Imperial College London
Department of Physics

Quantum Monte Carlo Simulations of Warm Dense Matter

Fionn Dara Malone

April 2017

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Physics of Imperial College London
and the Diploma of Imperial College London

Abstract

Recent experimental progress in laser technology has led to renewed interest in warm dense matter. Found in the interiors of gas giants and in inertial confinement fusion experiments, warm dense matter is relevant to problems of fundamental and technological importance but is a challenge to create experimentally and describe theoretically. Modern electronic structure theory, in the form of density functional theory coupled with molecular dynamics, in principle offers a route to describing realistic warm dense matter. However, until quite recently, no accurate exchange correlation free energy functionals existed and the accuracy of existing fits was largely unknown. In this thesis we extend the independent, systematically exact, density matrix quantum Monte Carlo method, to address these issues. Focussing on the warm dense uniform electron gas, we first outline how sampling issues present in the original formulation can be overcome and how numerical basis set corrections can significantly reduce the computational burden at high electronic temperatures. We next introduce a systematic approximation allowing larger system sizes to be tackled. In the process we resolve a controversy present between two competing path integral Monte Carlo methods, whose results for the exchange correlation energy of the uniform electron gas differ substantially in the warm dense regime. Finally, we develop a general procedure for deriving analytic finite size corrections in the warm dense regime, thus removing the final barrier to reaching the thermodynamic limit.

Copyright Declaration

Declaration of Originality

Except where otherwise stated, the work contained in this thesis is my own. This thesis has not been submitted in part or in whole for any other award at this or any other university. Parts of this work appear in the following publications:

- T. Dornheim, S. Groth, **F.D. Malone**, T. Schoof, T. Sjostrom, W.M.C. Foulkes, M. Bonitz, Phys. Plasm. 24, 056303 (2016)
- T. Dornheim, S. Groth, T. Sjostrom, **F.D. Malone**, W.M.C. Foulkes, M. Bonitz, Phys. Rev. Lett. 117, 156403 (2016)
- **F.D. Malone**, N.S. Blunt, E.W. Brown, D.K.K. Lee, J.S. Spencer, W.M.C. Foulkes, J.J. Shepherd, Phys. Rev. Lett. 117, 115701 (2016)
- **F.D. Malone**, N.S. Blunt, J.J. Shepherd, D.K.K. Lee, J.S. Spencer, W.M.C. Foulkes, J. Chem. Phys. 143, 044116 (2015)

The quantum Monte Carlo code developed during the course of this thesis is freely available as part of the HANDE QMC package (`http://www.hande.org.uk`), details of which are also published in:

- J.S. Spencer, N.S. Blunt, W.A. Vigor, **F.D. Malone**, W.M.C. Foulkes, J.J. Shepherd, A.J.W. Thom - J. Open Res. Soc. 3, 9 (2015)

Additional code for evaluating various analytic properties of the uniform electron gas are available at: `https://github.com/fdmalone/uegpy`.

Contents

List of Tables

List of Figures

Acknowledgements

First and foremost I want to thank my supervisors Matthew Foulkes and Derek Lee who have been a constant source of ideas and motivation for the past three and a half years. I owe a particular debt of gratitude to Matthew who carefully proofread this thesis and taught me many lessons in grammar in the process – any further errors are my own. I think most people realise how much they know about so little by the end of their PhD; it was really inspirational to have worked with people who know so much about so much. I also acknowledge support from Imperial College and the EPSRC for a President's PhD scholarship.

Two Jameses next require a special mention. James Spencer was like a third supervisor during my PhD and not only taught me how to code properly, but showed me a better and more open way to do science. James Shepherd was also a driving force behind this thesis and originally suggested looking into the uniform electron gas. His guidance, procurement of computer time, and numerous skype calls (and emails), played a large part in the eventual success of this thesis.

I have also been fortunate to have had the opportunity to work with some brilliant outside collaborators including Nick Blunt, Ethan Brown, Tim Schoof, Tobias Dornheim and Simon Groth. It is reassuring to find so many people interested in finding the right numbers and in collaborating rather than competing. I also wish to thank Prof. Dr. Michael Bonitz in Kiel for organising a successful QMC workshop which lead to this closer collaboration. A special thanks also goes out to the HANDE development team for numerous code reviews.

London can be an overwhelming place at the best of times and I feel it is only the variety of great people here which make putting up with it worthwhile. I particularly want to thank Domhnall & Emma, Matthias, Adam (and the rest of the CMTH group), Mongey, Ambra, Josep y Josep for making my time here so enjoyable.

To my wonderful wife Venina I express my eternal gratitude, firstly, for putting up with me and keeping me sane over the last two years; and secondly, for giving up the sun of Madrid for the rain of London. You make it all worthwhile.

Last but not least, I want to thank my parents, for teaching me how to learn.

Come on girl,
Let's sneak out of this party,
It's getting boring,
There's more to life than this.
– Björk

Introduction

> *We can begin by looking at the fundamental paradox of the many-body problem; namely that people who do not know how to solve the three-body problem are trying to solve the n-body problem.*
>
> – Harry J. Lipkin [1]

Many-body problems represent some of the most difficult and technologically relevant problems which need to be solved. In such problems we hope to understand and predict the properties of a system based solely off of a knowledge of its constituent components and how they interact with one another. A prime example is the problem of solving the many-electron Schrödinger equation for a collection of atoms and electrons, which governs the properties of most chemical, biological and physical systems on a microscopic level [2]. Unsurprisingly, like most many-body problems, the equations are generally impossible to solve analytically and we need to resort to approximate computer simulations to make any progress. It is the goal of *ab-initio* electronic structure methods to solve these equations.

Extensive algorithmic development, coupled with the rapid growth in computational power which occurred in the latter half of the twentieth century, now allows us to theoretically predict the properties of simple materials on a personal computer relatively accurately. However, the true power of electronic structure theory is in its potential ability to design new materials which do not yet exist and reach regimes where experiment is either currently impossible or extremely difficult. The latter of these points is perhaps seen best in the application of electronic structure methods to the problem of materials under extreme conditions, such as those found in the interiors of planets or in astrophysical contexts. Here, fundamental geophysical questions still exist regarding, e.g., the origin of the Earth's magnetic dynamo [3], while in planetary physics, reliable equations of state are required in order to constrain the ever increasing amount of exoplanet data [4]. Since we cannot journey to the centre of the Earth and our nearest neighbouring planets are many years away, reliable theoretical predictions are therefore much sought after.

One of the most challenging extreme settings which has emerged of late is that of warm dense matter (WDM) [5]. WDM is defined as matter existing at densities ranging from a tenth to a thousand times that of normal solids and at temperatues of $10^4 - 10^6$ K. WDM is thought to be found in nature in the interiors of gas giants such as Jupiter [6], and can be created in laboratory settings in inertial confinement fusion experiments [7] and in

laser irradiated solids [8]. Apart from its relevance to fundamental [9] and technological problems [10], WDM also represent a significant challenge for theory. In these regimes, matter is somewhere in between a degenerate, strongly coupled plasma, and a normal solid and exhibits characteristics of both. As such, non-perturbative theoretical methods capable of capturing strong coupling, quantum mechanical, many-body, thermal and material effects are required to describe phases ranging from a crystalline solid to a partially or fully ionised plasma.

Density functional theory (DFT) is the weapon of choice for ground state electronic structure theory and has emerged as a useful tool, coupled with molecular dynamics, in the description of WDM. At its heart lies an approximate form for the exchange correlation free energy functional the quality of which is generally unknown and varies from system to system. As DFT becomes more and more widely used in this field it is important to have alternative methods which can benchmark the approximate DFT results.

In this thesis we will outline how one such alternative method, called density Matrix Quantum Monte Carlo (DMQMC), can be adapted to simulate WDM. Taking up the challenge from the epigraph, we start from the 4-body problem before attempting to solve the N-body one.

Chapter 1 sets out the basics of statistical mechanics of many-electron systems. We then outline how conventional electronic structure methods need to be adapted to include the effects of temperature, highlighting some of the limitations and complications which arise in the process.

In Chapter 2 we introduce various quantum Monte Carlo methods which offer a stochastic approach to the solution of the many-electron Schrödinger equation. Beginning with the more familiar path integral quantum Monte Carlo methods we then go on to describe the origins and details of the DMQMC algorithm which grants access to the full thermal density matrix.

In Chapter 3 we detail how DMQMC can be adapted to simulate WDM, focussing on the warm dense uniform electron gas. We show why it is necessary to go beyond the original formulation of the algorithm when simulation the warm dense electron gas and how this new algorithm grants access to the evaluation of the Helmholtz free energy. We also develop new finite basis set correction techniques which allow us to overcome the prohibitively large plane wave bases required to converge the total energy at elevated temperatures. Finally, we compare to some accurate benchmark results for a small system of 4 spin-polarised electrons.

