculate the electronic wavefunction. The relative performance of several different exchange-correlation functionals were examined by comparing the equation of state produced for phase I ammonia by each of them with experimental data (see table 6.1). For all calculations we used Van-derbilt pseudopotentials [228] and a 35 hartree plane wave cutoff energy, and unless stated otherwise k -point sampling, used to calculate reciprocal space integrals over

Method	Equilibrium volume / $\text{cm}^3 \text{ mol}^{-1}$	Bulk modulus / GPa
LDA	15.8	15.8
PBE	21.6	6.1
PW	21.5	5.6
BLYP	24.3	3.9
Experiment [225]	19.365	5.8

Table 6.1: A table to show the relative performances of three common exchange-correlation functionals for optimized structure calculations of ammonia phase I. The abbreviations used are local density approximation (LDA) [118], Perdew-Burke-Ernzerhof (PBE) [119, 120], Perdew-Wang (PW) [227] and Becke-Lee-Yang-Parr (BLYP) [121, 122]. The experimental volume was obtained at 2 K, and the bulk modulus at 195 K.

Figure 6.4: Plot to show a representation of the momentum distribution of nitrogen (N) and hydrogen (H) in solid ammonia at 50 K, and the extent of the hydrogen bonding network. Positions were taken from the experimentally determined equilibrium configuration at 50 K. Hydrogen bonding is shown by a dotted line from the hydrogen in the direction of the NH bond. The ellipsoids centred on each atom (yellow corresponding to hydrogen, blue to nitrogen) show the components of the momentum distribution obtained from the MSD method. The major axes of each ellipsoid show the principle components of the kinetic energy tensor as determined by the Kabsch method, with their magnitudes being proportional to the corresponding components of the momentum distribution. The first plot gives a side-on view, the second looks down the C_3 symmetry axis.



the vibrational Brillouin zone, was done on a $4 \times 4 \times 4$ Monkhorst-Pack grid [229].

By performing a structure and energy minimization at several different simulation cell volumes (in the same way as described in section 6.3.1), we obtained energy vs volume data which could be fitted to a Birch-Murnaghan equation of state [156]. As the parameters used in this equation have physical interpretations, including equilibrium volume and bulk modulus, the value of these parameters which gives the best fit to the data also gives the best estimate for the corresponding observable property. The best agreement with experiment was found with the PW and PBE exchange-correlation functionals, which were more or less equivalent. As such, we decided to do all further calculations with the PBE functional.

QHA results were obtained from static calculations run using a single unit cell of 4 ammonia molecules. Convergence tests showed that the parameters used above were sufficient to converge the normal mode frequencies to within 15 cm^{-1} .

MSD, TAG and TNP data were calculated from PIMD simulations. The i-PI package of chapter 4 was used to propagate the NVT dynamics, and the PIGLET method of section 2.2 was used to increase the rate of convergence of the kinetic energy with respect to the number of beads. Thermostat parameters were obtained from ref. [139], with the value of the frequency scaling parameter, $\beta\hbar\omega$, chosen such that all the physical vibrations of the system would be sampled correctly. This happened to be $\beta\hbar\omega = 100$ for $T = 50 \text{ K}$, and $\beta\hbar\omega = 50$ for $T = 120, 180$ and 220 K . These temperatures were chosen so that our calculations would be at the same state points as the experiments being run concurrently.

All PIMD simulations were performed using 32 ammonia molecules at the experimental density [230–234] (see table 6.2), and periodic boundary conditions were used throughout. At temperatures for which no experimental data could be found, linear interpolation between data for the two nearest state points was used to estimate the

Temperature	Density / kg m^{-3}	Simulation box side length / a_0
50 K	875.8	19.105
120 K	848.8	19.305
180 K	826.3	19.479
220 K	705.7	20.531

Table 6.2: Basic thermodynamic data for the state points generated by our simulations.

density. It was found that only gamma point sampling was required to converge the reciprocal space integrals over the vibrational Brillouin zone.

The path integral dynamical algorithms were integrated using a 0.5 fs timestep and the trajectories were run for 10 ps. This proved to be sufficiently long to reduce statistical errors in the total kinetic energy to less than 1 meV, much less than the expected error due to the exchange-correlation functional. Simulations were run at 50 K with 4, 8, 12 and 16 replicas to converge the results with respect to n . From these results it was decided that 12 beads gave sufficiently good convergence (the total kinetic energy being within 2% of that obtained with 16 beads), and as such this number of beads was used at all other temperatures.

6.4.2 Results

The central property by which we will judge the performance of the various theoretical approaches is the momentum distribution itself, since this is the only quantity which can be reliably obtained experimentally. As such the momentum distributions given by each method over the different temperatures studied are compared in figures 6.5 to 6.8.

Firstly, here we see that, unsurprisingly, the classical Maxwell-Boltzmann distribution is hopelessly wrong at all the temperatures considered. This underlines the fact that NQE are very important and must be included when studying the dynamics of systems containing light atoms such as hydrogen at all but the highest

Figure 6.5: Plots showing the calculated spherically averaged $n(p)$ at 50 K for N and H respectively for a variety of methods.





Figure 6.6: As figure 6.5, but at 120 K.


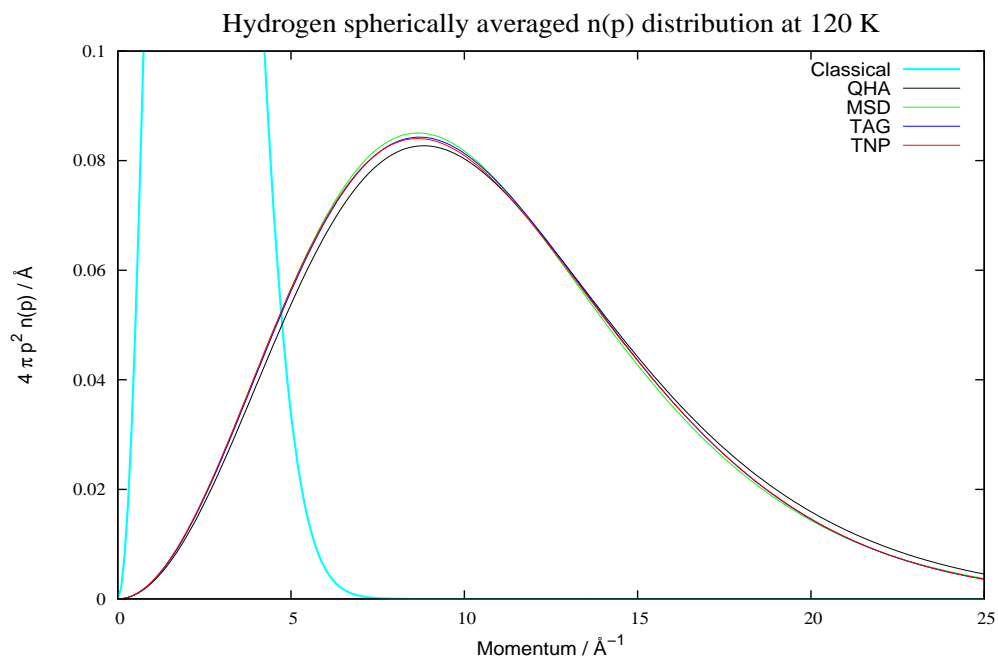


Figure 6.7: As figure 6.5, but at 180 K.

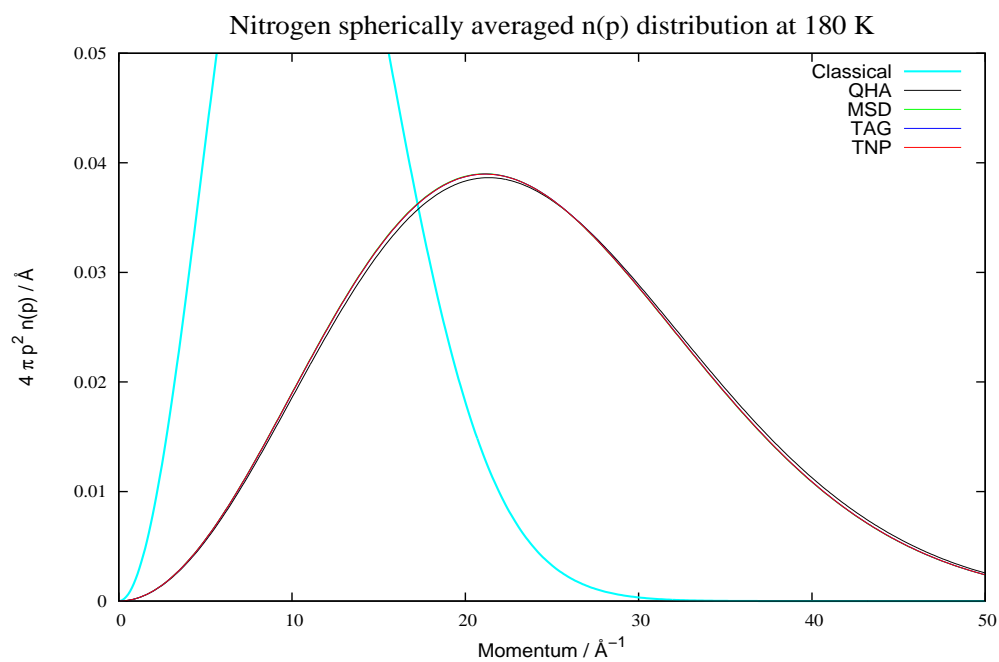

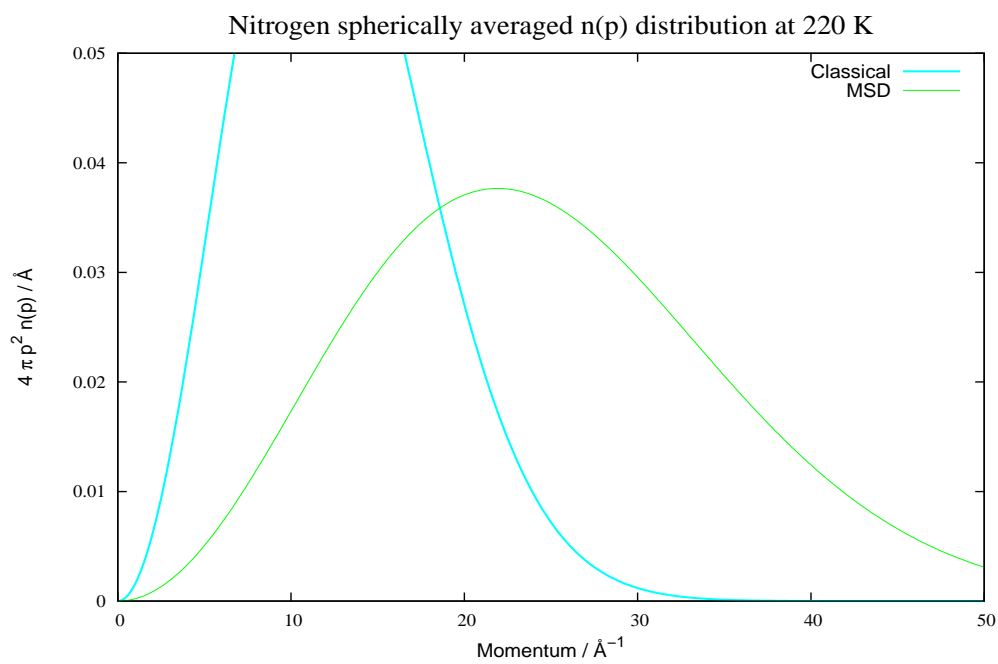
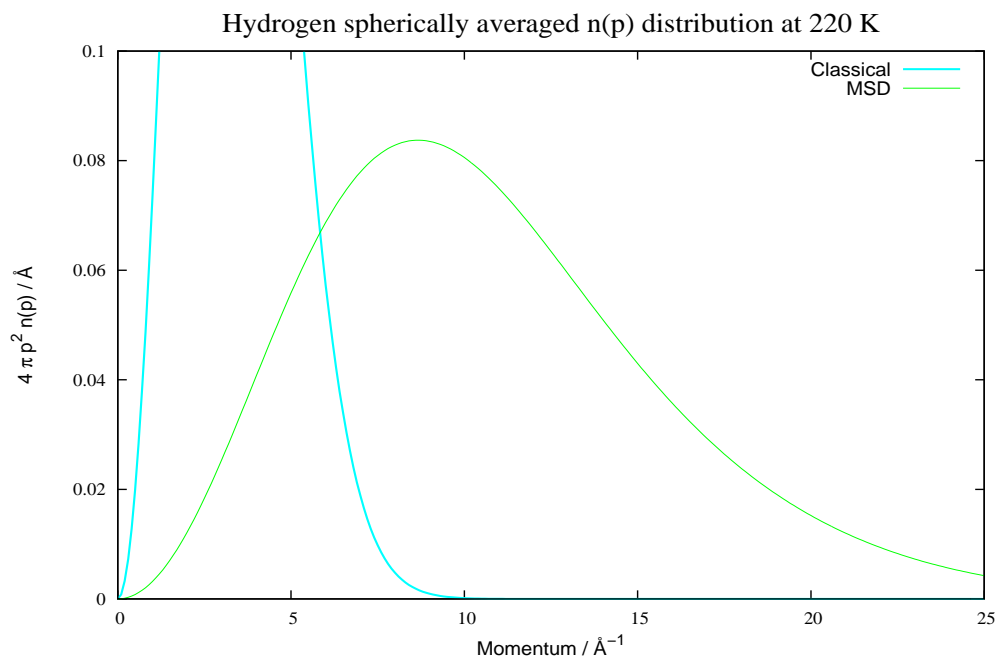


Figure 6.8: As figure 6.5, but at 220 K.



temperatures.

Secondly, we see that, while the various quantum mechanical methods give very similar results for solid ammonia, there is a much better agreement between the path integral methods than between these and the QHA. This discrepancy, while small, is too large to be explained by statistical uncertainty, and increases as the temperature decreases. This suggests that anharmonic effects, which are neglected by the QHA, are small but significant in this system, and especially at low temperatures.

To explain this, consider the detailed breakdown of the results as given in tables 6.3 and 6.4. Firstly, we note that the total kinetic energy given by the MSD, TAG and TNP methods are all in agreement, which is unsurprising given that they are all using the same data. In effect this simply shows that there is no significant numerical instability in any of these algorithms.