In Chapter 4 we discuss how a reliable, systematic approximation can be applied to DMQMC so that larger system sizes can be simulated. Using this approximation we resolve the significant disagreement present between exchange-correlation energies calculated using two different path integral Monte Carlo methods [11, 12]. In the pro-

cess we check, in some detail, the reliability of various approximations present in the algorithm in different regimes. We provide new data in the intermediate density regime for a spin polarised electron gas of $N = 33$ electrons which will be of use in future parametrisations of exchange-correlation free energy functionals necessary for thermal DFT.

Finally, in Chapter 5, we discuss how we can reach the thermodynamic limit using DMQMC data. In analysing the origin of existing analytic finite-size corrections we suggest an improved and general scheme of developing corrections which are valid in the warm dense regime. By performing extensive DMQMC simulations, we demonstrate the power of these corrections, before comparing to existing approximate functional fits for the thermodynamic properties of the uniform electron gas.

Chapter 1

Quantum Statistical Mechanics

This thesis is concerned with the description of matter far above the ground state. In this regime we no longer seek a single wavefunction but rather a statistical mixture of them, bringing us to the realm of quantum statistical mechanics [13–16]. In this chapter we will not aim to cover the field in its entirety but rather give the basics required to orient the reader and to provide a reference point for later discussion when certain methods will be applied to specific problems.

1.1 The Many-Electron Hamiltonian

Although most of the methods outlined in this thesis are fairly generic and can in principle be applied to a variety of model Hamiltonians, here we are interested in describing realistic warm dense matter made up of electrons and nuclei. The non-relativistic, many-particle Hamiltonian describing a system of electrons and nuclei interacting via the Coulomb interaction is given (in Hartree atomic units) as

$$\begin{aligned}\hat{H} = & -\frac{1}{2}\sum_i \hat{\nabla}_i^2 + \frac{1}{2}\sum_{i,j\neq i}\frac{1}{|\hat{\mathbf{r}}_i - \hat{\mathbf{r}}_j|} - \sum_{i,\alpha}\frac{Z_\alpha}{|\hat{\mathbf{r}}_i - \hat{\mathbf{R}}_\alpha|} \\ & -\frac{1}{2m_\alpha}\sum_\alpha \hat{\nabla}_\alpha^2 + \frac{1}{2}\sum_{\alpha,\beta\neq\alpha}\frac{Z_\alpha Z_\beta}{|\hat{\mathbf{R}}_\alpha - \hat{\mathbf{R}}_\beta|}\end{aligned} \tag{1.1}$$

$$\hat{H} = \hat{T} + \hat{V}_{ee} + \hat{V}_{en} + \hat{T}_n + \hat{V}_{nn}, \tag{1.2}$$

where Z_α and m_α are the charge and mass of nucleus α and $\mathbf{r}_i$ and $\mathbf{R}_\alpha$ are the positions of electron i and nucleus α respectively. Since the masses of the nuclei are roughly a few thousand times greater than that of the electrons, we can further make the Born-Oppenheimer approximation which allows us to effectively decouple the motion of the electrons from the nuclei [17, 18]. We can then solve Eq. (1.2) using quantum mechanics for a fixed set of nuclear coordinates, before moving the nuclei around, typically classically, in an external electronic potential energy landscape [19]. Thus, for the remainder of this thesis, we want to find methods capable of simulating a system of electrons at

fixed nuclear positions, $\{\mathbf{R}_\alpha\}$, governed by

$$\hat{H} = \hat{T} + \hat{V}_{ee} + \hat{V}_{en}(\{\mathbf{R}_\alpha\}) + \text{constant}, \tag{1.3}$$

at non-zero temperature.

1.2 The Density Matrix

In the standard interpretation of quantum mechanics, an isolated system can be completely described by a state vector $|\psi\rangle$ which lives in an abstract vector space $\mathcal{H}$ called the Hilbert space. Measurable observables are related to the eigenvalues, o_j, of Hermitian operators $\hat{O}$, i.e., $\hat{O}|o_j\rangle = o_j|o_j\rangle$, where $|o_j\rangle$ are the eigenstates of $\hat{O}$. If the system is in such a pure state, then we can thus write the expectation value of any operator as

$$O = \langle\psi|\hat{O}|\psi\rangle \tag{1.4}$$

$$= \sum_j o_j |\langle\psi|o_j\rangle|^2, \tag{1.5}$$

which represents the average value of many measurements of O on the state $|\psi\rangle$. Since the effects of temperature are unimportant in most experimental settings, it usually suffices to find the lowest energy eigenstate of the Hamiltonian $\hat{H}$ to describe the static properties of matter. Thus the problem is to find the wavefunction which minimises the energy functional

$$E[\psi] = \langle\psi|\hat{H}|\psi\rangle - \lambda(\langle\psi|\psi\rangle - 1), \tag{1.6}$$

where λ is a Lagrange multiplier imposing the constraint that the wavefunction is normalised.

At $T > 0$ we can no longer describe the system by a single wavefunction (a pure state) but rather need a statistical mixture of them. Such an ensemble of states can be described using the density operator[1]

$$\hat{\rho} = \sum_i p_i |\phi_i\rangle\langle\phi_i|, \tag{1.7}$$

where the sum is over any complete, orthonormal set of states $|\phi_i\rangle$, and p_i is the probability the system is in state $|\phi_i\rangle$ ($\sum_i p_i = 1$). From Eq. (1.7) we have that $\hat{\rho}^\dagger = \hat{\rho}$, $\mathrm{Tr}\,[\hat{\rho}] = 1$ and $\hat{\rho}$ is positive semidefinite (its eigenvalues are greater than or equal to zero). With this definition we can now calculate expectation values as

$$O = \langle\hat{O}\rangle = \sum_i p_i \langle\phi_i|\hat{O}|\phi_i\rangle \tag{1.8}$$

[1] The name density matrix and density operator will often be used interchangeably throughout this text.

$$= \mathrm{Tr}\left[\hat{\rho}\hat{O}\right]. \tag{1.9}$$

To go any further we need to specify the probabilities of a particular state being occupied in Eq. (1.7) which depends on the problem at hand. The canonical ensemble, which describes a system of fixed density in thermal equilibrium at temperature T, is well suited to describing warm dense matter. To determine the values of p_i we can try to find that $\hat{\rho}$ of the form Eq. (1.7) which maximises the von-Neumann entropy

$$S \equiv -k_B \sum_i p_i \log p_i \tag{1.10}$$

$$= -k_B \mathrm{Tr}\left[\hat{\rho}\log\hat{\rho}\right], \tag{1.11}$$

where k_B is Boltzmann's constant, subject to the constraints that the internal energy

$$U = \mathrm{Tr}\left[\hat{\rho}\hat{H}\right], \tag{1.12}$$

is constant and $\hat{\rho}$ has unit trace. Carrying this out, one finds the usual canonical density operator

$$\hat{\rho} = \frac{1}{Z}e^{-\beta\hat{H}}, \tag{1.13}$$

where $\beta = 1/k_B T$ is the inverse temperature and

$$Z = \mathrm{Tr}\left[e^{-\beta\hat{H}}\right], \tag{1.14}$$

is called the partition function. If we work in the basis of eigenstates of $\hat{H}$, $\{|\psi_i\rangle\}$, then we can identify the usual Boltzmann factors $p_i = e^{-\beta E_i}/Z$, where $\hat{H}|\psi_i\rangle = E_i|\psi_i\rangle$.

An important property of $\hat{\rho}$ is that it minimises the Helmholtz free energy

$$F = U - TS, \tag{1.15}$$

$$F[\hat{\rho}] = \mathrm{Tr}\left[\hat{\rho}\hat{H}\right] + \frac{1}{\beta}\mathrm{Tr}\left[\hat{\rho}\log\hat{\rho}\right]. \tag{1.16}$$

Thus, unlike at $T = 0$ where the internal energy is variational with respect to an approximate wavefunction, here it is now the free energy which is variational with respect to some trial density matrix $\hat{\rho}'$, i.e.,

$$F[\hat{\rho}] \leq F[\hat{\rho}']. \tag{1.17}$$

The goal is now to find ways of determining good approximate density matrices.

The most direct route to evaluating the matrix exponential appearing in Eq. (1.13) requires us to first diagonalise $\hat{H}$ by finding all its eigenstates and eigenvalues before we

can write $\hat{\rho}$ as

$$\hat{\rho} = \sum_i \frac{e^{-\beta E_i}}{Z} |\psi_i\rangle\langle\psi_i|. \tag{1.18}$$

One generic way to diagonalise $\hat{H}$ is to expand $|\psi\rangle$ in a many-particle basis as

$$|\psi\rangle = \sum_{\mathbf{i}} c_{\mathbf{i}} |D_{\mathbf{i}}\rangle, \tag{1.19}$$

where $\mathbf{i}$ is the index associated with some many-particle configuration, $c_{\mathbf{i}}$ are expansion coefficients and $|D_{\mathbf{i}}\rangle$ are many-particle states (strictly N-particle states if we work in the canonical ensemble). This allows us to write the eigenvalue problem in matrix form

$$\sum_{\mathbf{k}} H_{\mathbf{ik}} c_{\mathbf{k}} = E c_{\mathbf{i}}, \tag{1.20}$$

which can then be diagonalised on a computer. To do this numerically we need to define the many particle states appearing in Eq. (1.19), bringing us to the topic of second quantisation.

1.3 Second Quantisation

Electrons are (to the best of our knowledge) elementary, charged and, most importantly for the present discussion, spin-1/2 particles. It follows from the spin-statistics theorem in relativistic quantum mechanics[2] that they therefore obey Fermi-Dirac statistics and that the many-electron wavefunction is antisymmetric under the exchange of any two electrons. Electrons therefore obey the Pauli exclusion principle; i.e., any quantum state can only be occupied with at most one electron. For fermionic systems, Slater determinants form a basis of many-particle states which obey these two necessary conditions

$$\Psi(\mathbf{R}) = \langle \mathbf{R}|D\rangle = \frac{1}{\sqrt{N!}} \begin{vmatrix} \varphi_1(\mathbf{r}_1) & \varphi_1(\mathbf{r}_2) & \cdots \\ \varphi_2(\mathbf{r}_1) & \varphi_2(\mathbf{r}_2) & \cdots \\ \vdots & & \end{vmatrix}, \tag{1.21}$$

where $\mathbf{R} = \{\mathbf{r}_i\} = \{\mathbf{r}_1, \mathbf{r}_2, \cdots\}$ is a set of electronic coordinates and $\varphi_i(\mathbf{r})$ are some one-electron wavefunctions. The determinant structure ensures that $\Psi(\dots, \mathbf{r}_i, \mathbf{r}_j, \dots) = -\Psi(\dots, \mathbf{r}_j, \mathbf{r}_i, \dots)$ and $\Psi(\dots, \mathbf{r}_i, \mathbf{r}_j = \mathbf{r}_i, \dots) = 0$.

Slater determinants are somewhat unwieldy beasts, so throughout this thesis we typically use the language of second quantisation, where we instead work with the occupation number representation and introduce the fermionic creation ($\hat{c}_i^\dagger$) and annihilation ($\hat{c}_i$) operators which create or annihilate an electron in state i. Here i may label multi-

[2] That this follows over to the non-relativistic theory used here is largely justified as being consistent with observation.

ple quantum numbers, for instance the state's momentum $\mathbf{k}$ and spin σ. The fermionic nature of an electron is encoded in the anticommutation relations

$$\{\hat{c}_i^\dagger, \hat{c}_j\} = \delta_{ij} \tag{1.22}$$

$$\{\hat{c}_i, \hat{c}_j\} = \{\hat{c}_i^\dagger, \hat{c}_j^\dagger\} = 0, \tag{1.23}$$

where $\{\hat{A}, \hat{B}\} = \hat{A}\hat{B} + \hat{B}\hat{A}$ is the usual anti-commutator and δ_{ij} is the Kronecker delta symbol. We can then write

$$|D_\mathbf{i}\rangle = |n_0 n_1 \cdots\rangle \tag{1.24}$$

$$= \prod_i (\hat{c}_i^\dagger)^{n_i} |0\rangle, \tag{1.25}$$

where $n_i \in \{0, 1\}$ for fermions and $|0\rangle$ is the vacuum state with no particles. The action of the creation and annihilation operators on a particular state can be expressed as

$$\hat{c}_i^\dagger |n_1, \cdots, 0_i, \cdots\rangle = (-1)^{\sum_{j<i} n_j} |n_1, \cdots, 1_i, \cdots\rangle \tag{1.26}$$

$$\hat{c}_i |n_1, \cdots, 1_i, \cdots\rangle = (-1)^{\sum_{j<i} n_j} |n_1, \cdots, 0_i, \cdots\rangle \tag{1.27}$$

where the phase factor accounts for anticommuting the creation or annihilation operators through operators associated with occupied states before site i. Here we have assumed that creation operators are ordered from left to right in increasing value of i, although there are clearly other choices. The result of a creation (annihilation) operator on an occupied (unoccupied) state is zero.

With all this machinery we can write the many-electron Hamiltonian in second quantised form:

$$\hat{H} = \sum_{ij} t_{ij} \hat{c}_i^\dagger \hat{c}_j + \frac{1}{2} \sum_{ijkl} v_{ijkl} \hat{c}_i^\dagger \hat{c}_j^\dagger \hat{c}_l \hat{c}_k, \tag{1.28}$$

where

$$t_{ij} = \int d\mathbf{r}\ \varphi_i^*(\mathbf{r}) \left(-\frac{1}{2} \hat{\nabla}_\mathbf{r}^2 - \sum_\alpha \frac{Z_\alpha}{|\mathbf{r} - \mathbf{R}_\alpha|} \right) \varphi_j(\mathbf{r}), \tag{1.29}$$

$$v_{ijkl} = \int\int d\mathbf{r}\ d\mathbf{r}'\ \varphi_i^*(\mathbf{r}) \varphi_j^*(\mathbf{r}') \frac{1}{|\mathbf{r} - \mathbf{r}'|} \varphi_k(\mathbf{r}) \varphi_l(\mathbf{r}'), \tag{1.30}$$

we have dropped any constant factors and the sums are over all $ijkl$ in our single-particle basis set. Note that $\hat{H}$ can be applied to a state with an arbitrary number of electrons. All that remains is to work out the matrix elements appearing in Eq. (1.20). As the Hamiltonian only contains one and two-body interactions we just need to consider matrix elements between determinants which differ by at most two states. Keeping track of sign changes when evaluating matrix elements of the Hamiltonian is a bit tricky so

it helps to work out an explicit example.

Consider the matrix element $\langle D|\hat{H}|D_{ij}^{ab}\rangle$, where $|D_{ij}^{ab}\rangle = \hat{c}_a^\dagger \hat{c}_b^\dagger \hat{c}_j \hat{c}_i |D\rangle$, is a determinant where the electrons in orbitals $(ij) \in |D\rangle$ have been excited to the unoccupied orbitals (ab). Note that the ordering of creation and annihilation operators in this definition is very important. In this case the matrix element of the one-body term is zero so that it remains to work out

$$\langle D| \left(\frac{1}{2} \sum_{pqrs} v_{pqrs} \hat{c}_p^\dagger \hat{c}_q^\dagger \hat{c}_s \hat{c}_r \right) \hat{c}_a^\dagger \hat{c}_b^\dagger \hat{c}_j \hat{c}_i |D\rangle, \tag{1.31}$$

where we have relabelled the indices in the sum for clarity. As we work with orthogonal Slater determinants, the only terms in this sum which will yield a non-zero contribution are those for which the long string of operators yields ± 1. Inspecting the terms in Eq. (1.31) the only possible candidates are:

$$\begin{aligned} \langle D| \frac{1}{2} \Big(& v_{ijab} \hat{c}_i^\dagger \hat{c}_j^\dagger \hat{c}_b \hat{c}_a + v_{ijba} \hat{c}_i^\dagger \hat{c}_j^\dagger \hat{c}_a \hat{c}_b + \\ & v_{jiab} \hat{c}_j^\dagger \hat{c}_i^\dagger \hat{c}_b \hat{c}_a + v_{jiba} \hat{c}_j^\dagger \hat{c}_i^\dagger \hat{c}_a \hat{c}_b \Big) \hat{c}_a^\dagger \hat{c}_b^\dagger \hat{c}_j \hat{c}_i |D\rangle. \end{aligned} \tag{1.32}$$

Focussing on the first term we find that

$$\begin{aligned} v_{ijab} \hat{c}_i^\dagger \hat{c}_j^\dagger \hat{c}_b \hat{c}_a \hat{c}_a^\dagger \hat{c}_b^\dagger \hat{c}_j \hat{c}_i |D\rangle &= (-1)^6 (-1)^4 v_{ijab} \hat{c}_b \hat{c}_a \hat{c}_a^\dagger \hat{c}_b^\dagger \hat{c}_j^\dagger \hat{c}_j \hat{c}_i^\dagger \hat{c}_i |D\rangle & (1.33) \\ &= (-1)^2 v_{ijab} \hat{c}_a \hat{c}_a^\dagger \hat{c}_b \hat{c}_b^\dagger |D\rangle & (1.34) \\ &= v_{ijab} (1 - \hat{c}_a^\dagger \hat{c}_a)(1 - \hat{c}_b^\dagger \hat{c}_b) |D\rangle & (1.35) \\ &= v_{ijab} |D\rangle, & (1.36) \end{aligned}$$

where we have use the fact that $\hat{c}_a^\dagger \hat{c}_a |D\rangle = \hat{n}_a |D\rangle = 0$ since a and b are unoccupied in $|D\rangle$. Noting that $v_{ijab} = v_{jiba}$, $v_{ijba} = v_{jiab}$, $\hat{c}_i^\dagger \hat{c}_j^\dagger \hat{c}_a \hat{c}_b = -\hat{c}_i^\dagger \hat{c}_j^\dagger \hat{c}_b \hat{c}_a$ and $\hat{c}_j^\dagger \hat{c}_i^\dagger \hat{c}_a \hat{c}_b = (-1)^2 \hat{c}_i^\dagger \hat{c}_j^\dagger \hat{c}_b \hat{c}_a$, we find

$$\langle D|\hat{H}|D_{ij}^{ab}\rangle = v_{ijab} - v_{ijba}. \tag{1.37}$$

A similar analysis shows that

$$\langle D|\hat{H}|D\rangle = \sum_i t_{ii} + \sum_{i<j} \left(v_{ijij} - v_{ijji} \right), \tag{1.38}$$

and

$$\langle D|\hat{H}|D_i^a\rangle = t_{ia} + \sum_j \left(v_{ijaj} - v_{ijja} \right), \tag{1.39}$$

where $|D_i^a\rangle = \hat{c}_a^\dagger \hat{c}_i |D\rangle$ and the sums are over occupied states. Eqs. (1.37) to (1.39) are the usual Slater-Condon rules [20, 21].