The QHA on the other hand overestimates the kinetic energy, by an amount that decreases with increasing temperature and that is especially large for hydrogen. In fact, this trend alone is sufficient to qualitatively explain the discrepancies seen between the momentum distributions derived from the QHA and the PIMD data.

This is consistent with the earlier observation that the total kinetic energy is usually a much more important quantity in determining the shape of the spherically averaged momentum distribution than the exact values of the kinetic energy components. Since this is the only property that can be reliably determined experimentally, this suggests that the QHA will not give the best comparison with experiment, regardless of whether or not it properly describes the kinetic energy tensor eigenvalues.

However, the correct description of the kinetic energy components is important to the discussion of CQE, as these explicitly depend on how the kinetic energy of different phases is partitioned between the inter- and intramolecular modes of motion. To study this, the trend with respect to temperature of each of the kinetic energy

Temperature	Method	θ_x/meV	θ_y/meV	θ_z/meV	$\frac{\theta_y+\theta_z}{2}/\text{meV}$	$\langle K \rangle/\text{meV}$
50 K (solid)	QHA	10.7	15.3	15.4	15.4	41.4
	MSD	10.4	14.3	14.6	14.4	39.3
	TAG	10.4	13.9	15.0	14.5	39.3
	TNP	10.2	14.0	15.0	14.5	39.2
120 K (solid)	QHA	12.2	16.9	17.0	16.9	46.0
	MSD	11.8	16.3	16.4	16.3	44.4
	TAG	11.9	16.0	16.5	16.3	44.4
	TNP	11.8	16.0	16.6	16.3	44.4
180 K (solid)	QHA	14.0	18.8	18.9	18.8	51.6
	MSD	13.6	18.4	18.5	18.5	50.6
	TAG	13.8	18.1	18.8	18.4	50.6
	TNP	13.7	18.1	18.8	18.5	50.6
220 K (liquid)	MSD	15.1	19.6	19.7	19.6	54.3

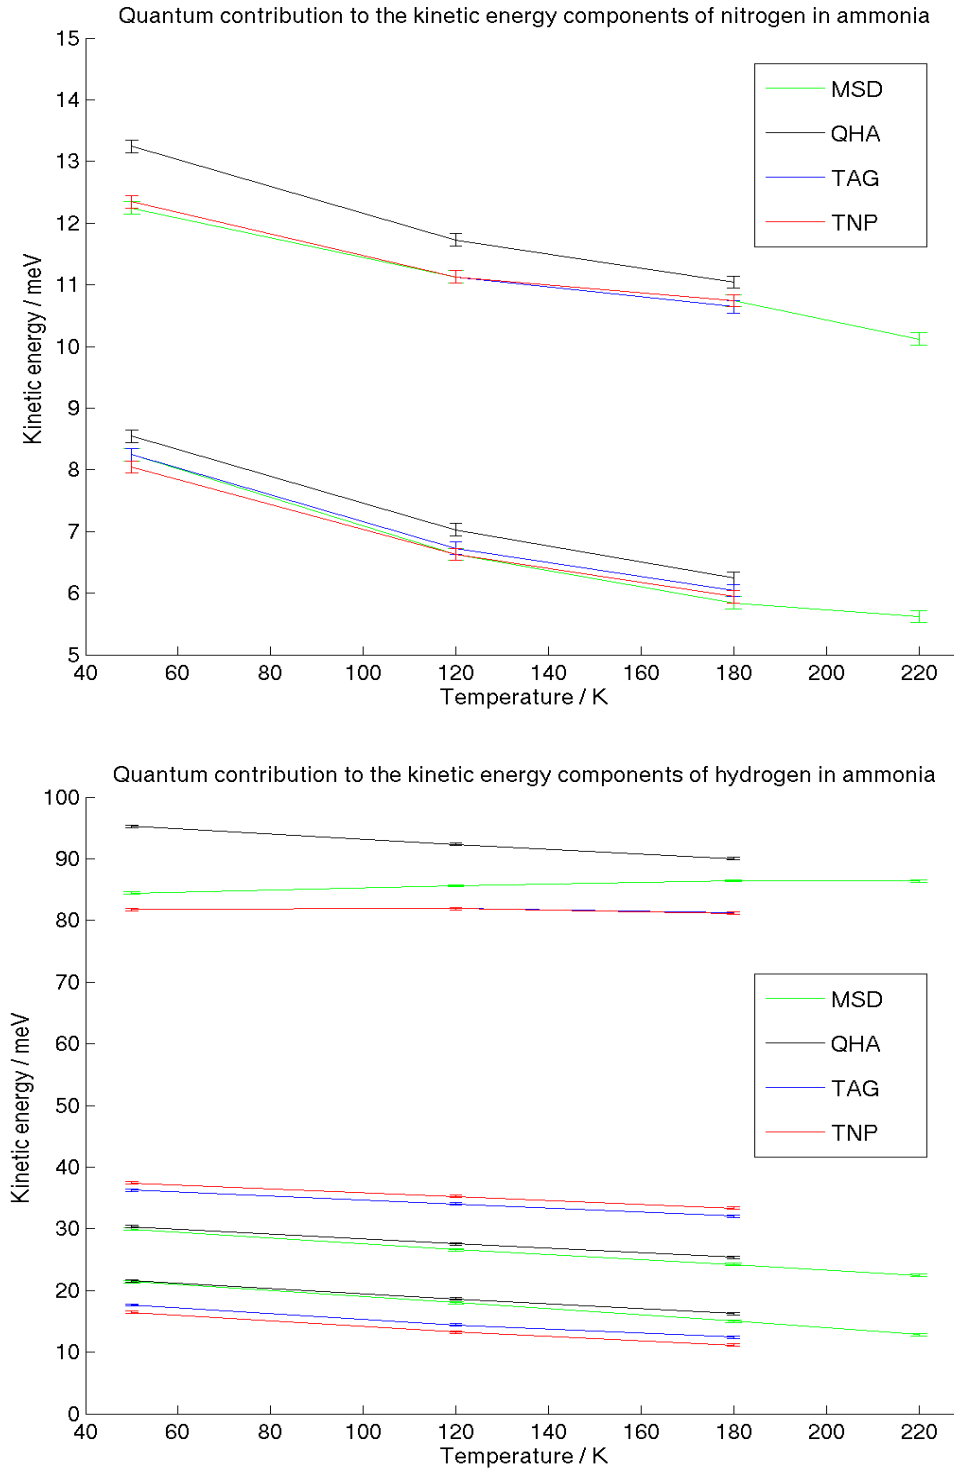
Table 6.3: Calculated principal components of the nitrogen kinetic energy tensor and total kinetic energy at all temperatures studied and for all methods. Also shown is the average of the estimates of the degenerate kinetic energy components, which is therefore the best estimate for both components. The TAG and TNP approximations have been omitted from the liquid data as no plateau was ever reached.

eigenvalues is shown in figure 6.9.

Firstly, we see that the QHA compares well with the molecular dynamics methods (especially MSD) at low frequencies, but gets both the average and the trend with respect to temperature wrong for the high frequency kinetic energy eigenvalue of hydrogen. This suggests that the leading order anharmonic contribution is due to the ZPE in the NH bond, for two reasons: anharmonic potential wells have a lower ZPE compared to harmonic wells with the same vibrational frequency, explaining why the QHA gives an overestimate of the total kinetic energy, and ZPE is expected to be most significant for the highest frequency kinetic energy tensor eigenvalues, and at lower temperatures.

This means that the QHA is unlikely to be able to describe CQE or KIE correctly, since these are very sensitive to the ZPE difference between each phase being considered; for example, fractionation ratios of two different isotopes are exponentially dependent on this energy difference. Since our motivation for studying this problem

Figure 6.9: Plots showing how the eigenvalues of the kinetic energy tensor vary with temperature for N and H respectively for a variety of methods. The classical contribution of $k_B T/2$ has been subtracted from each value. The average of the two degenerate eigenvalues are given for N.



Temperature	Method	θ_x/meV	θ_y/meV	θ_z/meV	$\langle K \rangle/\text{meV}$
50 K (solid)	QHA	23.7	32.5	97.4	153.6
	MSD	23.6	32.1	86.6	142.3
	TAG	19.8	38.4	83.9	142.1
	TNP	18.6	39.6	83.9	142.0
120 K (solid)	QHA	23.8	32.7	97.5	154.0
	MSD	23.2	31.8	90.8	145.7
	TAG	19.6	39.2	87.1	145.9
	TNP	18.6	39.6	83.9	145.9
180 K (solid)	QHA	24.0	33.1	97.8	154.9
	MSD	22.8	32.0	94.2	149.0
	TAG	20.2	39.8	89.0	149.0
	TNP	18.9	41.1	88.9	149.0
220 K (liquid)	MSD	22.3	31.9	95.9	150.1

Table 6.4: Calculated principal components of the hydrogen kinetic energy tensor and total kinetic energy at all temperatures studied and for all methods. The TAG and TNP approximations have been omitted from the liquid data as no plateau was ever reached.

was to investigate these kinds of phenomena, this suggests that the QHA is unsuitable for the calculation of momentum distributions in most relevant cases.

Intriguingly, there are also surprisingly large differences between the MSD and the TAG/TNP results when the individual kinetic energy eigenvalues are considered. In particular, the TAG and TNP results disagree with those of the MSD and QHA methods for the lowest hydrogen kinetic energy components. While it is not certain which of these two is correct; given that these components have the smallest ZPE, and thus the QHA should give a good description of these modes, this could be considered to be weak evidence in favour of MSD.

However, more problems with the TAG and TNP methods can be seen when looking at figures 6.10 to 6.13, which show how the results converge with respect to the window parameter used. Firstly, we see that both techniques have difficulty converging the two degenerate kinetic energy tensor eigenvalues for nitrogen. For the TAG method this is not surprising, since the same logic used in section 6.3.3 suggests that the issue illustrated in figure 6.2 will be most problematic when the distributions

Figure 6.10: Plots showing the convergence of the TAG and TNP methods with respect to window size (Δt in equation 6.3.18) at 50 K, for H and N respectively.





Figure 6.11: As figure 6.10, but at 120 K.





Figure 6.12: As figure 6.10, but at 180 K.

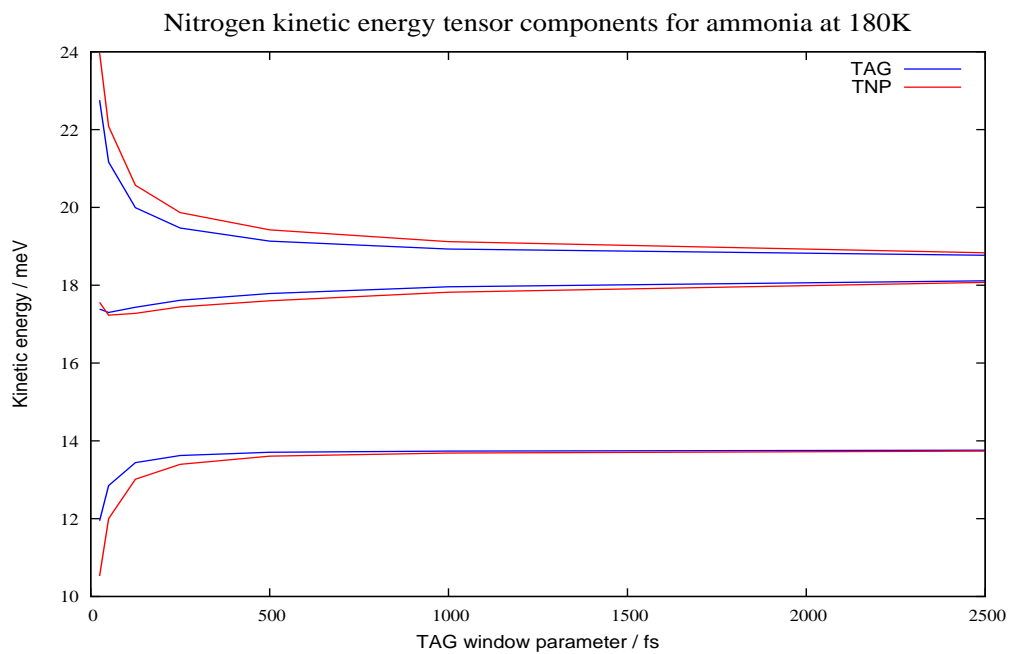
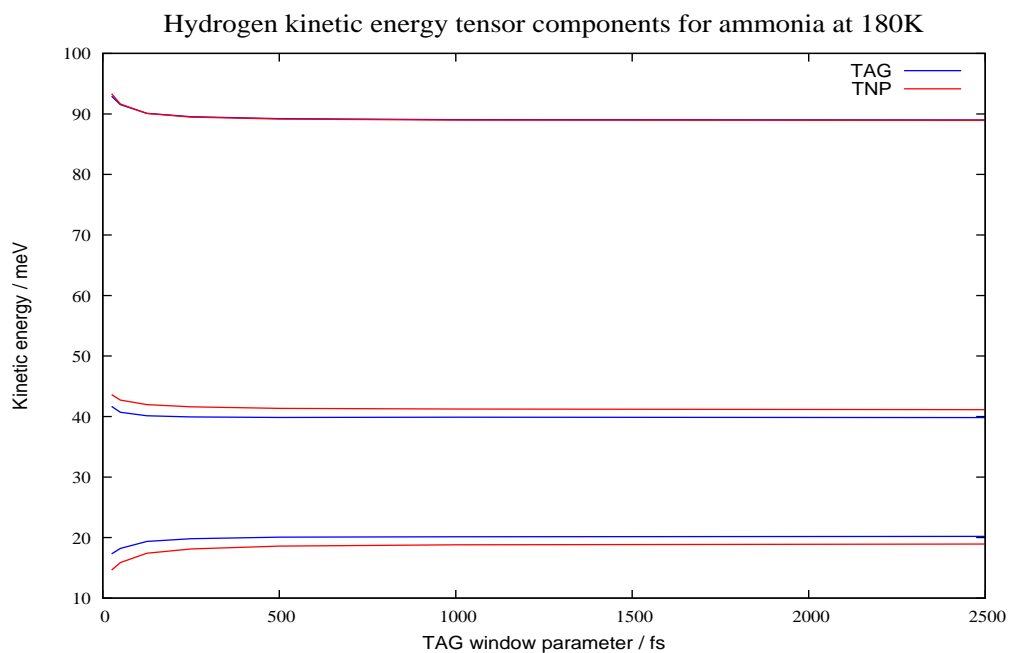
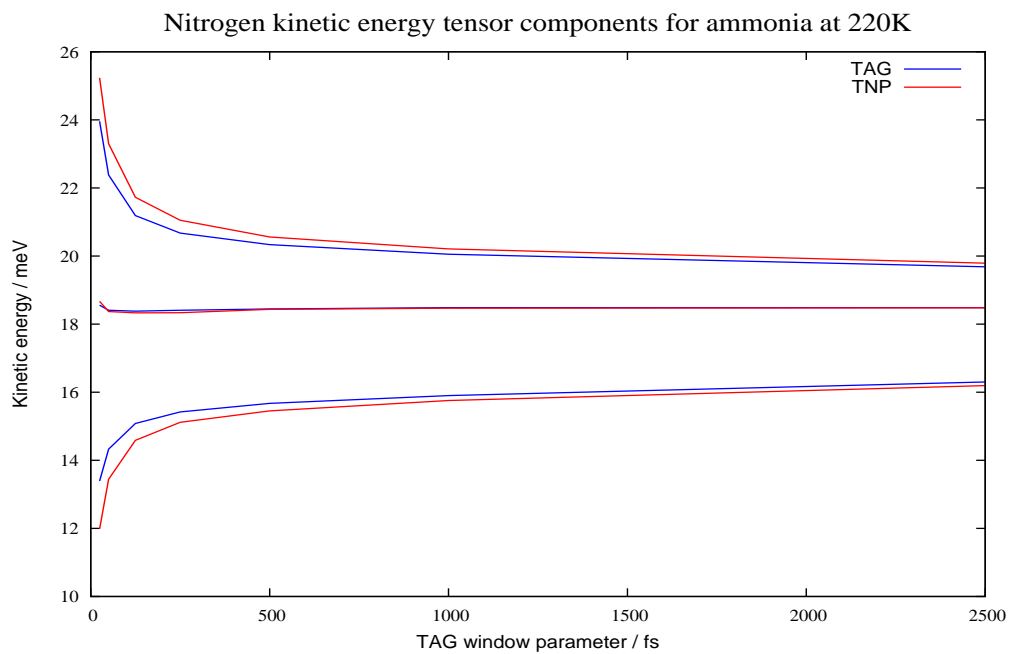