So, given some set of one-electron orbitals $\varphi_i(\mathbf{r})$, the challenge is then to work out the necessary integrals and diagonalise $\hat{H}$. From an implementation point of view, the occupation number representation has the added benefit in that the action of the creation and annihilation operators and the representation of the states themselves can be achieved by exploiting the bit representation of integers [22].

While this direct, exact diagonalisation (or full configuration interaction (FCI) in quantum chemistry [23]), route is extremely valuable for benchmarking purposes as it gives the exact result for an N-particle system in a finite basis set, it is extremely limited by system size. To see this, consider the expansion Eq. (1.19) for a system of N spin-polarised electrons using a single-particle basis of size M which will contain $\binom{M}{N}$ terms. So if we want to simulate, say, 20 electrons in a basis of 100 orbitals we would need to store about 5×10^{20} numbers to represent the wavefunction or roughly 7×10^7 terabytes of data (assuming double precision floating numbers), a good deal outside of current hard drive capacities. This exponentially scaling (in both M and N) brute force approach is manifestly hopeless if we wish to reach the thermodynamic limit[3] and alternative, generally approximate methods are necessary.

1.4 Free Electrons

The simplest approximation we can make is to ignore the fact that the electrons interact at all and study the free electron model which, despite its simplicity, is rather informative. We can write the Hamiltonian describing a collection of non-interacting electrons as

$$\hat{H}^0 = \sum_i \varepsilon_i \hat{c}_i^\dagger \hat{c}_i \tag{1.40}$$

$$= \sum_i \varepsilon_i \hat{n}_i, \tag{1.41}$$

where ε_i is some single-particle eigenvalue. Using the language of the previous section we can write the ideal (un-normalised) N-body density matrix in the canonical ensemble as

$$\hat{\rho}^0 = \sum_{\mathbf{i}} e^{-\beta E_{\mathbf{i}}} |D_{\mathbf{i}}(N)\rangle\langle D_{\mathbf{i}}(N)|, \tag{1.42}$$

which is now purely diagonal and where $E_{\mathbf{i}} = \langle D_{\mathbf{i}}|\hat{H}^0|D_{\mathbf{i}}\rangle = \sum_{i_{\text{occ}}} \varepsilon_i$. Eq. (1.42) is actually quite a complicated expression: its evaluation requires us to enumerate all possible $\binom{M}{N}$ configurations and sum up their single particle eigenvalues to evaluate the Boltzmann factors which soon becomes nontrivial.

It is normal to instead work in the grand canonical ensemble where now T and V are

[3] The thermodynamic limit is defined as the limit $N \to \infty$, $V \to \infty$, where V is the system volume, such that $N/V \to$ constant $= n$ the system density.

fixed but N is allowed to vary. This leads to the grand canonical density matrix

$$\hat{\rho}_{\mathrm{GC}} = e^{-\beta(\hat{H}-\mu\hat{N})}, \tag{1.43}$$

where μ is the chemical potential and $\hat{N} = \sum_i \hat{n}_i$ is the number operator. We can then write

$$\hat{\rho}^0_{\mathrm{GC}} = \sum_{N=0}^{\infty} e^{\beta\mu N} \sum_{\mathbf{i}} e^{-\beta E_{\mathbf{i}}(N)} |D_{\mathbf{i}}(N)\rangle\langle D_{\mathbf{i}}(N)|, \tag{1.44}$$

which actually simplifies matters considerably so that the partition function can be written as

$$\begin{aligned} Z^0_{\mathrm{GC}} &= \sum_N \sum_{\mathbf{i}} e^{-\beta(E_{\mathbf{i}}(N)-\mu N)} && (1.45)\\ &= \sum_{n_1}\sum_{n_2}\sum_{n_3} \ldots e^{-\beta(\varepsilon_1-\mu)n_1} e^{-\beta(\varepsilon_2-\mu)n_2} e^{-\beta(\varepsilon_3-\mu)n_3} \ldots && (1.46)\\ &= \prod_i \sum_{n_i=0}^{1} \left(e^{-\beta(\varepsilon_i-\mu)}\right)^{n_i} && (1.47)\\ &= \prod_i \left(1 + e^{-\beta(\varepsilon_i-\mu)}\right), && (1.48) \end{aligned}$$

which is a simple product of contributions from the one-electron states. This simplification is so considerable that most approximate theories we will discuss work in the grand canonical ensemble as a result. Although observables calculated in either the canonical or grand canonical ensemble will yield identical results in the thermodynamic limit, for finite systems they will generally differ.

With Eq. (1.48) we can readily derive all of the thermodynamic properties of the system, usually from the grand potential

$$\begin{aligned} \Omega &= -k_B T \log Z_{\mathrm{GC}} && (1.49)\\ &= F - \mu N. && (1.50) \end{aligned}$$

The mean number of particles, $N = \langle\hat{N}\rangle$, can be found from usual thermodynamic relations [14] and Eq. (1.50) as

$$N(T, V, \mu) = -\left(\frac{\partial\Omega}{\partial\mu}\right)_{TV}, \tag{1.51}$$

which for the simple case of Eq. (1.48) yields

$$N = \sum_i \frac{1}{e^{\beta(\varepsilon_i-\mu)} + 1}, \tag{1.52}$$

where we can identify the usual Fermi factor

$$f_i = \frac{1}{e^{\beta(\varepsilon_i - \mu)} + 1}. \tag{1.53}$$

As we are normally interested in systems at a fixed density, it is typical to eliminate μ in favour of the density $n = N/V$ which requires inversion of Eq. (1.51). This is rather simply done for the non-interacting case by finding the roots of Eq. (1.52), but is more involved for systems of interacting particles [14, 24–27].

Two important limits can be investigated using Eq. (1.53). As $T \to 0$, $f_i \to \theta(E_F - \varepsilon_i)$, where θ is the usual Heaviside step function and E_F is the Fermi energy. In the case of a three dimensional free electron gas we have[4]

$$E_F = \frac{1}{2}\left(\frac{9\pi}{2(2-\zeta)}\right)^{2/3}\frac{1}{r_s^2}, \tag{1.54}$$

where $\zeta = (N_\uparrow - N_\downarrow)/N$ is the spin polarisation, N_σ is the number of spin σ electrons and we have introduced the Wigner-Seitz parameter

$$r_s = \left(\frac{3}{4\pi n a_0^3}\right)^{1/3}, \tag{1.55}$$

as a measure of the mean interparticle spacing $\tilde{r}_s = r_s a_0$ and a_0 is the Bohr radius. Thus at zero temperature all states below E_F are occupied and those above it are unoccupied, defining the usual Fermi surface. In this regime quantum effects are important and the system is said to be degenerate.

In the opposite limit of $T \to \infty$ we have $f_i \approx e^{-\beta(\varepsilon_i - \mu)}$ and the system obeys classical Boltzmann statistics where quantum effects are no longer important. It is typical to measure the degree of degeneracy in terms of the degeneracy temperature

$$\Theta = T/T_F, \tag{1.56}$$

where $T_F = E_F/k_B$ is the Fermi temperature. Typical values of T_F for normal metals are about 10^4 K [29]. The degenerate and non-degenerate limits correspond to $\Theta \ll 1$ and $\Theta \gg 1$ respectively.

Apart from serving as a pedagogical tool, the free gas with Fermi-Dirac statistics accounts for a number of important phenomena. A non-exhaustive list includes: the stability of matter [30], the linear dependence at low temperature of the specific heat capacity of normal metals and the existence of a non-zero 'Fermi' pressure which is responsible for the stability of white dwarf stars [15]. What is even more remarkable is that even in the presence of electron-electron interactions, most normal Fermi systems

[4] Note that the factor $(2-\zeta)$ is understood to be valid for the fully spin polarised $\zeta = 1$ and unpolarised $\zeta = 0$ gases only, since these are the only cases we will address in this thesis. The general case can be found in Ref. [28].

behave *qualitatively* like the ideal system, with even the concept of a Fermi surface persisting [28, 31]. This Landau Fermi liquid theory [32, 33] has formed the bedrock of most our attempts to describe realistic matter, which are generally based upon a one-electron picture. Indeed, any metal which behaves other than a normal Fermi liquid usually results in exotic physics.

Despite these successes in describing the qualitative features of the problem, electrons do interact with one another and the electron-electron energy is typically the same order of magnitude as the other terms in Eq. (1.3). We can measure the relative importance of the Coulomb interaction in terms of the classical Coulomb coupling parameter[5]

$$\Gamma = \frac{1}{r_s a_0 k_B T}, \tag{1.57}$$

which is the ratio of the classical Coulomb energy and kinetic energy of the electrons. The system is said to be strongly coupled if $\Gamma \geq 1$. (At $T = 0$, Γ becomes meaningless, and here r_s takes over as a measure of electron-electron interaction strength relative to the quantum mechanical kinetic energy.) In the warm dense regime we have $\Theta \approx \Gamma \approx r_s \approx 1$ so that quantum mechanical and interaction effects are important and no small coupling parameter exists. Thus, if we want to have quantitative agreement with experiment, we need methods capable of treating both thermal and electron-electron effects.

1.5 Hartree–Fock Theory

Hartree–Fock theory is perhaps the simplest approach to treating interacting electronic systems, and forms the basis of most sophisticated quantum chemical methods (which are often described as 'post Hartree–Fock' methods). At $T = 0$ the problem is to find the single Slater determinant, from a search over all possible candidates, which minimises the total energy. In practice, the equations which result from this problem are non-local integral equations which are not easy to solve, except in a few cases. The lesser known thermal Hartree–Fock approximation [14, 35, 36] poses a similar challenge: given a trial density matrix of the form of Eq. (1.44), find the one which minimises the Helmholtz free energy (grand potential)[6].

Inserting Eq. (1.44) into Eq. (1.16) with $\hat{H}$ given by Eq. (1.3) and carrying out the variation one finds the self-consistency conditions for the eigenvalues and orbitals

[5] There is a similar coupling parameter for the nuclei. A good discussion of the various regimes can be found in Ref. [34].

[6] The thermal Hartree–Fock equations can also be derived from a Green's function picture [14] much like for $T = 0$. Mermin showed in Ref. [35] that the two pictures were equivalent.

(assuming they are eigenfunctions of σ_z)

$$\begin{aligned}\varepsilon_{i\sigma}\varphi_{i\sigma}(\mathbf{r}) = &\left(-\frac{1}{2}\nabla^2 + V_{en}(\mathbf{r})\right)\varphi_{i\sigma}(\mathbf{r}) + \sum_{j\neq i,\sigma'} f_{j\sigma'}\int d\mathbf{r}'\,\frac{|\varphi_{j\sigma'}(\mathbf{r}')|^2}{|\mathbf{r}-\mathbf{r}'|}\varphi_{i\sigma}(\mathbf{r}) \\ &- \sum_{j\neq i,\sigma'}\delta_{\sigma\sigma'}f_{j\sigma'}\int d\mathbf{r}'\,\frac{\varphi^*_{j\sigma'}(\mathbf{r}')\varphi_{i\sigma}(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|}\varphi_{j\sigma'}(\mathbf{r}),\end{aligned} \tag{1.58}$$

where, to add to the complexity, the Fermi factors are evaluated using the thermal Hartree–Fock eigenvalues. Note that both $\varepsilon_{i\sigma}$ and $\varphi_{i\sigma}(\mathbf{r})$ now depend on temperature.

Given the gross failings of Hartree–Fock theory for the ground state, there is not much hope for $T > 0$. Nevertheless, its variational nature as the best possible single-particle-like density matrix, marks it as a useful starting point for comparison purposes, although its application to realistic systems is in its relative infancy [37, 38].

1.6 Density Functional Theory

Mean field theories, such as Hartree–Fock, are bound to exhibit shortcomings in describing the many-body nature of matter. One notable exception to this is density functional theory (DFT), which is unusual in that, in principle, one can find the exact free energy of the system by solving a mean field problem. Between 1965 and 1967 Hohenberg, Kohn and Sham [39, 40] demonstrated how this is possible and in between Mermin [41] generalised the procedure for $T \geq 0$. Here, in keeping with the theme of this chapter, we will only discuss the Mermin-Kohn-Sham (MKS) formalism of which $T = 0$ is a special case[7,8].

Following the Lieb-Levy constrained search route to thermal DFT [47–49] we begin by defining the functional

$$O[n(\mathbf{r})] \equiv \min_{\hat{\rho}\to n(\mathbf{r})} \mathrm{Tr}\left[\hat{O}\hat{\rho}\right], \tag{1.59}$$

for any reasonable operator $\hat{O}$ and here the minimisation is carried out by searching over all positive definite, unit trace, density operators $\hat{\rho}$ which result in the density $n(\mathbf{r})$. Note that the density can be determined from

$$n(\mathbf{r}) = \mathrm{Tr}\left[\hat{\rho}\hat{\psi}^\dagger(\mathbf{r})\hat{\psi}(\mathbf{r})\right], \tag{1.60}$$

[7] Note that MKS DFT is often used when applying DFT to the problem of ground state metallic systems. Although the system of interest is at $T = 0$, it is often easier to converge the DFT calculations with respect to k-points by using a small, but finite value of T and then extrapolating the results to $T = 0$ [42, 43].

[8] For some recent theoretical developments in thermal DFT see Refs. [44–46].

the fermionic field operators are defined as

$$\hat{\psi}(\mathbf{r}) \equiv \sum_i \varphi_i(\mathbf{r})\hat{c}_i, \tag{1.61}$$

and the trace over spin indices in Eq. (1.60) is implied. Mermin showed that from the Gibbs variational principle Eq. (1.17), the grand potential functional[9]

$$\Omega[n(\mathbf{r})] = \min_{\hat{\rho}\to n(\mathbf{r})} \mathrm{Tr}\left[\hat{\rho}\left(\hat{T} + \hat{V}_{ee} + \hat{V}_{en} + \frac{1}{\beta}\log\hat{\rho} - \mu\hat{N}\right)\right], \tag{1.62}$$

is minimised when $n(\mathbf{r})$ is equal to true density with the true grand canonical density matrix determined with $\hat{V}_{en}$. It follows then that we can write

$$\Omega[n(\mathbf{r})] = F[n(\mathbf{r})] + \int d\mathbf{r}\ (V_{en}(\mathbf{r}) - \mu)n(\mathbf{r}), \tag{1.63}$$

where

$$F[n(\mathbf{r})] = \min_{\hat{\rho}\to n(\mathbf{r})} \mathrm{Tr}\left[\hat{\rho}\left(\hat{T} + \hat{V}_{ee} + \frac{1}{\beta}\log\hat{\rho}\right)\right], \tag{1.64}$$

is a *universal* functional of the density which does not (explicitly) depend on $V_{en}(\mathbf{r})$. This means that if $F[n(\mathbf{r})]$ were known we could solve for the properties of any material exactly just by varying the density, which is a function of 3 variables (rather than $6N$ in the case of the many-electron density matrix).

Despite the profound nature of these theorems, they are not constructive as the universal functional is unknown and determining it would require a full solution of the many-electron problem which leaves us where we were. The power of density functional theory comes from the fact that very simple approximations to the unknown functional work remarkably well and, from a practical point of view, the problem can be recast as a simple eigenvalue problem due to Kohn and Sham which can readily be solved on a computer.

The crucial step in developing the MKS scheme is to first separate out bits of the universal functional which are significant and straightforward to calculate in the hope that what remains is small and can be approximated. To achieve this, note that we can write

$$\begin{aligned}\Omega[n(\mathbf{r},\beta)] =& F_s[n(\mathbf{r})] + U_H[n(\mathbf{r})] + F_{\mathrm{xc}}[n(\mathbf{r})] \\ &+ \int d\mathbf{r}\ (V_{en}(\mathbf{r}) - \mu)n(\mathbf{r}),\end{aligned} \tag{1.65}$$

where

$$F_{\mathrm{xc}}[n(\mathbf{r})] = F[n(\mathbf{r})] - U_H[n(\mathbf{r})] - F_s[n(\mathbf{r})], \tag{1.66}$$

[9] It is also possible to formulate thermal DFT in the canonical ensemble [47].

is called the exchange-correlation free energy functional,

$$F_s[n(\mathbf{r})] = \min_{\hat{\rho}\to n(\mathbf{r})} \mathrm{Tr}\left[\hat{\rho}\left(\hat{T} + \frac{1}{\beta}\log\hat{\rho}\right)\right], \tag{1.67}$$

is the Helmholtz free energy of a non-interacting reference system with density $n(\mathbf{r})$ and

$$U_H[n(\mathbf{r})] = \frac{1}{2}\int\int d\mathbf{r}\; d\mathbf{r}'\; \frac{n(\mathbf{r})n(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|} \tag{1.68}$$

is the classical Hartree energy.