Figure 6.13: As figure 6.10, but at 220 K.



are the same. In fact, the difference between the averages of these eigenvalues will be proportional to the standard deviation of the θ_α^j distributions, and so will only vanish once $\bar{\mathbf{K}}_n^j$ has been converged.

The failure of the TNP method might not have been predicted so easily, but it appears that the problem with inverting the $n(p)$ distribution is particularly difficult in the case of two degenerate eigenvalues. This is probably because the shape of the momentum distribution changes much less when one eigenvalue is increased and another decreased if they are close together in energy, so the inversion becomes very sensitive to statistical noise in the limit as two eigenvalues become equal.

In this particular case the problem is easily solved, since the calculated $\langle \theta_\alpha^j \rangle$ values corresponding to the degenerate eigenvalues can simply be averaged to give the best guess. More rigorously, the data could instead be fit to an anisotropic Gaussian in which $\sigma_x \neq \sigma_y = \sigma_z$, including the degeneracy explicitly. However, this does suggest that, in general, if two eigenvalues are close together rather than rigorously degenerate it will be necessary to use a very large window parameter in order to reach a plateau.

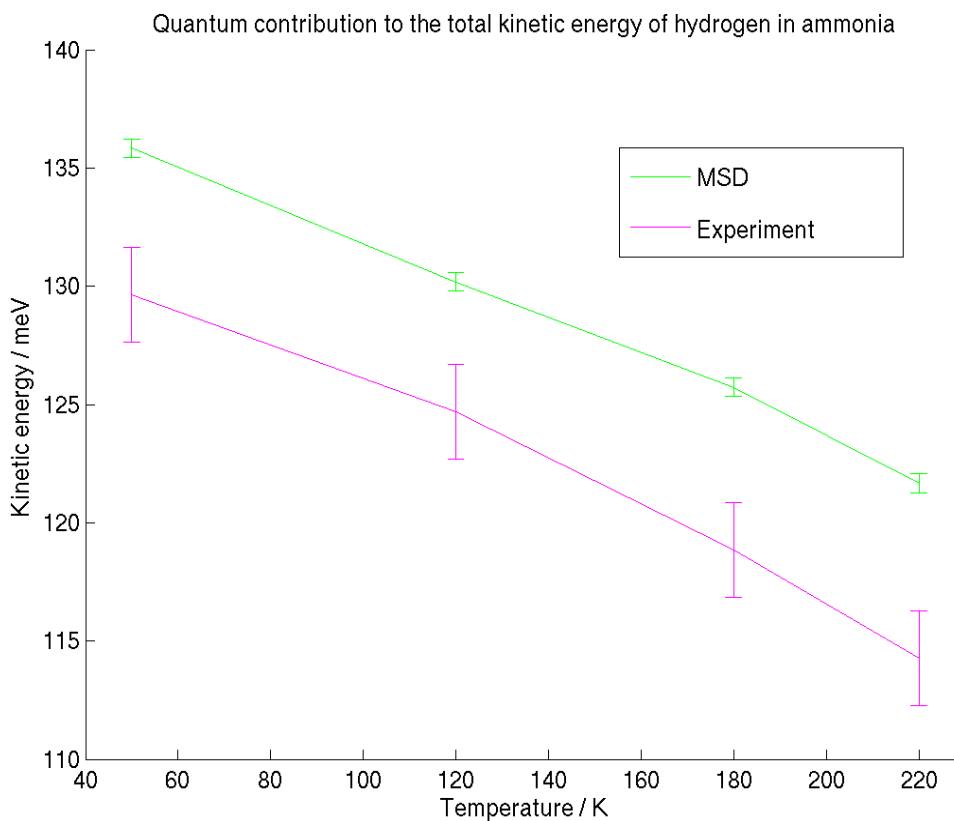
A more serious problem is that in the liquid case no plateau is reached, which calls into question whether the TAG and TNP results shown in figure 6.13 are even meaningful. It is for this reason that we have omitted the TAG and TNP data for liquid ammonia from the previous tables and graphs. This is consistent with the fact that the timescale for rotational diffusion is much shorter for ammonia than water (the rotational diffusional constant being $0.42 \text{ rad}^2\text{ps}^{-1}$ for ammonia at 220 K compared to $0.12 \text{ rad}^2\text{ps}^{-1}$ for water at 298K [235]), such that the time window needed to converge each of the kinetic energy tensor eigenvalues is either comparable to or greater than this.

Fortunately, there do not appear to be any problems with the data produced by the MSD technique. This is perhaps not surprising, since it is not affected by rota-

tional diffusion as long as the molecular geometry is retained, and does not require any approximate ordering procedures. As such, this is the method we will use to compare the results of the MD simulations to the harmonic analysis and the experimental results.

Comparison with experiment

Figure 6.14: A comparison of the total kinetic energy of hydrogen between experiment (carried out by Andrew Seel at the Rutherford Appleton Laboratory) and MD simulations. The classical contribution of $3k_B T/2$ has been subtracted from each value.



The results of the DINS experiment against the MSD data are shown in figure 6.14. From this graph it is obvious that the absolute value of the kinetic energy is not accurately calculated by the PIMD simulations. In fact, across the whole

temperature range there is a discrepancy of about 5-6 meV, or about 4-5 %. This is most likely to be due to inaccuracies in the chosen exchange-correlation functional, and is consistent with the results of a previous study of water [59], which showed similar errors.

Fortunately, an absolute measurement of the kinetic energy is not usually necessary, as only an accurate description of kinetic energy differences is necessary to understand CQE and for isotope fractionation ratios. Since the trend with respect to temperature found computationally agrees quantitatively with experiments to within statistical accuracy, this suggests that these properties will be given correctly by PIMD simulations.

On the other hand, there is very little in our data that suggests that CQE are very important in ammonia. Figure 6.9 shows that there is almost no change in the slope of the kinetic energy tensor components with respect to temperature at the melting point (which is 195.4 K for ammonia at standard pressure [236]). While there is a statistically significant reduction of the largest kinetic energy tensor component upon melting for both N and H compared to an extrapolation of the solid data and vice-versa, it is not nearly so significant as the large repartitioning of energy seen between water and ice [192].

This is again consistent with the experimental data, which shows that there does not appear to be any significant change in the vibrational spectrum of ammonia upon melting [237]. This could have been predicted by the fact that we expect the hydrogen bond network, which underlies most of the CQE in water, to be much weaker in ammonia due to the geometry of the system.

In conclusion, we have found that the MSD method is the most reliable theoretical method for the current system, as it can cope gracefully with degeneracy in the kinetic energy tensor eigenvalues and works even in the limit as the timescale of rotational

diffusion tends to zero. Comparison with a harmonic analysis has shown that the leading order anharmonic effect in ammonia is seen in the reduction of the ZPE of the highly quantized NH bond. Finally, the MSD method compared favourably with experiments, matching the trend of total kinetic energy versus temperature perfectly, but did not show any evidence in favour of large CQE in the momentum distributions of ammonia.

Chapter 7

Conclusions

In this thesis we have discussed several computational and theoretical methods used to facilitate the study of nuclear quantum effects in condensed phase hydrogen-bonded systems. In particular, we presented i-PI, a new piece of software developed to make it easier to run path integral molecular dynamics simulations with *ab initio* electronic structure calculations.

This was designed on the principle that the nuclear dynamics should be as independent from the particular choice of electronic structure code as possible, both so that i-PI is compatible with a wide range of different programs and to minimize the implementation effort required by the user. To achieve this, the i-PI dynamics routines run separately to the electronic structure program, with the only data communication between them being via internet sockets.

Many of the latest innovations in path integral molecular dynamics have been included in i-PI, from techniques to reduce the number of path integral replicas required for numerical convergence to estimators and algorithms used to efficiently calculate static and dynamical properties. The second half of this thesis discussed in detail some specific applications of these capabilities.

Firstly, we introduced a novel integrator for constant pressure dynamics in chap-

ter 5. This is a particularly important application for the combination of path integral molecular dynamics and *ab initio* electronic structure calculations, since it is not often possible to get accurate experimental density measurements of systems under extremes of pressure and temperature. It also makes direct comparison of experimental and computational results simpler, and often leads to a much more natural way of calculating properties which involve fluctuations in the volume.

To test its efficacy, this algorithm was applied to a system of water at 750 K and 10 GPa. The main result of this work was to show that, while the influence of nuclear quantum effects on global properties such as number density is generally small, this is not true of local bond fluctuations or dissociations. By using a simple geometric criterion as a measure of the frequency of transient auto-ionization events, namely changes in the number of hydrogens for which a particular oxygen atom is their nearest neighbour, it was shown that the concentration of short-lived ion pairs increased by an order of magnitude due to nuclear quantum effects. Similarly, instances of genuine bond dissociation were studied in a qualitative manner, and seen to increase by about two orders of magnitude.

This is in agreement with previous computational studies of water at room temperature and pressure which showed that the distribution of OH bond lengths has a much larger tail in quantum simulations compared to classical ones, even when many other properties agree closely with each other. The effect of pressure, which reduces O-O nearest neighbour distance and hence the size of the fluctuations required for auto-ionization, is therefore expected to be exaggerated in quantum simulations, as is observed.

On the other hand, a quantitative measure of the frequency of bond dissociation would require the use of quantum rate theory techniques such as RPMD or CMD, which were not applied in this study. Further work would therefore be required to

make any quantitative statements about proton mobility in this system.

Chapter 6 included a comparative study of several different methods used to calculate single particle momentum distributions. We are interested in this because the classical distribution is known *a priori*, so the importance of the nuclear quantum effects is easily calculated, and also because it gives a direct measure of the degree of quantum delocalization for any given particle, giving useful information about its chemical environment.

Ammonia was studied for two reasons: firstly because it was a natural extension to previous work on bulk water, and secondly since experimental measurements of this system were being done concurrently at the Rutherford Appleton Laboratory. While both the solid and liquid phases were investigated, the liquid phase is most interesting, since one of the most important aspects of the momentum distribution is its local anisotropy, which tends to get averaged out over time during simulations.

Two of the methods used to calculate momentum distributions, the transient anisotropic Gaussian method and the transient momentum distribution method, worked on the principle that, while the momentum distribution will eventually be spherically averaged due to rotational diffusion, over sufficiently small time scales the local frame of reference will stay constant. These methods then require that the time needed to converge the components of the momentum distribution is less than the rotational diffusion time scale.

Another technique commonly used for these calculations is the mean square displacement method. This relies on the existence of a consistent molecular geometry between different time steps, so that each step can be compared on a level footing. For each molecular configuration generated by the dynamics, an orthogonal rotation matrix is found which gives the least-squares fit to a reference geometry and this rotated frame used for all statistical averages.

Finally, as we are also interested in the contribution of anharmonicity, we compared the path integral results to those given by a quasi-harmonic approximation, which uses the normal modes of the zero-temperature optimized configuration to approximate the finite-temperature momentum distribution components.

Looking at the results, all methods were broadly in agreement that the hydrogen momentum distribution is much broader than that of the nitrogen nuclei, and that one component of the hydrogen momentum distribution has a much larger contribution to the kinetic energy than any of the others. This is expected to be due to the OH bond stretch, as was explicitly shown using the mean-square displacement method by taking the eigenvector corresponding to the largest component of the momentum distribution and depicting it in the reference frame.

However, while the different path integral based methods were generally in agreement, the transient anisotropic Gaussian and transient momentum distribution techniques both failed for the liquid system. This suggests that the time scale of rotational diffusion is shorter than that needed to converge the momentum components, in agreement with experimental studies of the rotational diffusion constant of ammonia which show that it is much larger than that of water. The mean square displacement method showed no such problems, and so was used as the baseline against which to compare the experimental results and the other theoretical methods.

In particular, it was found that the quasi-harmonic approximation showed a small but consistent over-estimation of all the momentum distribution components, and particularly the one corresponding to the OH bond stretch. This is to be expected, simply because the magnitude of the kinetic energy in this mode is much greater than the others, since this has the greatest zero-point energy.

Similarly, while there was a discrepancy of about 5 % between the absolute value of the total kinetic energy found by experiments and that found using the mean-

square displacement method, the trend was the same to well within the observed error bars. The difference can be easily explained by the errors in the density functional exchange-correlation functional, and is not particularly worrying since most properties of interest are only dependent on energy differences, rather than on their absolute values.