Carrying out the constrained minimisation of $F_s[n(\mathbf{r})]$ amounts to determining when

$$\mathrm{Tr}\left[\hat{\rho}\left(\hat{T} + \frac{1}{\beta}\log\hat{\rho}\right)\right] - \int d\mathbf{r}\; (V_{\mathrm{eff}}(\mathbf{r}) - \mu)n(\mathbf{r}), \tag{1.69}$$

is minimised. Here $V_{\mathrm{eff}}(\mathbf{r})$ is the Lagrange multiplier imposing the constraint that $n(\mathbf{r})$ is the exact equilibrium density of the interacting system while μ fixes the average particle number. Eq. (1.69) is of the form of the expression for the grand potential of a set of non-interacting electrons in the presence of an effective external potential $V_{\mathrm{eff}}(\mathbf{r})$. We know from the discussion of a free electron gas that a functional of the form of Eq. (1.69) is minimal when $\hat{\rho}$ is of the form of the ideal grand canonical density matrix, so that we can immediately write

$$n_s(\mathbf{r}) = \sum_i f_i|\varphi_i(\mathbf{r})|^2, \tag{1.70}$$

$$T[n_s(\mathbf{r})] = \sum_i f_i \int d\mathbf{r}\varphi_i^*(\mathbf{r})\left(-\frac{1}{2}\nabla_{\mathbf{r}}^2\varphi_i(\mathbf{r})\right), \tag{1.71}$$

and

$$S[n_s(\mathbf{r})] = -k_B\sum_i \left[f_i\log f_i + (1-f_i)\log(1-f_i)\right], \tag{1.72}$$

where the single-particle states are determined from the solving the Schrödinger-like equation

$$\left(-\frac{1}{2}\nabla_i^2 + V_{\mathrm{eff}}[n(\mathbf{r})]\right)\varphi_i(\mathbf{r}) = \varepsilon_i\varphi_i(\mathbf{r}). \tag{1.73}$$

V_{eff} is determined by requiring $\delta\Omega = 0$, and carrying this out one finds

$$V_{\mathrm{eff}}[n(\mathbf{r})] = V_H[n(\mathbf{r})] + V_{en}(\mathbf{r}) - \mu + V_{\mathrm{xc}}[n(\mathbf{r})], \tag{1.74}$$

where we have

$$\frac{\delta F_{\mathrm{xc}}[n(\mathbf{r})]}{\delta n(\mathbf{r})} = V_{\mathrm{xc}}[n(\mathbf{r})], \tag{1.75}$$

and so on.

So, after all that, we arrive at the MKS scheme for thermal DFT, wherein we can specify a guess for $n(\mathbf{r})$, from which we can construct $V_{\mathrm{eff}}[n(\mathbf{r})]$ from Eq. (1.74) insert this into Eq. (1.73) and solve for the Kohn-Sham eigenvalues and eigenvectors before

finally constructing the density via Eq. (1.70). This procedure can be repeated by using this output density as a new input and iterating until self-consistency is reached.

One of the primary limitations of DFT is that the above all relies on a good approximation for $F_{\rm xc}[n(\mathbf{r})]$ about which we know little. One of the earliest and still widely used approaches at $T = 0$ was to assume this was given by the exchange correlation energy of a uniform electron gas (UEG). This is the local density approximation (LDA) where, at $T \geq 0$, one has

$$F_{\rm xc}[n(\mathbf{r})] = \int d\mathbf{r} f_{\rm xc}^{\rm UEG}(n, \beta) n(\mathbf{r}), \tag{1.76}$$

where $f_{\rm xc}^{\rm UEG}$ is the exchange-correlation free energy per-electron of a uniform electron gas (UEG) of the same density and temperature as the system we wish to study. Until quite recently [11, 12, 50–55] no accurate data for $f_{\rm xc}^{\rm UEG}$ existed over the whole temperature-density plane, so a common approximation was to replace $f_{\rm xc}$ with a ground state exchange-correlation functional where the only temperature dependence arrived through the density. This amounts to an uncontrolled approximation[10] which introduces errors of unknown magnitude. For applications of DFT in most normal settings it is expected that this approximation is of little or no importance. However, in the warm dense regime, these electronic thermal contributions can grow significant relative to the total free energy and the approximation is of questionable validity [59, 60].

Another, more practical, issue with MKS-DFT is that as the temperature increases, the number of single-particle basis functions required to converge the density grows rapidly which can render it impractical at very high temperatures, although progress is being made in removing this limitation [61]. An alternative approach is orbital free DFT (OFDFT), which, as the name suggests, does not require an explicit diagonalisation of the Kohn-Sham Hamiltonian so that arbitrarily high temperatures can be treated. Unfortunately, applications of OFDFT have been somewhat limited by a lack of accurate kinetic energy density functionals. See Refs. [62, 63] for some recent developments.

1.7 Beyond Mean Field Theory

In this chapter we have attempted to outline some of the most familiar[11] methods available at $T > 0$ which could in principle be applied to the problem of warm dense matter. Going beyond mean field theories like DFT is difficult to do. One option is thermal many-body perturbation theory (MBPT) which offers in principle a systematic[12] route

[10] Note that this uncontrolled approximation is on top of another uncontrolled approximation in that $T = 0$ DFT results can often depend quite strongly on which approximate functional is used [56–58].

[11] From an electronic structure theorist's perspective.

[12] By systematic I mean that there is some complexity knob allowing us to tune how good the answer we get is. For perturbation theory this is usually not the case if done in a brute force manner. Systematically improvability is an appealing feature to have in a method, although breakthroughs in physics usually come through insight rather than brute force.

to higher accuracy and also grants access to spectral quantities which DFT cannot provide [64]. The basic idea is to systematically build up approximations (typically) to the one-particle Green's function which can in turn be related to the grand potential via coupling constant integration [14].

While MBPT has been extremely important in the development of our understanding of many-body quantum systems [24, 31, 65] it is fraught with complications. First, the number of terms one has to sum at each order in the perturbation series grows increasingly large and more complicated even with the help of Feynman diagrams leading one to truncate the series at a certain order. Truncating the series amounts to assuming that the electron-electron interaction represents a small perturbation to the non-interacting Hamiltonian which is certainly not generally the case, particularly in the warm dense regime.

Secondly, for metallic systems, one typically finds that certain terms appearing in the series are highly divergent (even at second order) so any truncated series is of dubious utility. The issue of divergences can be avoided by resumming certain classes of diagrams to all orders to obtain contributions that cancel some of the divergent terms. Again, however, this is possible analytically only in simple cases [65].

Finally, actually numerically evaluating the Green's function for an arbitrary system is rather an intricate process given the analytic structure of the Green's function. Nevertheless, the one-particle Green's function allows one to determine spectral properties of the system which are typically of the most interest to experimentalists [64]. With the advent of new algorithms and computational power there has been a renewed interest of late in applying perturbative techniques such as the random phase and GW approximation [66] in solid-state and quantum chemical contexts [67–72].

Naturally, there are many other techniques out there which could be used, including classical approximations such as Debye-Hückel theory valid at high T [14], dieletric theories [73] which are related to MBPT, the thermal analogues of quantum chemical approaches [74–77] and various other novel approaches [78, 79]. One of the main issues with most of the semi-analytical methods is that they are perturbative in some parameter. In the warm dense regime there is no small parameter so that none of these methods are valid in all regimes (see Fig. 1.1). DFT, on the other hand, lacks accurate input for reliable exchange-correlation free energy functionals. In the next chapter, we will discuss Quantum Monte Carlo (QMC) methods which offer, in principle, a non-perturbative, explicitly N-body approach to address both of these issues.

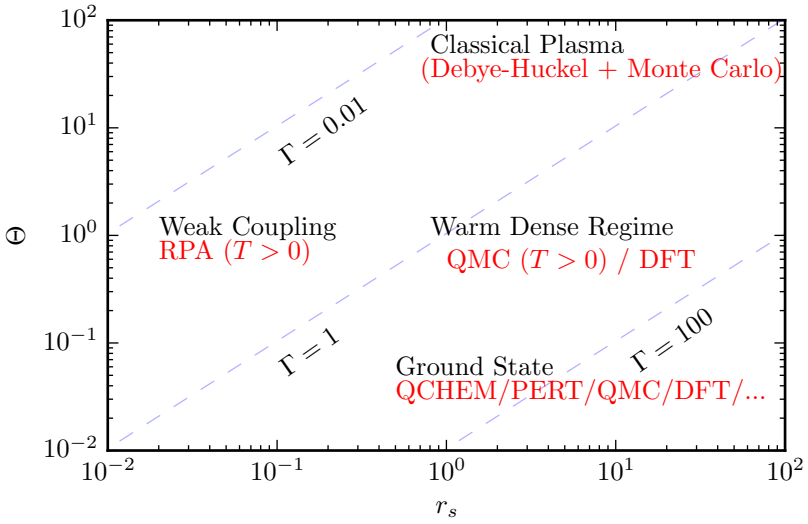

Figure 1.1: Phases of matter in the temperature density-plane and a selection of methods at our disposal. In the ground state there are a hierarchy of methods, starting from quantum chemistry (QCHEM) which provide increasing levels of accuracy. For weak coupling, but strong degeneracy, perturbation theories (such as the RPA) offer a good description, while in the classical regime one has Debye-Hückel at weak coupling and classical Monte Carlo simulations [80, 81] at strong coupling. In the warm dense regime non-perturbative approaches such as DFT and QMC are required. The dashed lines are values of constant $\Gamma \sim r_s/\Theta$.