While the experimental trend was reproduced accurately by our simulations, this trend showed no evidence for the hydrogen bond network having a strong influence on the dynamics of ammonia: there was no obvious repartitioning of energy between the different momentum distribution components upon melting, nor a significant change in the trend of any component with respect to temperature. This is in agreement IR spectra, which show no major change in any the vibrational spectra of solid vs liquid ammonia, and suggests that competing quantum effects, which are very important in water, are unlikely to have a huge influence on the dynamics of ammonia.

Appendix A

Stochastic Liouvillian operators

In this appendix we will discuss how to derive Liouvillian operators for stochastic dynamical equations. We will first adapt the argument of sections 3.4 and 7.2 of ref. [100] to derive a closed-form expression for the Liouvillian operator. Then, we will show how to calculate this operator for a given set of equations of motion, in a similar way to the calculation of Kramers-Moyal expansion coefficients in ref. [89].

However, note that this is not meant to be a fully rigorous proof, since stochastic calculus has several mathematical technicalities which go far beyond the scope of this thesis. In particular, it can be shown that the time integral of the random noise gives a quantity which is non-differentiable, meaning that the Langevin equation defined in this thesis does not technically exist.

Various methods such as Itô calculus can be used to put this theory on a mathematically sound basis, but these will not be discussed here.

A.1 General formalism

To determine the Liouvillian operator for a Markovian stochastic process, the time evolution of phase space must be studied. In general, this is governed by the Chapman-

Kolmogorov equation, given by

$$\rho(x_3, t_3|x_1, t_1) = \int dx_2 \rho(x_3, t_3|x_2, t_2) \rho(x_2, t_2|x_1, t_1) \quad (\text{A.1.1})$$

where $t_1 \leq t_2 \leq t_3$ and $\rho(x_3, t_3|x_1, t_1)$ is the transition probability for going from state x_1 at time t_1 to state x_3 at time t_3 , normalized by

$$\int dx_3 \rho(x_3, t_3|x_1, t_1) = 1 \quad (\text{A.1.2})$$

In this notation, a phase space average of a general function $f(x)$ is defined as

$$\langle f(y) \rangle = \int dx f(x) \rho(x, t|y, t') \quad (\text{A.1.3})$$

To study the evolution of the phase space, we take the time derivative to give

$$\frac{\partial}{\partial t} \langle f(y) \rangle = \int dx f(x) \frac{\partial \rho(x, t|y, t')}{\partial t} \quad (\text{A.1.4})$$

$$= \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \int dx f(x) (\rho(x, t + \Delta t|y, t') - \rho(x, t|y, t')) \quad (\text{A.1.5})$$

Using equation A.1.1 on the first term and equation A.1.2 on the second gives

$$\begin{aligned} \frac{\partial}{\partial t} \langle f(y) \rangle = \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \int dx f(x) & \left(\int dz \rho(x, t + \Delta t|z, t) \rho(z, t|y, t') \right. \\ & \left. - \rho(x, t|y, t') \int dz \rho(z, t + \Delta t|x, t) \right) \quad (\text{A.1.6}) \end{aligned}$$

$$= \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \int dx \int dz (f(x) - f(z)) \rho(x, t + \Delta t|z, t) \rho(z, t|y, t') \quad (\text{A.1.7})$$

where in the last term the dummy indices x and z have been interchanged.

For applications to MD we are only interested in stochastic equations which produce continuous trajectories, i.e. those for which $|x - z| = |\Delta x| \rightarrow 0$ in the limit as

$\Delta t \rightarrow 0$. We can therefore expand the function $f(x)$ using a Taylor series to give

$$f(x) - f(z) = \Delta x \frac{\partial f(z)}{\partial z} + \frac{\Delta x^2}{2} \frac{\partial^2 f(z)}{\partial z^2} + \mathcal{O}(\Delta x^3) \quad (\text{A.1.8})$$

As we will show later, no terms of $\mathcal{O}(\Delta x^3)$ contribute to the integral, so we define

$$A(z) = \lim_{\Delta t \rightarrow 0} \int dx \frac{\Delta x}{\Delta t} \rho(x, t + \Delta t | z, t) \quad (\text{A.1.9})$$

$$B(z) = \lim_{\Delta t \rightarrow 0} \int dx \frac{\Delta x^2}{2\Delta t} \rho(x, t + \Delta t | z, t) \quad (\text{A.1.10})$$

to give

$$\frac{\partial}{\partial t} \langle f(y) \rangle = \int dz \left(A(z) \frac{\partial f(z)}{\partial z} + B(z) \frac{\partial^2 f(z)}{\partial z^2} \right) \rho(z, t | y, t') \quad (\text{A.1.11})$$

Integrating this by parts, and assuming that the surface terms go to zero, gives

$$\frac{\partial}{\partial t} \langle f(y) \rangle = \int dz f(z) \left(-\frac{\partial}{\partial z} (A(z) \rho(z, t | y, t')) + \frac{\partial^2}{\partial z^2} (B(z) \rho(z, t | y, t')) \right) \quad (\text{A.1.12})$$

Comparing equations A.1.4 and A.1.12, we see that

$$\frac{\partial \rho(z, t | y, t')}{\partial t} = \left(-\frac{\partial}{\partial z} (A(z) \rho(z, t | y, t')) + \frac{\partial^2}{\partial z^2} (B(z) \rho(z, t | y, t')) \right) \quad (\text{A.1.13})$$

Finally, the definition of the Liouvillian operator in equation 3.2.1 gives

$$\mathcal{L} = \frac{\partial}{\partial z} (A(z) \cdot) - \frac{\partial^2}{\partial z^2} (B(z) \cdot) \quad (\text{A.1.14})$$

This is equivalent to a Kramers-Moyal expansion [238, 239], truncated to second order.

A.2 Drift and diffusion constants

$A(z)$ and $B(z)$ are the drift and diffusion constants respectively, as discussed in the multi-dimensional case in equation 2.2.2. To calculate these we note that equations A.1.9 and A.1.10 can be rewritten using equation A.1.3 to give

$$A(x) = \lim_{\Delta t \rightarrow 0} \frac{\langle \Delta x(\Delta t) \rangle}{\Delta t} \quad (\text{A.2.1})$$

$$B(x) = \lim_{\Delta t \rightarrow 0} \frac{\langle \Delta x(\Delta t)^2 \rangle}{2\Delta t} \quad (\text{A.2.2})$$

where $\Delta x(\Delta t) = x(t + \Delta t) - x(t)$

Let us assume that the equation of motion for x can be written as

$$\dot{x}(t) = a(x(t), t) + b(x(t), t)\xi(t) \quad (\text{A.2.3})$$

This gives

$$\Delta x(\Delta t) = \int_t^{t+\Delta t} dt' \dot{x}(t') \quad (\text{A.2.4})$$

$$= \int_t^{t+\Delta t} dt' a(x(t'), t') + b(x(t'), t')\xi(t') \quad (\text{A.2.5})$$

$$= \int_t^{t+\Delta t} dt' a(x(t), t') + \Delta x(t' - t) \frac{\partial a(x(t), t')}{\partial x} + \\ b(x(t), t')\xi(t') + \Delta x(t' - t) \frac{\partial b(x(t), t')}{\partial x} \xi(t') + \dots \quad (\text{A.2.6})$$

Now we note three things. Firstly, $\Delta x(t' - t)$ has the same functional form as $\Delta x(\Delta t)$, and so it can also be represented using equation A.2.6. This procedure can then be repeated with the new variable $\Delta x(t'' - t)$ and so on, until only functions of $x(t)$ remain. Secondly, in the limit as $\Delta t \rightarrow 0$ only terms involving a single integral will survive. Finally, when taking the phase space average of $\Delta x(\Delta t)$ we can make

use of the properties of $\xi(t)$, i.e. $\langle \xi(t) \rangle = 0$ and $\langle \xi(t)\xi(t') \rangle = \delta(t' - t)$.

Using all these facts and equation A.2.1 gives

$$A(x) = \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \left[\int_t^{t+\Delta t} dt' \langle a(x(t), t') \rangle + \int_t^{t+\Delta t} dt' \int_t^{t'} dt'' \left\langle b(x(t), t'') \xi(t'') \frac{\partial b(x(t), t')}{\partial x} \xi(t') \right\rangle + \dots \right] \quad (\text{A.2.7})$$

$$= \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \left[\int_t^{t+\Delta t} dt' \langle a(x(t), t') \rangle + \int_t^{t+\Delta t} dt' \int_t^{t'} dt'' \left\langle b(x(t), t'') \frac{\partial b(x(t), t')}{\partial x} \right\rangle \delta(t' - t'') + \dots \right] \quad (\text{A.2.8})$$

$$= \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \int_t^{t+\Delta t} dt' \langle a(x(t), t') \rangle + \left\langle b(x(t), t) \frac{\partial b(x(t), t)}{\partial x} \right\rangle + \mathcal{O}(\Delta t) \quad (\text{A.2.9})$$

$$= \left\langle a(x(t), t) + b(x(t), t) \frac{\partial b(x(t), t)}{\partial x} \right\rangle \quad (\text{A.2.10})$$

Repeating for $B(x)$, there is only one leading term, given by

$$B(x) = \lim_{\Delta t \rightarrow \infty} \frac{1}{2\Delta t} \int_t^{t+\Delta t} dt' \int_t^{t+\Delta t} dt'' \langle \dot{x}(t') \dot{x}(t'') \rangle \quad (\text{A.2.11})$$

$$= \lim_{\Delta t \rightarrow \infty} \frac{1}{2\Delta t} \int_t^{t+\Delta t} dt' \int_t^{t+\Delta t} dt'' \langle b(x(t), t') \xi(t') b(x(t), t'') \xi(t'') \rangle + \dots \quad (\text{A.2.12})$$

$$= \lim_{\Delta t \rightarrow \infty} \frac{1}{2\Delta t} \int_t^{t+\Delta t} dt' \int_t^{t+\Delta t} dt'' \langle b(x(t), t') b(x(t), t'') \rangle \delta(t' - t'') + \dots \quad (\text{A.2.13})$$

$$= \lim_{\Delta t \rightarrow \infty} \frac{1}{2\Delta t} \int_t^{t+\Delta t} dt' \langle b(x(t), t')^2 \rangle + \mathcal{O}(\Delta t) \quad (\text{A.2.14})$$

$$= \frac{1}{2} \langle b(x(t), t)^2 \rangle \quad (\text{A.2.15})$$

Earlier, we argued that terms of $\mathcal{O}(\Delta x^3)$ in the Taylor series of equation A.1.8 can be ignored, as they do not give any contribution to the Liouvillian. The reason for this can be seen clearly in the above derivation, as the only terms in the expansion which survive in the limit as $\Delta t \rightarrow 0$ are those which involve a single integral. Since a

combination of n noise functions can only reduce the number of integrals by $n/2$ [100], all terms in functions of order $\mathcal{O}(\Delta x^3)$ and above must involve at least two integrals, and so will give zero in the limit as $\Delta t \rightarrow 0$.

Therefore the Liouvillian operator for equations of motion specified by equation A.2.3 is

$$\mathcal{L} = \frac{\partial}{\partial x} \left(\left\langle a(x(t), t) + b(x(t), t) \frac{\partial b(x(t), t)}{\partial x} \right\rangle \cdot \right) - \frac{1}{2} \frac{\partial^2}{\partial x^2} \left(\left\langle b(x(t), t)^2 \right\rangle \cdot \right) \quad (\text{A.2.16})$$

Finally, we can consider the example of a white noise Langevin equation, for which $a(x) = -\gamma x$ and $b(x) = \sqrt{2m\gamma/\beta}$. The corresponding Liouvillian operator is then given by

$$\mathcal{L} = -\gamma \frac{\partial}{\partial x} (x \cdot) - \frac{m\gamma}{\beta} \frac{\partial^2}{\partial x^2}. \quad (\text{A.2.17})$$

Appendix B

Constant pressure integration algorithm

Here we discuss how to integrate the equations of motion given in section 5.3. The equations of motion are split into four groups which can be integrated analytically, as will now be shown. In all that follows, we will use the symbols a, b, c and A to represent integration constants.

B.1 Applying the Langevin thermostat

Equations 5.3.1a and 5.3.1g describe a system of $3N + 1$ independent and equivalent stochastic differential equations, of the form

$$\dot{x}(t) = \sqrt{\frac{2m\gamma_x}{\beta_n}} \xi_x(t) - \gamma_x x(t) \quad (\text{B.1.1})$$

Multiplying both sides of this equation by $\exp(\gamma_x t)$ gives

$$\exp(\gamma_x t)\dot{x}(t) + \gamma_x \exp(\gamma_x t)x(t) = \sqrt{\frac{2m\gamma_x}{\beta_n}} \exp(\gamma_x t)\xi_x(t) \quad (\text{B.1.2})$$

$$\therefore \frac{d}{dt} (\exp(\gamma_x t)x(t)) = \sqrt{\frac{2m\gamma_x}{\beta_n}} \exp(\gamma_x t)\xi_x(t) \quad (\text{B.1.3})$$

$$\therefore \exp(\gamma_x t)x(t) = a + \int_0^t \sqrt{\frac{2m\gamma_x}{\beta_n}} \exp(\gamma_x t')\xi_x(t')dt' \quad (\text{B.1.4})$$

Taking this equation at $t = 0$ gives us $a = x(0)$, from which we have

$$x(t) = x(0)\exp(-\gamma_x t) + \int_0^t \sqrt{\frac{2m\gamma_x}{\beta_n}} \exp(\gamma_x(t' - t))\xi_x(t')dt' \quad (\text{B.1.5})$$

$$= x(0)\exp(-\gamma_x t) + I \quad (\text{B.1.6})$$

where I is a stochastic integral which we must now investigate.

By the definition of an integral, I can be replaced with an infinite sum using

$$I = \sqrt{\frac{2m\gamma_x}{\beta_n}} \exp(-\gamma_x t) \lim_{\Delta t \rightarrow 0} \left[\sum_{s=0}^{t/\Delta t} \exp(\gamma_x s\Delta t)\xi_x(s\Delta t)\Delta t \right] \quad (\text{B.1.7})$$

This then corresponds to the mean of an infinite number of independent observations of the same random variable, and so by the central limit theorem this must have the same distribution as a single Gaussian random variable [240].

Since Gaussian distributions are characterized by their variance and mean, we will now calculate these quantities for I . To do this, we first remind ourselves once again that the Gaussian random noise $\xi_x(t)$ has a zero mean, unit variance, and is uncorrelated in time, such that $\langle \xi_x(t) \rangle = 0$ and $\langle \xi_x(t)\xi_x(t') \rangle = \delta(t - t')$.

The mean is then found using

$$\langle I \rangle = \int_0^t \sqrt{\frac{2m\gamma_x}{\beta_n}} \exp(\gamma_x(t' - t)) \langle \xi_x(t') \rangle dt' = 0 \quad (\text{B.1.8})$$

Similarly, the variance is given by

$$\langle I^2 \rangle - \langle I \rangle^2 = \langle I^2 \rangle \tag{B.1.9}$$

$$= \int_0^t \int_0^t \frac{2m\gamma_x}{\beta_n} \exp(\gamma_x(t' - t)) \exp(\gamma_x(t'' - t)) \langle \xi_x(t') \xi_x(t'') \rangle dt' dt'' \tag{B.1.10}$$

$$= \int_0^t \int_0^t \frac{2m\gamma_x}{\beta_n} \exp(\gamma_x(t' - t)) \exp(\gamma_x(t'' - t)) \delta(t' - t'') dt' dt'' \tag{B.1.11}$$

$$= \int_0^t \frac{2m\gamma_x}{\beta_n} \exp(2\gamma_x(t' - t)) dt' \tag{B.1.12}$$

$$= \frac{m}{\beta_n} [\exp(2\gamma_x(t' - t))]_0^t \tag{B.1.13}$$

$$= \frac{m}{\beta_n} (1 - \exp(-2\gamma_x t)) \tag{B.1.14}$$

We can therefore avoid the evaluation of the integral I , and instead draw a single value from a Gaussian random number with zero mean and a variance of $m(1 - \exp(-2\gamma_x t)) / \beta_n$, which then gives

$$x(t) = x(0)\exp(-\gamma_x t) + \sqrt{\frac{m}{\beta_n} (1 - \exp(-2\gamma_x t))} \xi_x(0) \tag{B.1.15}$$

B.2 Propagating the forces

Consider the equations of motion 5.3.1b and 5.3.1h, reproduced below

$$\dot{\mathbf{p}}^{(k)}(t) = \tilde{\mathbf{f}}_j^{(k)} \tag{B.2.1}$$

$$\dot{\eta}(t) = 3 \left[nV(P_{\text{int}}(t) - P_{\text{ext}}) + \frac{1}{\beta_n} \right] \tag{B.2.2}$$

For clarity, let us rewrite this last equation as

$$\dot{\eta}(t) = A + \sum_{j=1}^N \frac{|\tilde{\mathbf{p}}_j^{(0)}(t)|^2}{m_j} \tag{B.2.3}$$

where

$$A = \frac{3}{\beta_n} - 3nVP_{\text{ext}} - 3V \sum_{i=0}^n \frac{dU(\mathbf{r}_1^{(i)}, \dots, \mathbf{r}_N^{(i)})}{dV} + \sum_{i=0}^n \sum_{j=1}^N (\mathbf{r}_j^{(i)} - \mathbf{r}_j^c) \cdot \frac{\partial U(\mathbf{r}_1^{(i)}, \dots, \mathbf{r}_N^{(i)})}{\partial \mathbf{r}_i^{(j)}} \quad (\text{B.2.4})$$

To solve this, first we integrate the nuclear momentum equation of motion, which gives

$$\tilde{\mathbf{p}}^{(k)}(t) = \tilde{\mathbf{f}}_j^{(k)} t + a \quad (\text{B.2.5})$$

$$\therefore \tilde{\mathbf{p}}^{(k)}(0) = a \quad (\text{B.2.6})$$

$$\therefore \tilde{\mathbf{p}}^{(k)}(t) = \tilde{\mathbf{f}}_j^{(k)} t + \tilde{\mathbf{p}}^{(k)}(0) \quad (\text{B.2.7})$$

Then, taking the time derivative of the cell momentum equation of motion, we see that

$$\ddot{\eta}(t) = 2 \sum_{j=1}^N \frac{\tilde{\mathbf{p}}_j^{(0)}(t) \cdot \dot{\tilde{\mathbf{p}}}_j^{(0)}(t)}{m_j} \quad (\text{B.2.8})$$

$$= 2 \sum_{j=1}^N \frac{\tilde{\mathbf{p}}_j^{(0)}(t) \cdot \tilde{\mathbf{f}}_j^{(0)}}{m_j} \quad (\text{B.2.9})$$

$$= 2 \sum_{j=1}^N \frac{\tilde{\mathbf{p}}_j^{(0)}(0) \cdot \tilde{\mathbf{f}}_j^{(0)}}{m_j} + 2t \sum_{j=1}^N \frac{|\tilde{\mathbf{f}}_j^{(0)}|^2}{m_j} \quad (\text{B.2.10})$$

$$\therefore \eta(t) = b + ct + t^2 \sum_{j=1}^N \frac{\tilde{\mathbf{p}}_j^{(0)}(0) \cdot \tilde{\mathbf{f}}_j^{(0)}}{m_j} + \frac{t^3}{3} \sum_{j=1}^N \frac{|\tilde{\mathbf{f}}_j^{(0)}|^2}{m_j} \quad (\text{B.2.11})$$

Finally, the constants of integration can be found using the boundary conditions

$\eta(0) = b$ and $\dot{\eta}(0) = c$, to give

$$\eta(t) = \eta(0) + t \left(A + \sum_{j=1}^N \frac{|\tilde{\mathbf{p}}_j^{(0)}(0)|^2}{m_j} \right) + t^2 \sum_{j=1}^N \frac{\tilde{\mathbf{p}}_j^{(0)}(0) \cdot \tilde{\mathbf{f}}_j^{(0)}}{m_j} + \frac{t^3}{3} \sum_{j=1}^N \frac{|\tilde{\mathbf{f}}_j^{(0)}|^2}{m_j} \quad (\text{B.2.12})$$

$$\begin{aligned} &= \eta(0) + 3t \left[nV(P_{\text{int}}(0) - P_{\text{ext}}) + \frac{1}{\beta_n} \right] + \\ & t^2 \sum_{j=1}^N \frac{\tilde{\mathbf{p}}_j^{(0)}(0) \cdot \tilde{\mathbf{f}}_j^{(0)}}{m_j} + \frac{t^3}{3} \sum_{j=1}^N \frac{|\tilde{\mathbf{f}}_j^{(0)}|^2}{m_j} \end{aligned} \quad (\text{B.2.13})$$

B.3 Propagating the centroid motion

For the step applying the effect of the barostat, we use equations 5.3.1c, 5.3.1e, 5.3.1f and the centroid mode of 5.3.1d, reproduced below

$$\dot{\tilde{\mathbf{p}}}_j^{(0)}(t) = -\frac{\tilde{\mathbf{p}}_j^{(0)}(t)\eta}{\mu} \quad (\text{B.3.1})$$

$$\dot{\tilde{\mathbf{r}}}_j^{(0)}(t) = \frac{\tilde{\mathbf{p}}_j^{(0)}(t)}{m_j} + \frac{\tilde{\mathbf{r}}_j^{(0)}(t)\eta}{\mu} \quad (\text{B.3.2})$$

$$\dot{V}(t) = \frac{3V(t)\eta}{\mu} \quad (\text{B.3.3})$$

Since equations B.3.1 and B.3.3 involve first-order differential equations of one variable only they can be trivially solved, giving

$$\tilde{\mathbf{p}}_j^{(0)}(t) = \tilde{\mathbf{p}}_j^{(0)}(0) \exp\left(-\frac{\eta t}{\mu}\right) \quad (\text{B.3.4})$$

$$V(t) = V(0) \exp\left(\frac{3\eta t}{\mu}\right) \quad (\text{B.3.5})$$

Multiplying both sides of equation B.3.2 by $\exp(-\eta t/\mu)$ gives

$$\exp\left(-\frac{\eta t}{\mu}\right) \dot{\tilde{\mathbf{r}}}_j^{(0)}(t) - \frac{\eta}{\mu} \exp\left(-\frac{\eta t}{\mu}\right) \tilde{\mathbf{r}}_j^{(0)}(t) = \frac{1}{m_j} \exp\left(-\frac{\eta t}{\mu}\right) \tilde{\mathbf{p}}_j^{(0)}(t) \quad (\text{B.3.6})$$

$$\therefore \frac{d}{dt} \left(\exp\left(-\frac{\eta t}{\mu}\right) \tilde{\mathbf{r}}_j^{(0)}(t) \right) = \frac{1}{m_j} \exp\left(-\frac{2\eta t}{\mu}\right) \tilde{\mathbf{p}}_j^{(0)}(0) \quad (\text{B.3.7})$$

$$\therefore \exp\left(-\frac{\eta t}{\mu}\right) \tilde{\mathbf{r}}_j^{(0)}(t) = -\frac{\mu}{2\eta m_j} \exp\left(-\frac{2\eta t}{\mu}\right) \tilde{\mathbf{p}}_j^{(0)}(0) + b \quad (\text{B.3.8})$$

$$\begin{aligned} \therefore \tilde{\mathbf{r}}_j^{(0)}(t) &= -\frac{\mu}{2\eta m_j} \exp\left(-\frac{\eta t}{\mu}\right) \tilde{\mathbf{p}}_j^{(0)}(0) \\ &\quad + b \exp\left(\frac{\eta t}{\mu}\right) \end{aligned} \quad (\text{B.3.9})$$

Next, we use the boundary condition to give

$$\tilde{\mathbf{r}}_j^{(0)}(0) = -\frac{\mu}{2\eta m_j} \tilde{\mathbf{p}}_j^{(0)}(0) + b \quad (\text{B.3.10})$$

$$\therefore b = \tilde{\mathbf{r}}_j^{(0)}(0) + \frac{\mu}{2\eta m_j} \tilde{\mathbf{p}}_j^{(0)}(0) \quad (\text{B.3.11})$$

Finally, substituting this into the original equation gives

$$\tilde{\mathbf{r}}_j^{(0)}(t) = \frac{\mu}{2\eta m_j} \left(\exp\left(\frac{\eta t}{\mu}\right) - \exp\left(-\frac{\eta t}{\mu}\right) \right) \tilde{\mathbf{p}}_j^{(0)}(0) + \tilde{\mathbf{r}}_j^{(0)}(0) \exp\left(\frac{\eta t}{\mu}\right) \quad (\text{B.3.12})$$

$$= \frac{\mu}{\eta m_j} \sinh\left(\frac{\eta t}{\mu}\right) \tilde{\mathbf{p}}_j^{(0)}(0) + \tilde{\mathbf{r}}_j^{(0)}(0) \exp\left(\frac{\eta t}{\mu}\right) \quad (\text{B.3.13})$$

B.4 Propagating the free ring polymer motion

The free ring polymer motion is described by equation 5.3.1d, as reproduced below

$$\dot{\tilde{\mathbf{r}}}_j^{(k)}(t) = \frac{\tilde{\mathbf{p}}_j^{(k)}(t)}{m_j} \quad (\text{B.4.1})$$

$$\dot{\tilde{\mathbf{p}}}_j^{(k)}(t) = -m_j \omega_k^2 \tilde{\mathbf{r}}_j^{(k)}(t) \quad (\text{B.4.2})$$

Taking the time derivative of equation B.4.2 gives

$$\ddot{\tilde{\mathbf{p}}}_j^{(k)}(t) = -m_j \omega_k^2 \tilde{\mathbf{r}}_j^{(k)}(t) = -\omega_k^2 \tilde{\mathbf{p}}_j^{(k)}(t) \quad (\text{B.4.3})$$

The solution to this simple second-order ordinary differential equation is well known, and is given by

$$\therefore \tilde{\mathbf{p}}_j^{(k)}(t) = a \cos(\omega_k t) + b \sin(\omega_k t) \quad (\text{B.4.4})$$

$$\therefore \tilde{\mathbf{p}}_j^{(k)}(0) = a \quad (\text{B.4.5})$$

$$\therefore \tilde{\mathbf{p}}_j^{(k)}(t) = \tilde{\mathbf{p}}_j^{(k)}(0) \cos(\omega_k t) + b \sin(\omega_k t) \quad (\text{B.4.6})$$

$$\therefore \dot{\tilde{\mathbf{p}}}_j^{(k)}(t) = -\omega_k \tilde{\mathbf{p}}_j^{(k)}(0) \sin(\omega_k t) + b \omega_k \cos(\omega_k t) \quad (\text{B.4.7})$$

Substituting this into equation B.4.2 gives

$$-m_j \omega_k^2 \tilde{\mathbf{r}}_j^{(k)}(t) = -\omega_k \tilde{\mathbf{p}}_j^{(k)}(0) \sin(\omega_k t) + b \omega_k \cos(\omega_k t) \quad (\text{B.4.8})$$

$$\therefore \tilde{\mathbf{r}}_j^{(k)}(t) = \frac{1}{m_j \omega_k} \left(\tilde{\mathbf{p}}_j^{(k)}(0) \sin(\omega_k t) - b \cos(\omega_k t) \right) \quad (\text{B.4.9})$$

$$\therefore \tilde{\mathbf{r}}_j^{(k)}(0) = -\frac{b}{m_j \omega_k} \quad (\text{B.4.10})$$

$$\therefore b = -m_j \omega_k \tilde{\mathbf{r}}_j^{(k)}(0) \quad (\text{B.4.11})$$

The solutions are then given by

$$\therefore \tilde{\mathbf{r}}_j^{(k)}(t) = \frac{1}{m_j \omega_k} \tilde{\mathbf{p}}_j^{(k)}(0) \sin(\omega_k t) + \tilde{\mathbf{r}}_j^{(k)}(0) \cos(\omega_k t) \quad (\text{B.4.12})$$

$$\therefore \tilde{\mathbf{p}}_j^{(k)}(t) = \tilde{\mathbf{p}}_j^{(k)}(0) \cos(\omega_k t) - m_j \omega_k \tilde{\mathbf{r}}_j^{(k)}(0) \sin(\omega_k t) \quad (\text{B.4.13})$$

Bibliography

- [1] M. P. Allen and D. J. Tildesley. *Computer Simulation of Liquids*. Oxford University Press, New York, 1987.
- [2] D. Frenkel and B. Smit. *Understanding Molecular Simulation From Algorithms to Applications*. Academic Press, San Diego, 2001.
- [3] J. Chocholoušová and M. Feig. Implicit solvent simulations of DNA and DNA-protein complexes: Agreement with explicit solvent vs experiment. *J. Phys. Chem. B*, 110:17240, 2006.
- [4] T. W. Allen, T. Baštuž, S. Kuyucak, and S.-H. Chung. Gramicidin A channel as a test ground for molecular dynamics force fields. *Biophys. J.*, 84:2159, 2003.
- [5] R. B. Best, N.-V. Buchete, and H. Gerhard. Are current molecular dynamics force fields too helical? *Biophys. J.*, 95:L7, 2008.
- [6] S. Habershon, T. E. Markland, and D. E. Manolopoulos. Competing quantum effects in the dynamics of a flexible water model. *J. Chem. Phys.*, 131:024501, 2009.
- [7] Z. Cao and J. Wang. A comparative study of two different force fields on structural and thermodynamics character of H1 peptide via molecular dynamics simulations. *J. Biomol. Struct. Dyn.*, 27:651, 2010.