Chapter 2

Quantum Monte Carlo Methods

Monte Carlo methods encompass a wide variety of stochastic algorithms used in nearly all areas of science to solve high dimensional, deterministic problems. In physics, we are often interested in evaluating multidimensional integrals of the type

$$I = \int \cdots \int d\mathbf{p} d\mathbf{q} \ldots d\mathbf{s} w(\mathbf{p}, \ldots, \mathbf{s}) f(\mathbf{p} \ldots \mathbf{s}), \tag{2.1}$$

for some weighting function w and some observable f. For instance, I could be the classical ensemble average of some function with $w = e^{-\beta H(\mathbf{p}\ldots\mathbf{s})}/Z$ given some classical Hamiltonian H, or a quantum mechanical expectation value with a normalised $w = |\Psi(\mathbf{R})|^2$. In both cases the dimensionality of the integral often renders it impossible to evaluate exactly using conventional algorithms. The Monte Carlo method works instead by randomly generating configurations according to the weight function w assuming it can be interpreted in some way probabilistically. If the configurations generated sample the probability distribution then, by the central limit theorem, the integral can be estimated by the mean value of N_s samples

$$I \approx \bar{I} = \frac{1}{N_s} \sum_{s}^{N_s} f(\mathbf{X}_s), \tag{2.2}$$

where $\mathbf{X}_s = (\mathbf{p}_s, \cdots, \mathbf{s}_s)$. The virtue of Monte Carlo methods is that the estimate for the standard error in $\bar{I}$

$$\sigma_{\bar{I}} = \frac{\sigma_I}{\sqrt{N_s}} = \sqrt{\frac{1}{N_s(N_s-1)} \sum_s (f(\mathbf{X}_s) - \bar{I})^2} \tag{2.3}$$

decays like $\frac{1}{\sqrt{N_s}}$ and does not (explicitly) depend on the dimensionality of the problem, unlike traditional quadrature methods.

Quantum Monte Carlo (QMC) methods utilise Monte Carlo sampling to attempt to solve the many-body problem exactly. Many different algorithms exist (and they are in some way related), but they are generally categorised as being either a variational method, where we try to optimise the energy using a trial wave function [82, 83] or density matrix [84]; a projector method, where we try to project out the ground state

wavefunction [83]; or a path integral method, where we try to sample contributions to the partition function [85].

In this chapter we will first discuss path integral Monte Carlo (PIMC), which is perhaps the most natural way of developing a QMC method capable of simulating matter at non-zero temperature before outlining how ideas from ground-state projector Monte Carlo can be used in developing an approach to solving for the thermodynamic properties of a system.

2.1 Path Integral methods

Consider the un-normalised N-body canonical density matrix in coordinate space

$$\rho(\mathbf{R}, \mathbf{R}', \beta) = \langle \mathbf{R}| e^{-\beta \hat{H}} |\mathbf{R}'\rangle, \tag{2.4}$$

where $\mathbf{R} = (\mathbf{r}_1, \mathbf{r}_2, \ldots, \mathbf{r}_N)$ is a set of electronic coordinates. Noting that

$$\hat{\rho} = \left(e^{-\frac{\beta}{P}\hat{H}}\right)^P, \tag{2.5}$$

we can write the density matrix as

$$\rho(\mathbf{R}, \mathbf{R}', \beta) = \int d\mathbf{R}_1 \; d\mathbf{R}_2 \; \ldots \mathbf{R}_{P-1} \; \rho(\mathbf{R}, \mathbf{R}_1, \tau)\rho(\mathbf{R}_1, \mathbf{R}_2, \tau) \ldots \rho(\mathbf{R}_{P-1}, \mathbf{R}', \tau), \tag{2.6}$$

where $\tau = \beta/P$ and we have inserted $P-1$ complete sets of states $\int d\mathbf{R}|\mathbf{R}\rangle\langle\mathbf{R}| = \mathbb{1}$ to split the full density matrix into $P-1$ short-time contributions. The idea here is that these short-time density matrices can be evaluated using the Suzuki-Trotter approximation [86, 87]

$$e^{-\tau\hat{H}} = e^{-\tau(\hat{H}^0 + \hat{H}')} = e^{-\tau\hat{H}^0} e^{-\tau\hat{H}'} + \mathcal{O}\left(\tau^2\right), \tag{2.7}$$

so that we can write

$$\rho(\mathbf{R}_i, \mathbf{R}_j, \tau) = \int d\mathbf{R}' \; \langle \mathbf{R}_i| e^{-\tau\hat{H}^0} |\mathbf{R}'\rangle\langle\mathbf{R}'| e^{-\tau\hat{H}'} |\mathbf{R}_j\rangle. \tag{2.8}$$

With Eq. (2.6), the partition function can then be evaluated as

$$Z = \int d\mathbf{R} \; \rho(\mathbf{R}, \mathbf{R}, \beta) \tag{2.9}$$

$$= \int d\mathbf{X} \; w(\mathbf{X}), \tag{2.10}$$

for a configuration $\mathbf{X} = (\mathbf{R}, \mathbf{R}_1, \ldots, \mathbf{R}_{P-1})$ and

$$w(\mathbf{X}) = \prod_{i=0}^{P-1} \rho(\mathbf{R}_i, \mathbf{R}_{i+1}, \tau), \tag{2.11}$$

where $\mathbf{R}_0 = \mathbf{R}_P = \mathbf{R}$. The closed nature of the partition function (state $|\mathbf{R}\rangle$ appears at either end of the expression) leads to the interpretation that Eq. (2.9) is an integral over all closed paths running from imaginary $\tau = 0$ to $\tau = \beta$.

With ρ in this form we can now evaluate expectation values (for an operator $\hat{O}$ which is diagonal in position space) as

$$\langle \hat{O} \rangle = \int d\mathbf{X} O(\mathbf{X}) \pi(\mathbf{X}), \tag{2.12}$$

for $\pi(\mathbf{X}) = w(\mathbf{X})/Z$. This $3PN$ dimensional integral is of the form of Eq. (2.1) so we turn to the Monte Carlo method to evaluate it. Since determining Z is generally impossible we use the Metropolis-Hastings algorithm, which provides a general way of sampling a distribution without knowing its normalisation [88, 89].

The Metropolis algorithm works by generating a sequence of configurations $\mathbf{X}_1 \to \mathbf{X}_2 \to \ldots$ in such a way that any new configuration only depends on the previous one and no others. Such a sequence is called a Markov chain. One can show [90] that any Markov chain which is ergodic and obeys the detailed balance condition[1] is guaranteed to converge to the equilibrium stationary distribution $\pi(\mathbf{X})$. Here ergodicity means that any configuration can be reached from any other in a finite number of steps, while the detailed balance condition states that

$$\pi(\mathbf{X}_i) P(\mathbf{X}_i, \mathbf{X}_j) = \pi(\mathbf{X}_j) P(\mathbf{X}_j, \mathbf{X}_i), \tag{2.13}$$

where $P(\mathbf{X}_i, \mathbf{X}_j)$ is the transition probability from state $\mathbf{X}_i$ to state $\mathbf{X}_j$. Thus, by choosing our transition probabilities carefully we can arrive at the desired canonical distribution function. Writing the transition probability as

$$P(\mathbf{X}_i, \mathbf{X}_j) = T(\mathbf{X}_i, \mathbf{X}_j) \times A(\mathbf{X}_i, \mathbf{X}_j), \tag{2.14}$$

where $T(\mathbf{X}_i, \mathbf{X}_j)$ is the probability of selecting the move $\mathbf{X}_i \to \mathbf{X}_j$ and $A(\mathbf{X}_i, \mathbf{X}_j)$ is the probability of accepting it, one can show that the choice of

$$A(\mathbf{X}_i, \mathbf{X}_j) = \min\left(1, \frac{\pi(\mathbf{X}_j) T(\mathbf{X}_j, \mathbf{X}_i)}{\pi(\mathbf{X}_i) T(\mathbf{X}_i, \mathbf{X}_j)}\right), \tag{2.15}$$

satisfies the detailed balance condition and hence allows us to reach the desired distribution. Importantly, only the ratio of values of the probability density π are required; a knowledge of the normalisation is unnecessary.