- [8] M. Boero, M. Parrinello, and K. Terakura. First principles molecular dynamics study of Ziegler-Natta heterogeneous catalysis. *J. Am. Chem. Soc.*, 120:2746, 1998.
- [9] L. Sun, K. Song, and W. L. Hase. A S_N2 reaction that avoids its deep potential energy minimum. *Science*, 296:875, 2002.
- [10] M. Ø. Jensen, D. W. Borhani, K. Lindorff-Larsen, P. Maragakis, V. Jogini, M. P. Eastwood, R. O. Dror, and D. E. Shaw. Principles of conduction and hydrophobic gating in K^+ channels. *Proc. Natl. Acad. Sci. U.S.A.*, 107:5833, 2010.
- [11] C. Cavazzoni. *Large Scale First-Principles Simulations of Water and Ammonia at High Pressure and Temperature*. PhD thesis, Scuola Internazionale Superiore di Studi Avanzati, Trieste, Italy, 1998.
- [12] C. Cavazzoni, G. L. Chiarotti, S. Scandolo, E. Tosatti, M. Bernasconi, and M. Parinello. Superionic and metallic states of water and ammonia at giant planet conditions. *Science*, 283:44, 1999.
- [13] M. R. Manaa, L. E. Fried, C. F. Melius, M. Elstner, and T. Frauenheim. Decomposition of HMX at extreme conditions: A molecular dynamics simulation. *J. Phys. Chem. A*, 106:9024, 2002.
- [14] L. Koči, A. B. Belonoshko, and R. Ahuja. Molecular dynamics study of liquid iron under high pressure and high temperature. *Phys. Rev. B*, 73:1098, 2006.
- [15] O. Isayev, L. Gorb, M. Qasim, and J. Leszczynski. Ab initio molecular dynamics study on the initial chemical events in nitramines: Thermal decompositions of CL-20. *J. Phys. Chem. B*, 112:11005, 2008.

- [16] J. M. McMahon, M. A. Morales, C. Pierleoni, and D. M. Ceperley. The properties of hydrogen and helium under extreme conditions. *Rev. Mod. Phys.*, 84:1607, 2012.
- [17] M. Parrinello and A. Rahman. Study of an F center in molten KCl. *J. Chem. Phys.*, 80:860, 1984.
- [18] B. de Raedt, M. Sprik, and M. L. Klein. Computer simulation of muonium in water. *J. Chem. Phys.*, 80:5719, 1984.
- [19] S. A. Chin. Symplectic integrators from composite operator factorizations. *Phys. Lett. A*, 226:344, 1997.
- [20] T. E. Markland and D. E. Manolopoulos. An efficient ring polymer contraction scheme for imaginary time path integral simulations. *J. Chem. Phys.*, 129:024105, 2008.
- [21] M. Ceriotti, D. E. Manolopoulos, and M. Parrinello. Accelerating the convergence of path integral dynamics with a generalized Langevin equation. *J. Chem. Phys.*, 134:084104, 2011.
- [22] J. Cao and G. A. Voth. A new perspective on quantum time correlation functions. *J. Chem. Phys.*, 99:10070, 1993.
- [23] J. Cao and G. 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, 1994.
- [24] 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.

- [25] S. Habershon, D. E. Manolopoulos, T. E. Markland, and T. 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, 2013.
- [26] M. Rossi, M. Ceriotti, and D. E. Manolopoulos. How to remove the spurious resonances from ring polymer molecular dynamics. *J. Chem. Phys.*, 140:234116, 2014.
- [27] D. Marx and M. Parrinello. Ab initio path-integral molecular dynamics. *Z. Phys. B*, 95:143, 1994.
- [28] R. N. Barnett, H.-P. Cheng, H. Hakkinen, and U. Landman. Studies of excess electrons in sodium chloride clusters and of excess protons in water clusters. *J. Phys. Chem.*, 99:7731, 1995.
- [29] D. Marx and M. Parrinello. Structural quantum effects and three-centre two-electron bonding in CH_5^+ . *Nature*, 375:216, 1995.
- [30] D. Marx and M. Parrinello. Ab initio path integral molecular dynamics: Basic ideas. *J. Chem. Phys.*, 104:4077, 1996.
- [31] S. Biermann, D. Hohl, and D. Marx. Proton quantum effects in high pressure hydrogen. *J. Low Temp. Phys.*, 110:97, 1998.
- [32] D. Marx, M. E. Tuckerman, and M. Parrinello. Solvated excess protons in water: quantum effects on the hydration structure. *J. Phys.: Condens. Matter*, 12:A153, 2000.
- [33] M. E. Tuckerman, D. Marx, M. L. Klein, and M. Parrinello. Efficient and general algorithms for path integral Car-Parrinello molecular dynamics. *J. Chem. Phys.*, 104:5579, 1996.

- [34] T. von Rosenvinge, M. E. Tuckerman, and M. L. Klein. Ab initio molecular dynamics study of crystal hydrates of HCl including path integral results. *Faraday Discuss.*, 106:273, 1997.
- [35] H. S. Mei, M. E. Tuckerman, D. E. Sagnella, and M. L. Klein. Quantum nuclear ab initio molecular dynamics study of water wires. *J. Phys. Chem. B*, 102:10446, 1998.
- [36] S. Miura, M. E. Tuckerman, and M. L. Klein. An ab initio path integral molecular dynamics study of double proton transfer in the formic acid dimer. *J. Chem. Phys.*, 109:5290, 1998.
- [37] B. Walker and A. Michaelides. Direct assessment of quantum nuclear effects on hydrogen bond strength by constrained-centroid ab initio path integral molecular dynamics. *J. Chem. Phys.*, 133:174306, 2010.
- [38] X.-Z. Li, B. Walker, and A. Michaelides. Quantum nature of the hydrogen bond. *Proc. Natl. Acad. Sci. U.S.A.*, 108:6369, 2011.
- [39] J. A. Morrone and R. Car. Nuclear quantum effects in water. *Phys. Rev. Lett.*, 101:017801, 2008.
- [40] J. A. Morrone, L. Lin, and R. Car. Tunneling and delocalization effects in hydrogen bonded systems: A study in position and momentum space. *J. Chem. Phys.*, 130:204511, 2009.
- [41] C. Zhang and A. Michaelides. Quantum nuclear effects on the location of hydrogen above and below the palladium (100) surface. *Surf. Sci.*, 605:689, 2011.

- [42] X.-Z. Li, B. Walker, M. I. J. Probert, C. J. Pickard, R. J. Needs, and A. Michaelides. Classical and quantum ordering of protons in cold solid hydrogen under megabar pressures. *J. Phys.: Condens. Matter*, 25, 2013.
- [43] J. Chen, X.-Z. Li, Q. Zhang, M. I. J. Probert, C. J. Pickard, R. J. Needs, A. Michaelides, and E. Wang. Quantum simulation of low-temperature metallic liquid hydrogen. *Nat. Commun.*, 4:2064, 2013.
- [44] M. Ceriotti, J. N. More, and D. E. Manolopoulos. i-PI: a Python interface for ab initio path integral molecular dynamics simulations. *Comput. Phys. Commun.*, 185:1019, 2014.
- [45] M. Ceriotti, M. Parinello, T. E. Markland, and D. E. Manolopoulos. Efficient stochastic thermostating of path integral molecular dynamics. *J. Chem. Phys.*, 133:124104, 2010.
- [46] H. C. Andersen. Molecular dynamics simulations at constant pressure and/or temperature. *J. Chem. Phys.*, 72:2384, 1980.
- [47] M. Parrinello. Polymorphic transitions in single crystals: A new molecular dynamics method. *J. Appl. Phys.*, 52:7182, 1981.
- [48] G. J. Martyna, A. Hughes, and M. E. Tuckerman. Molecular dynamics algorithms for path integrals at constant pressure. *J. Chem. Phys.*, 110:3275, 1999.
- [49] G. Bussi, T. Zykova-Timan, and M. Parinello. Isothermal-isobaric molecular dynamics using stochastic velocity rescaling. *J. Chem. Phys.*, 130:074101, 2009.
- [50] M. Born and J. R. Oppenheimer. Zur quantentheorie der molekeln. *Ann. Phys.*, 389:457, 1927.

- [51] P. A. M. Dirac. *The Principles of Quantum Mechanics*. The International Series of Monographs on Physics 27. Oxford University Press, New York, 1958.
- [52] D. Chandler. *Introduction to Modern Statistical Mechanics*. Oxford University Press, New York, 1987.
- [53] R P Feynman and A R Hibbs. *Quantum Mechanics and Path Integrals*. McGraw-Hill, New York, 1964.
- [54] H. F. Trotter. On the product of semi-groups of operators. *Proc. Am. Math. Soc.*, 10:545, 1959.
- [55] M. Tuckerman, B. J. Berne, and G. J. Martyna. Reversible multiple time scale molecular dynamics. *J. Chem. Phys.*, 97:1990, 1992.
- [56] R. W. Hall and B. J. Berne. Nonergodicity in path integral molecular dynamics. *J. Chem. Phys.*, 81:3641, 1984.
- [57] M. F. Herman, E. J. Bruskin, and B. J. Berne. On path integral Monte Carlo simulations. *J. Chem. Phys.*, 76:5150, 1982.
- [58] J. Cao and B. J. Berne. On energy estimators in path integral Monte Carlo simulations: Dependence of accuracy on algorithm. *J. Chem. Phys.*, 91:6359, 1989.
- [59] M. Ceriotti and D. E. Manolopoulos. Efficient first principles calculation of the quantum kinetic energy and momentum distribution of nuclei. *Phys. Rev. Lett.*, 109:100604, 2012.
- [60] M. Takahashi and M. Imada. Monte Carlo calculation of quantum systems. II. higher order correction. *J. Phys. Soc. Japan*, 53:3765, 1984.

- [61] X. Li and J. Q. Broughton. High order correction to the Trotter expansion for use in computer simulation. *J. Chem. Phys.*, 86:5094, 1987.
- [62] M. Suzuki. Hybrid exponential product formulas for unbounded operators with possible applications to Monte Carlo simulations. *Phys. Lett. A*, 201:425, 1995.
- [63] S. Jang, S. Jang, and G. A. Voth. Applications of higher order composite factorization schemes in imaginary time path integral simulations. *J. Chem. Phys.*, 115:7832, 2001.
- [64] T. 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.
- [65] A. Pérez and M. 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.
- [66] G. S. Fanourgakis, G. K. Schenter, and S. S. Xantheas. A quantitative account of quantum effects in liquid water. *J. Chem. Phys.*, 125:141102, 2006.
- [67] M. Ceriotti, G. a. R. Brain, O. Riordan, and D. E. Manolopoulos. The inefficiency of re-weighted sampling and the curse of system size in high-order path integration. *Proc. R. Soc. A*, 468(2137):2, 2011.
- [68] T. E. Markland and D. E. Manolopoulos. A refined ring polymer contraction scheme for systems with electrostatic interactions. *Chem. Phys. Lett.*, 464:256, 2008.
- [69] G. S. Fanourgakis, T. E. Markland, and D. E. Manolopoulos. A fast path integral method for polarizable force fields. *J. Chem. Phys.*, 131:094102, 2009.

- [70] P. Ewald. Die berechnung optischer und elektrostatischer gitterpotentiale. *Ann. Phys.*, 64:253, 1921.
- [71] J. W. Eastwood, R. W. Hockney, and D. Lawrence. P3M3DP - three dimensional periodic particle-particle/particle-mesh program. *Comput. Phys. Commun.*, 19:215, 1980.
- [72] T. E. Markland and B. J. Berne. Unraveling quantum mechanical effects in water using isotopic fractionation. *Proc. Natl. Acad. Sci. U.S.A.*, 109:7988, 2012.
- [73] J. Liu, R. S. Andino, C. M. Miller, X. Chen, D. M. Wilkins, M. Ceriotti, and D. E. Manolopoulos. A surface-specific isotope effect in mixtures of light and heavy water. *J. Phys. Chem. C*, 117:2944, 2013.
- [74] E. Anglada, J. Junquera, and J. M. Soler. Efficient mixed-force first-principles molecular dynamics. *Phys. Rev. E*, 68:055701, 2003.
- [75] M. Guidon, F. Schiffmann, J. Hutter, and J. VandeVondele. Ab initio molecular dynamics using hybrid density functionals. *J. Chem. Phys.*, 128:214104, 2008.
- [76] R. P. Steele. Communication: Multiple-timestep ab initio molecular dynamics with electron correlation. *J. Chem. Phys.*, 139:011102, 2013.
- [77] N. Luehr, T. E. Markland, and T. J. Martínez. Multiple time step integrators in ab initio molecular dynamics. *J. Chem. Phys.*, 140:084116, 2014.
- [78] M. Ceriotti, G. Bussi, and M. Parinello. Nuclear quantum effects in solids using a colored-noise thermostat. *Phys. Rev. Lett.*, 103:030603, 2009.
- [79] M. Ceriotti, G. Bussi, and M. Parinello. Colored-noise thermostats à la carte. *J. Chem. Theory Comput.*, 6:1170, 2010.

- [80] P. Langevin. Sur la théorie du mouvement brownien. *C. R. Acad. Sci.*, 146:530, 1908.
- [81] D. L. Ermak and Y. Yeh. Equilibrium electrostatic effects on the behavior of polyions in solution: polyion-mobile ion interaction. *Chem. Phys. Lett.*, 24:243, 1974.
- [82] T. Schneider and E. Stoll. Molecular-dynamics study of a three-dimensional one-component model for distortive phase transitions. *Phys. Rev. B.*, 17:1302, 1978.
- [83] S. Nosé. A molecular dynamics method for simulations in the canonical ensemble. *Mol. Phys.*, 52:255, 1984.
- [84] W. G. Hoover. Canonical dynamics: Equilibrium phase-space distributions. *Phys Rev. A*, 31:1695, 1985.
- [85] G. J. Martyna, M. L. Klein, and M. Tuckerman. Nosé-Hoover chains: The canonical ensemble via continuous dynamics. *J. Chem. Phys.*, 97:2635, 1992.
- [86] G. Bussi, D. Donadio, and M. Parrinello. Canonical sampling through velocity rescaling. *J. Chem. Phys.*, 126:014101, 2007.
- [87] G. Bussi and M. Parrinello. Stochastic thermostats: comparison of local and global schemes. *Comput. Phys. Commun.*, 179:26, 2008.
- [88] G. Bussi and M. Parrinello. Accurate sampling using Langevin dynamics. *Phys. Rev. E*, 75:056707, 2007.
- [89] H. Risken. *The Fokker-Planck Equation*. Springer, New York, 1989.

- [90] M. Ceriotti, G. Bussi, and M. Parrinello. Langevin equation with colored noise for constant-temperature molecular dynamics simulations. *Phys. Rev. Lett.*, 102:020601, 2009.
- [91] M. Ceriotti and M. Parrinello. The delta-thermostat: selective normal-modes excitation by colored-noise Langevin dynamics. *Procedia Comput. Sci.*, 1:1607, 2010.
- [92] R. Zwanzig. *Nonequilibrium Statistical Mechanics*. Springer, Berlin, 2001.
- [93] G. E. Uhlenbeck and L. S. Ornstein. On the theory of Brownian motion. *Phys. Rev.*, 36:823, 1930.
- [94] S. Habershon and D. E. Manolopoulos. Zero point energy leakage in condensed phase dynamics: An assessment of quantum simulations for liquid water. *J. Chem. Phys.*, 131:244518, 2009.
- [95] B. N. Finkelstein. Über den virialsatz in der wellenmechanik. *Z. Phys. A*, 50:293, 1928.
- [96] E. Hückel. Quantentheoretische beiträge zum benzolproblem. *Z. Phys.*, 70:204, 1931.
- [97] J. W. Cooley and J. W. Tukey. An algorithm for the machine calculation of complex Fourier series. *Math. Comput.*, 19:297, 1965.
- [98] M. E. Tuckerman, Y. Liu, G. Ciccotti, and G. J. Martyna. Non-Hamiltonian molecular dynamics: Generalizing Hamiltonian phase space principles to non-Hamiltonian systems. *J. Chem. Phys.*, 115:1678, 2001.
- [99] V. E. Tarasov. Stationary solutions of Liouville equations for non-Hamiltonian systems. *Ann. Phys.*, 316:393, 2005.

- [100] C. W. Gardiner. *Handbook of Stochastic Mechanics*. Springer, Berlin, 2003.
- [101] J. C. Sexton and D. H. Weingarten. Hamiltonian evolution for the hybrid Monte Carlo algorithm. *Nucl. Phys. B*, 380:665, 1992.
- [102] H. J. C. Berendsen, J. R. Grigera, and T. P. Straatsma. The missing term in effective pair potentials. *J. Phys. Chem.*, 91:6269, 1987.
- [103] J. E. Lennard-Jones. Cohesion. *Proc. Phys. Soc.*, 43:461, 1931.
- [104] E. Sanz, C. Vega, J. L. F. Abascal, and L. G. MacDowell. Phase diagram of water from computer simulation. *Phys. Rev. Lett.*, 92:255701, 2004.
- [105] J. L. F. Abascal and C. Vega. A general purpose model for the condensed phases of water: TIP4P/2005. *J. Chem. Phys.*, 123:234505, 2005.
- [106] Y. Wu, H. L. Tepper, and G. A. Voth. Flexible simple point-charge water model with improved liquid-state properties. *J. Chem. Phys.*, 124:024503, 2006.
- [107] C. Vega and J. L. F. Abascal. Relation between the melting temperature and the temperature of maximum density for the most common models of water. *J. Chem. Phys.*, 123:144504, 2005.
- [108] G. S. Fanourgakis and S. S. Xantheas. The flexible, polarizable Thole-type interaction potential for water (TTM2-F) revisited. *J. Phys. Chem. A*, 110:4100, 2006.
- [109] F. Paesani, W. Zhang, D. A. Case, T. E. Cheatham III, and G. A. Voth. An accurate and simple quantum model for liquid water. *J. Chem. Phys.*, 125:184507, 2006.
- [110] P. Hohenberg and W. Kohn. Inhomogeneous electron gas. *Phys. Rev.*, 136:B864, 1964.

- [111] M. Levy. Universal variational functionals of electron densities, first-order density matrices, and natural spin-orbitals and solution of the v -representability problem. *Proc. Natl. Acad. Sci. U.S.A.*, 76:6062, 1979.
- [112] W. Kohn. Density functional and density matrix method scaling linearly with the number of atoms. *Phys. Rev. Lett.*, 76:3168, 1996.
- [113] S. Goedecker. Linear scaling electronic structure methods. *Rev. Mod. Phys.*, 71:1085, 1999.
- [114] T. Spura, H. Elgabarty, and T. D. Kühne. “On the fly” coupled cluster path-integral molecular dynamics: impact of nuclear quantum effects on the protonated water dimer. *Phys. Chem. Chem. Phys.*, 17:14355, 2015.
- [115] R. G. Parr and W. Yang. Density-functional theory of the electronic structure of molecules. *Annu. Rev. Phys. Chem.*, 46:701, 1995.
- [116] W. Kohn. Density functional theory for systems of very many atoms. *Int. J. Quantum Chem.*, 56:229, 1995.
- [117] W. Kohn and L. J. Sham. Self-consistent equations including exchange and correlation effects. *Phys. Rev.*, 140:A1133, 1965.
- [118] J. P. Perdew and Y. Wang. Accurate and simple analytic representation of the electron-gas correlation energy. *Phys. Rev. B*, 45:13244, 1992.
- [119] J. P. Perdew, K. Burke, and M. Ernzerhof. Generalized gradient approximation made simple. *Phys. Rev. Lett.*, 77:3865, 1996.
- [120] J. P. Perdew, K. Burke, and M. Ernzerhof. Erratum: Generalized gradient approximation made simple. *Phys. Rev. Lett.*, 78:1396, 1997.

- [121] A. D. Becke. Density-functional exchange-energy approximation with correct asymptotic behavior. *Phys. Rev. A*, 38:3098, 1988.
- [122] C. Lee, W. Yang, and R. G. Parr. Development of the Colle-Salvetti correlation-energy formula into a functional of the electron density. *Phys. Rev. B*, 37:785, 1988.
- [123] Y. Zhao and D. G. Truhlar. The M06 suite of density functionals for main group thermochemistry, thermochemical kinetics, noncovalent interactions, excited states, and transition elements: two new functionals and systematic testing of four M06-class functionals and 12 other functionals. *Theor. Chem. Account.*, 120:215, 2008.
- [124] B. Santra, A. Michaelides, and M. Scheffler. On the accuracy of density-functional theory exchange-correlation functionals for H bonds in small water clusters: Benchmarks approaching the complete basis set limit. *J. Chem. Phys.*, 127:184104, 2007.
- [125] R. G. Parr. Density functional theory. *Annu. Rev. Phys. Chem.*, 34:631, 1983.
- [126] J. S. Tse. Ab initio molecular dynamics with density functional theory. *Annu. Rev. Phys. Chem.*, 53:249, 2002.
- [127] J. Klimeš and A. Michaelides. Perspective: Advances and challenges in treating van der Waals dispersion forces in density functional theory. *J. Chem. Phys.*, 137:120901, 2012.
- [128] K. Burke. Perspective on density functional theory. *J. Chem. Phys.*, 136:150901, 2012.
- [129] A. D. Becke. Perspective: Fifty years of density-functional theory in chemical physics. *J. Chem. Phys.*, 140:18A301, 2014.

- [130] R. P. Feynman. Forces in molecules. *Phys. Rev.*, 56:340, 1939.
- [131] P. Pulay. Ab initio calculation of force constants and equilibrium geometries in polyatomic molecules. *Mol. Phys.*, 17:197, 1969.
- [132] Z. Lin. Pulay forces in density functional theory for periodic and molecular systems. *Phys. Lett. A*, 299:413, 2002.
- [133] P. Giannozzi, S. Baroni, N. Bonini, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, G. L. Chiarotti, M. Cococcioni, I. Dabo, A. Dal Corso, S. de Gironcoli, S. Fabris, G. Fratesi, R. Gebauer, U. Gerstmann, C. Gougoussis, A. Kokalj, M. Lazzeri, L. Martin-Samos, N. Marzari, F. Mauri, R. Mazzarello, S. Paolini, A. Pasquarello, L. Paulatto, C. Sbraccia, S. Scandolo, G. Sclauszero, A. P. Seitsonen, A. Smogunov, P. Umari, and R. M. Wentzcovitch. QUANTUM ESPRESSO: a modular and open-source software project for quantum simulations of materials. *J. Phys. Condens. Matt.*, 21:395502, 2009.
- [134] D. Alfè. Ab initio molecular dynamics, a simple algorithm for charge extrapolation. *Comput. Phys. Commun.*, 118:31, 1999.
- [135] V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter, and M. Scheffler. Ab initio molecular simulations with numeric atom-centered orbitals. *Comput. Phys. Commun.*, 180:2175, 2009.
- [136] CP2K. Copyright IBM Corp 1990-2006, Copyright MPI für Festkörperforschung Stuttgart 1997-2001.
- [137] CP2K. <http://cp2k.berlios.de>.
- [138] S. Plimpton. Fast parallel algorithms for short-range molecular dynamics. *J. Comput. Phys.*, 117:1, 1995.

- [139] GLE4MD. <http://epfl-cosmo.github.io/gle4md/>.
- [140] M. Ceriotti and T. E. Markland. Efficient methods and practical guidelines for simulating isotope effects. *J. Chem. Phys.*, 138:014112, 2013.
- [141] L. Lin, J. A. Morrone, R. Car, and M. Parrinello. Displaced path integral formulation for the momentum distribution of quantum particles. *Phys. Rev. Lett.*, 105:110602, 2010.
- [142] T. D. Hone, P. J. Rossky, and G. A. Voth. A comparative study of imaginary time path integral based methods for quantum dynamics. *J. Chem. Phys.*, 124:154103, 2006.
- [143] S. Habershon, G. S. Fanourgakis, and D. E. Manolopoulos. Comparison of path integral molecular dynamics methods for the infrared absorption spectrum of liquid water. *J. Chem. Phys.*, 129:074501, 2008.
- [144] I. R. Craig and D. E. Manolopoulos. Chemical reaction rates from ring polymer molecular dynamics. *J. Chem. Phys.*, 122:084106, 2005.
- [145] I. R. Craig and D. E. Manolopoulos. A refined ring polymer molecular dynamics theory of chemical reaction rates. *J. Chem. Phys.*, 123:034102, 2005.
- [146] J. O. Richardson and S. C. Althorpe. Ring-polymer molecular dynamics rate-theory in the deep-tunneling regime: Connection with semiclassical instanton theory. *J. Chem. Phys.*, 131:214106, 2009.
- [147] T. J. H. Hele and S. 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.

- [148] S. C. Althorpe and T. J. H. Hele. Derivation of a true ($t \rightarrow 0_+$) quantum transition-state theory. II. Recovery of the exact quantum rate in the absence of recrossing. *J. Chem. Phys.*, 139:084115, 2013.
- [149] T. J. H. Hele and S. C. Althorpe. On the uniqueness of $t \rightarrow 0_+$ quantum transition-state theory. *J. Chem. Phys.*, 139:084116, 2013.
- [150] B. J. Braams and D. E. Manolopoulos. On the short-time limit of ring polymer molecular dynamics. *J. Chem. Phys.*, 125(12):124105, 2006.
- [151] I. R. Craig and D. E. Manolopoulos. Inelastic neutron scattering from liquid para-hydrogen by ring polymer molecular dynamics. *Chem. Phys.*, 322:236, 2006.
- [152] S. Jang and G. 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.
- [153] A. Witt, S. D. Ivanov, M. Shiga, H. Forbert, and D. Marx. On the applicability of centroid and ring polymer path integral molecular dynamics for vibrational spectroscopy. *J. Chem. Phys.*, 130:194510, 2009.
- [154] S. D. Ivanov, A. Witt, M. Shiga, and D. Marx. Communications: On artificial frequency shifts in infrared spectra obtained from centroid molecular dynamics: Quantum liquid water. *J. Chem. Phys.*, 132:031101, 2010.
- [155] J. C. Owicki and H. A. Scheraga. Monte Carlo calculations in the isothermal-isobaric ensemble. 2. Dilute aqueous solution of methane. *J. Am. Chem. Soc.*, 99:7413, 1977.
- [156] F. D. Murnaghan. The compressibility of media under extreme pressures. *Proc. Natl. Acad. Sci.*, 30:244, 1944.

- [157] A. D. Fortes, J. P. Brodholt, I. G. Wood, and L. Vočadlo. Hydrogen bonding in solid ammonia from ab initio calculations. *J. Chem. Phys.*, 118:5987, 2003.
- [158] D. P. Cruikshank. *Neptune and Triton*. University of Arizona Press, Tuscon, AZ, 1995.
- [159] A. C. Mitchell and W. J. Nellis. Equation of state and electrical conductivity of water and ammonia shocked to the 100 GPa (1 Mbar) pressure range. *J. Chem. Phys.*, 76:6273, 1982.
- [160] W. J. M. Rankine. On the thermodynamic theory of waves of finite longitudinal disturbances. *Philos. T. Roy. Soc.*, 160:277, 1870.
- [161] H. R. Wendt and F. F. Abraham. Empirical criterion for the glass transition region based on Monte Carlo simulations. *Phys. Rev. Lett.*, 41:1244, 1978.
- [162] M. Sprik, R. W. Impey, and M. L. Klein. Second-order elastic constants for the Lennard-Jones solid. *Phys. Rev. B*, 29:4368, 1984.
- [163] P. Attard. On the density of volume states in the isobaric ensemble. *J. Chem. Phys.*, 103:9884, 1995.
- [164] I. F. Silvera and V. V. Goldman. The isotropic intermolecular potential for H₂ and D₂ in the solid and gas phases. *J. Chem. Phys.*, 69:4209, 1978.
- [165] E. Schwegler, G. Galli, F. Gygi, and R. Q. Hood. Dissociation of water under pressure. *Phys. Rev. Lett.*, 87:265501, 2001.
- [166] S. Goedecker, M. Teter, and J. Hutter. Separable dual-space Gaussian pseudopotentials. *Phys. Rev. B*, 54:1703, 1996.

- [167] S. Grimme, J. Antony, S. Ehrlich, and H. Krieg. A consistent and accurate ab initio parametrization of density functional dispersion correction (DFT-D) for the 94 elements H-Pu. *J. Chem. Phys.*, 132:154104, 2010.
- [168] J. Schmidt, J. VandeVondele, I.-F. W. Kuo, D. Sebastiani, J. I. Siepmann, J. Hutter, and C. J. Mundy. Isobaric-isothermal molecular dynamics simulations utilizing density functional theory: An assessment of the structure and density of water at near-ambient conditions. *J. Phys. Chem. B*, 113:11959, 2009.
- [169] I.-C. Lin, A. P. Seitsonen, I. Tavernelli, and U. Rothlisberger. Structure and dynamics of liquid water from ab initio molecular dynamics - comparison of BLYP, PBE, and revPBE density functionals with and without van der Waals corrections. *J. Chem. Theory Comput.*, 8:3902, 2012.
- [170] A. Goncharov, N. Goldman, L. Fried, J. Crowhurst, I.-F. Kuo, C. Mundy, and J. Zaug. Dynamic ionization of water under extreme conditions. *Phys. Rev. Lett.*, 94:125508, 2005.
- [171] M. Ceriotti, J. Cuny, M. Parrinello, and D. E. Manolopoulos. Nuclear quantum effects and hydrogen bond fluctuations in water. *Proc. Natl. Acad. Sci. U.S.A.*, 110:15591, 2013.
- [172] S. D. Hamann. *Electrolyte solutions at high pressure*. Springer, 1974.
- [173] D. Homouz, G. Reiter, J. Eckert, J. Mayers, and R. Blinc. Measurement of the 3D Born-Oppenheimer potential of a proton in a hydrogen-bonded system via deep inelastic neutron scattering: The superprotonic conductor $\text{Rb}_3\text{H}(\text{SO}_4)_2$. *Phys. Rev. Lett.*, 98:115502, 2007.

- [174] C. A. Chatzidimitriou-Dreismann, M. Krzystyniak, T. Abdul-Redah, and M. Lerch. Neutron Compton scattering from the super proton conductor $\text{H}_3\text{OSbTeO}_6$ and polyethylene: a comparison of proton momentum distributions and reduced cross-sections. *J. Phys.: Condens. Matter*, 19:436231, 2007.
- [175] D. Laage and J. T. Hynes. A molecular jump mechanism of water reorientation. *Science*, 311:832, 2006.
- [176] D. Laage and J. T. Hynes. On the molecular mechanism of water reorientation. *J. Phys. Chem. B*, 112:14230, 2008.
- [177] D. Laage, G. Stirnemann, F. Sterpone, R. Rey, and J. T. Hynes. Reorientation and allied dynamics in water and aqueous solutions. *Annu. Rev. Phys. Chem.*, 62:395, 2011.
- [178] G. N. Lewis and R. T. Macdonald. Some properties of pure $\text{H}^2\text{H}^2\text{O}$. *J. Am. Chem. Soc.*, 55:3057, 1933.
- [179] D. N. Glew and N. S. Rath. Aqueous nonelectrolyte solutions. Part XIV. D-ice and D-hydrate freezing points of deuterium oxide - ethylene oxide solutions and the formula of congruent ethylene oxide D-hydrate. *Can. J. Chem.*, 74:1830, 1996.
- [180] G. S. Bordonskii and S. D. Krylov. Dielectric properties of D_2O ice in the microwave range. *Russ. J. Phys. Chem. A*, 75:834, 2001.
- [181] W. S. Price, H. Ide, and Y. Arata. Self-diffusion of supercooled water to 238 K using PGSE NMR diffusion measurements. *J. Phys. Chem. A*, 103:448, 1999.
- [182] W. S. Price, H. Ide, Y. Arata, and O. Söderman. Temperature dependence of the self-diffusion of supercooled heavy water to 244 K. *J. Phys. Chem. B*, 104:5874, 2000.

- [183] J. Vaníček and W. H. Miller. Efficient estimators for quantum instanton evaluation of the kinetic isotope effects: Application to the intramolecular hydrogen transfer in pentadiene. *J. Chem. Phys.*, 127:114309, 2007.
- [184] G. L. Richmond. Molecular bonding and interactions at aqueous surfaces as probed by vibrational sum frequency spectroscopy. *Chem. Rev.*, 102:2693, 2002.
- [185] Y. R. Shen and V. Ostroverkhov. Sum-frequency vibrational spectroscopy on water interfaces: Polar orientation of water molecules at interfaces. *Chem. Rev.*, 106:1140, 2006.
- [186] J. McGuire and Y. R. Shen. Ultrafast vibrational dynamics at water interfaces. *Science*, 313:1945, 2006.
- [187] S. Gopalakrishnan, D. Liu, H. C. Allen, M. Kuo, and M. J. Shultz. Vibrational spectroscopic studies of aqueous interfaces: Salts, acids, bases and nanodrops. *Chem. Rev.*, 106:1155, 2006.
- [188] H. Wolff and E. Wolff. Über den dampfdruck-isotopie-effekt von wasser und eis. *Ber. Bunsen-Ges. Phys. Chem.*, 73:393, 1969.
- [189] W. M. Jones. Vapor pressures of tritium oxide and deuterium oxide. interpretation of the isotope effects. *J. Chem. Phys.*, 48:207, 1967.
- [190] W. A. Van Hook. Vapor pressure isotope effect in aqueous systems. III. the vapor pressure of HOD (-60 to 200 C). *J. Phys. Chem.*, 76:3040, 1972.
- [191] J. Pupezin, G. Jakli, G. Jancso, and W. A. Van Hook. The vapor pressure isotope effect in aqueous systems. I. H₂O-D₂O (-64 C to 100 C) and H₂¹⁶O-H₂¹⁸O (-17 C to 16 C); ice and liquid. II. alkali metal chloride solution in H₂O and D₂O (-5 C to 100 C). *J. Phys. Chem.*, 76:743, 1972.

- [192] G. Romanelli, M. Ceriotti, D. E. Manolopoulos, C. Pantalei, R. Senesi, and C. Andreani. Direct measurement of competing quantum effects on the kinetic energy of heavy water upon melting. *J. Phys. Chem. Lett.*, 4:3251, 2013.
- [193] P. C. Hohenberg and P. M. Platzman. High-energy neutron scattering from liquid He⁴. *Phys. Rev.*, 152:198, 1966.
- [194] 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, 2005.
- [195] J. Mayers and G. Reiter. The VESUVIO, electron volt neutron spectrometer. *Meas. Sci. Technol.*, 23:045902, 2012.
- [196] M. Cohen and R. P. Feynman. Theory of inelastic scattering of cold neutrons from liquid helium. *Phys. Rev.*, 107:13, 1957.
- [197] G. I. Watson. Neutron Compton scattering. *J. Phys.: Condens. Matter*, 8:5955, 1996.
- [198] J. Radon. On the determination of functions from their integral values along certain manifolds. *IEEE Trans. Med. Imag.*, 5:170, 1986.
- [199] C. Andreani, G. Baciocco, R. S. Holt, and J. Mayers. Resolution in deep inelastic neutron scattering using pulsed neutron sources. *Nucl. Instrum. Meth. A*, 276:297, 1989.
- [200] V. F. Sears. Scaling and final-state interactions in deep-inelastic neutron scattering. *Phys. Rev. B*, 30:44, 1984.
- [201] G. F. Reiter, J. Mayers, and J. Noreland. Momentum-distribution spectroscopy using deep inelastic neutron scattering. *Phys. Rev. B*, 65:104305, 2002.

- [202] G. Reiter and R. Silver. Measurement of interionic potentials in solids using deep-inelastic neutron scattering. *Phys. Rev. Lett.*, 54:1047, 1985.
- [203] L. Lin, J. A. Morrone, R. Car, and M. Parrinello. Momentum distribution, vibrational dynamics, and the potential of mean force in ice. *Phys. Rev. B*, 83:220302, 2011.
- [204] M. R. Hestenes and E. Stiefel. Methods of conjugate gradients for solving linear systems. *J. Res. Nat. Bur. Stand.*, 49:409, 1952.
- [205] C. G. Broyden. The convergence of a class of double-rank minimization algorithms 1. general considerations. *J. Appl. Math.*, 6:76, 1970.
- [206] R. Fletcher. A new approach to variable metric algorithms. *Comput. J.*, 13:317, 1970.
- [207] D. Goldfarb. A family of variable-metric methods derived by variational means. *Math. Comput.*, 24:23, 1970.
- [208] D. F. Shanno. Conditioning of quasi-Newton methods for function minimization. *Math. Comput.*, 24:647, 1970.
- [209] D. F. Shanno and P. C. Kettler. Optimal conditioning of quasi-Newton methods. *Math. Comput.*, 24:657, 1970.
- [210] M. Ceriotti, G. Miceli, A. Pietropaolo, D. Colognesi, A. Nale, M. Catti, M. Bernasconi, and M. Parrinello. Nuclear quantum effects in ab initio dynamics: Theory and experiments for lithium imide. *Phys. Rev. B*, 82:174306, 2010.

- [211] J. A. Morrone, V. Srinivasan, D. Sebastiani, and R. Car. Proton momentum distribution in water: an open path integral molecular dynamics study. *J. Chem. Phys.*, 126:234504, 2007.
- [212] W. Kabsch. A solution for the best rotation to relate two sets of vectors. *Acta Cryst.*, A32:922, 1976.
- [213] W. Kabsch. A discussion of the solution for the best rotation to relate two sets of vectors. *Acta Cryst.*, A34:827, 1978.
- [214] G. Golub and W. Kahan. Calculating the singular values and pseudo-inverse of a matrix. *J. SIAM Numer. Anal. Ser. B*, 2:205, 1965.
- [215] K. Momma and F. Izumi. VESTA 3 for three-dimensional visualization of crystal, volumetric and morphology data. *J. Appl. Crystallogr.*, 44:1272, 2011.
- [216] A. J. Birch. 117. reduction by dissolving metals. Part I. *J. Chem. Soc.*, page 430, 1944.
- [217] A. J. Birch. 212. reduction by dissolving metals. Part II. *J. Chem. Soc.*, page 809, 1945.
- [218] A. J. Birch. 119. reduction by dissolving metals. Part III. *J. Chem. Soc.*, page 593, 1946.
- [219] A. J. Birch. 25. reduction by dissolving metals. Part IV. *J. Chem. Soc.*, page 102, 1947.
- [220] A. J. Birch. 327. reduction by dissolving metals. Part V. *J. Chem. Soc.*, page 1642, 1947.
- [221] A. J. Birch and S. M. Mukherji. 532. reduction by dissolving metals. Part VI. *J. Chem. Soc.*, page 2531, 1949.

- [222] P. P. Edwards. From solvated electrons to metal anions: Electronic structure and dynamics. *J. Sol. Chem.*, 14:187, 1985.
- [223] J. Olovsson and D. H. Templeton. X-ray study of solid ammonia. *Acta Crys.*, 12:832, 1959.
- [224] J. W. Reed and P. M. Harris. Neutron diffraction study of solid deuterioammonia. *J. Chem. Phys.*, 35:1730, 1961.
- [225] A. W. Hewat and C. Riekel. The crystal structure of deuterioammonia between 2 and 180 K by neutron powder profile refinement. *Acta Crys. A*, 35:569, 1979.
- [226] F. Leclercq, P. Damay, and M. Foukani. Structure of powder deuterioammonia between 2 and 180 K revisited: A refinement of the neutron diffraction pattern taking into account molecular reorientations: analysis of the diffuse intensity. *J. Chem. Phys.*, 102:4400, 1995.
- [227] J. P. Perdew, J. A. Chevary, S. H. Vosko, K. A. Jackson, M. R. Pederson, D. J. Singh, and C. Fiolhais. Atoms, molecules, solids, and surfaces: Applications of the generalized gradient approximation for exchange and correlation. *Phys. Rev. B*, 46:4978, 1993.
- [228] D. Vanderbilt. Soft self-consistent pseudopotentials in a generalized eigenvalue formalism. *Phys. Rev. B*, 41:7892, 1990.
- [229] H. J. Monkhorst and J. D. Pack. Special points for Brillouin-zone integrations. *Phys. Rev. B*, 13:5188, 1976.
- [230] V. G. Manzhelii and A. M. Tolkachev. Density of ammonia and methane in the solid state. *Fiz. Tverd. Tela S. Peterburg*, 5:3413, 1963.

- [231] A. M. Tolkachev and V. G. Manzhelii. Density of solidificated gases. *Fiz. Tverd. Tela S. Peterburg*, 7:2125, 1965.
- [232] L. Haar and J. S. Gallagher. Thermodynamic properties of ammonia. *J. Chem. Phys. Ref. Data*, 7:635, 1978.
- [233] I. B. Sladkov and I. Y. Kleoshkina. Prediction of density of molecular inorganic liquids. *Zh. Prikl. Khim.*, 64:1424, 1991.
- [234] G. Liessmann, W. Schmidt, and S. Reiffarth. *Recommended Thermophysical Data*. Data compilation of the Saechsische Olefinwerke Boehlen Germany, 1995.
- [235] E. H. Hardy, A. Zygar, M. D. Zeidler, M. Holz, and F. D. Sacher. Isotope effect on the translational and rotational motion in liquid water and ammonia. *J. Chem. Phys.*, 114:3174, 2001.
- [236] W. M. Haynes, editor. *CRC Handbook of Chemistry and Physics*. Taylor and Francis, Boca Roton, 92 edition, 2011.
- [237] A. Bromberg, S. Kimel, and A. Ron. Infrared spectrum of liquid and crystalline ammonia, 1977.
- [238] H. A. Kramers. Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica A*, 7:284, 1940.
- [239] J. E. Moyal. Stochastic processes and statistical physics. *J. R. Stat. Soc. London, Ser. B*, 11:150, 1949.
- [240] W. Feller. *An introduction to probability theory and its applications*. Wiley, New York, 1974.