[1] The detailed balance criterion is sufficient to ensure convergence but not necessary.

A PIMC simulation proceeds by specifying some initial configuration $\mathbf{X}_0 = (\mathbf{R}(\tau_0), \mathbf{R}(\tau_1), \cdots, \mathbf{R}(\tau_M))$ and successively generating new configurations which are accepted or rejected with a probability proportional to the ratio of the short-time density matrices given in Eq. (2.8). The choice of update scheme and short-time propagator can significantly affect the efficiency of the algorithm and we refer the reader to the literature for specific examples [85].

2.1.1 The Sign Problem

In the discussion so far we have not assumed that the particles obey any particular statistics. To correctly account for Fermi statistics we have to explicitly sum over all possible permutations of the N particles in the expression for the partition function

$$Z = \frac{1}{N!} \sum_{\mathcal{P} \in S_N} (-1)^{\mathcal{P}} \int d\mathbf{R} \langle \mathbf{R} | e^{-\beta \hat{H}} | \mathcal{P}\mathbf{R} \rangle \tag{2.16}$$

$$= \sum_{\mathcal{P} \in S_N} \int d\mathbf{X} w(\mathbf{X}, \mathcal{P}), \tag{2.17}$$

where S_N is the permutation group and paths are now β-periodic up to a permutation of the end points. The above expression is problematic as negative weights can appear so that we can no longer interpret the integrand as a probability measure. This issue is avoided by taking the absolute values of the weights and then re-introducing the sign when evaluating expectation values. To see this, note we can write

$$\langle \hat{O} \rangle = \frac{\sum_{\mathcal{P}} \int d\mathbf{X} w(\mathbf{X}, \mathcal{P}) O(\mathbf{X}, \mathcal{P})}{\sum_{\mathcal{P}} \int d\mathbf{X} w(\mathbf{X}, \mathcal{P})} \tag{2.18}$$

$$= \frac{\sum_{\mathcal{P}} \int d\mathbf{X} |w(\mathbf{X}, \mathcal{P})| O(\mathbf{X}, \mathcal{P}) \frac{w(\mathbf{X}, \mathcal{P})}{|w(\mathbf{X}, \mathcal{P})|}}{\sum_{\mathcal{P}} \int d\mathbf{X} |w(\mathbf{X}, \mathcal{P})| \frac{w(\mathbf{X}, \mathcal{P})}{|w(\mathbf{X}, \mathcal{P})|}} \tag{2.19}$$

$$= \frac{\langle OS \rangle_+}{\langle S \rangle_+}, \tag{2.20}$$

where the average is over the positive distribution function $|w(\mathbf{X}, \mathcal{P})|$ and

$$S = \frac{w(\mathbf{X}, \mathcal{P})}{|w(\mathbf{X}, \mathcal{P})|}, \tag{2.21}$$

is the sign of a given configuration $(\mathbf{X}, \mathcal{P})$.

The notorious fermion sign problem emerges when the number of positive and negative contributions to the sum appearing above are comparable so that the average sign is nearly zero. At higher temperatures the particles' statistics become less important and we would expect that the average sign is not much different from one. However, as $\beta \to \infty$, or as $N \to \infty$, permutations become more likely, and thus $\langle S \rangle_+$ tends to zero.

In fact, one can show that for fermions the average sign decays exponentially as

$$\langle S \rangle_+ \propto e^{-\beta N (f_F - f_B)}, \tag{2.22}$$

where $f_{F/B}$ is the free energy per particle of the Fermi/Bose system. Given that the statistical error decays like $N_s^{-1/2}$, we see that directly simulating interacting fermions is restricted to very small systems and relatively high temperatures.

2.1.2 Restricted Paths

One way to avoid the sign problem is to use the restricted path approximation to PIMC (RPIMC) [91]. The basic idea is to split configuration space into 'nodal cells' inside of which the fermion density matrix is always of the same sign before stitching these pockets together to form the full solution. Since the nodes of the true density matrix are unknown, one usually resorts to replacing them with the nodes of some trial density matrix, most typically with the nodes of a collection of free electrons, although success with other choices has recently been achieved [92]. The simulation proceeds as normal except now moves which cross the nodes of the trial density matrix are rejected.

RPIMC with free-particle nodes has been applied to many different problems but the restriction amounts to an uncontrolled approximation the quality of which is not well known. Indeed, the author of Ref. [93] suggests that RPIMC cannot reproduce the exact solution for a system of non-interacting electrons and that the quality of RPIMC results should degrade as the density of the system increases.

2.1.3 Other approaches

Real-space RPIMC has to date been the most successful and widely used path-integral approach for studying continuum fermion systems, but has some limitations including the uncontrolled nodal approximation. Two very different and promising approaches, namely permutation-blocking [94] and configuration PIMC [95, 96], have emerged recently which address some of these limitations.

Permutation-blocking (PB-PIMC) is essentially conventional PIMC in coordinate space but uses a high-order Trotter-decomposition, incorporates exchange directly using determinantal density matrices and uses more efficient Metropolis sampling. The combination of these developments leads to a significant improvement over conventional PIMC in treating fermions without resorting to uncontrolled approximations but, this approach is still limited by the fermion sign problem at low temperatures.

Configuration path integral Monte Carlo (CPIMC) takes quite a different route and instead uses the perturbative expansion for the canonical density operator (see, for

example, Ref. [14])

$$\hat{\rho} = e^{-\beta \hat{H}^0} \sum_{K=0}^{\infty} (-1)^K \int_0^\beta d\tau_1 \int_{\tau_1}^\beta d\tau_2 \cdots \int_{\tau_{K-1}}^\beta d\tau_K \hat{H}'_I(\tau_K) \ldots \hat{H}'_I(\tau_1), \tag{2.23}$$

where

$$\hat{O}_I(\tau) = e^{\tau \hat{H}^0} \hat{O} e^{-\tau \hat{H}^0}, \tag{2.24}$$

is the definition of an operator in the interaction picture. With this form, an analogous path integral formulation can be built up, except now the trace in the partition function is carried out using a basis of many-particle Slater determinants. This yields a continuous time algorithm free from time step errors very similar to diagrammatic Monte Carlo approaches which have been successful in treating lattice models [97–99]. Working in the space of explicitly anti-symmetric Slater determinants allows exchange processes to be incorporated naturally, which is otherwise difficult to achieve using real space methods. CPIMC is thus best suited at high densities (low r_s) where quantum degeneracy effects are most important. Much like conventional PIMC, CPIMC suffers from the sign problem at low temperatures which can be seen immediately from the factors of $(-1)^K$ appearing in Eq. (2.23). One can attempt to restrict the expansion order K (also called the number of 'kinks') to alleviate the sign problem and then try to extrapolate the residual bias [12]. However, this can only be achieved up to a point [53].

While promising, the sign problem still restricts all approaches which seek an exact solution at low temperatures, while RPIMC results in biases of unknown magnitude. Recently, a very promising approach emerged in the context of quantum chemistry called full configuration interaction QMC (FCIQMC), which allowed for the simulation of larger systems than would otherwise have been possible without introducing uncontrolled approximations. In the last sections of this chapter we will discuss how ideas from the ground state theory can be incorporated into an independent approach for simulating systems at non-zero temperatures.

2.2 Projector Methods

One way to find the ground state of the many-electron Hamiltonian is via the imaginary time Schrödinger equation

$$-\frac{d}{d\tau}|\Psi\rangle = \hat{H}|\Psi\rangle, \tag{2.25}$$

the formal solution of which is given as

$$|\Psi(\tau)\rangle = e^{-\hat{H}\tau}|\Psi(0)\rangle. \tag{2.26}$$

Writing $|\Psi(0)\rangle$ in a basis of eigenstates of $\hat{H}$ as

$$|\Psi(0)\rangle = \sum_i c_i |E_i\rangle, \tag{2.27}$$

it follows that in the limit $\tau \to \infty$,

$$|\Psi(\tau \to \infty)\rangle \approx c_0 e^{-E_0 \tau} |E_0\rangle, \tag{2.28}$$

since all higher excited states will decay exponentially faster. Splitting the exponential into P short time operators as in Section 2.1 we can write Eq. (2.26) as

$$|\Psi(\tau)\rangle = \left(e^{-\Delta\tau \hat{H}}\right)^P |\Psi(0)\rangle, \tag{2.29}$$

where $\Delta\tau = \tau/P$. Thus, by repeatedly applying the evolution operator $e^{-\Delta\tau\hat{H}}$ we can *project* out the ground state as in Eq. (2.28), provided $\langle E_0|\Psi(0)\rangle \neq 0$.

This is the basis for projector QMC methods of which many variants exist, such as Green's function [82], auxiliary field [100, 101] and diffusion QMC [83]. As in the case of PIMC, projector QMC methods also suffer from the sign problem. To overcome these issues, various uncontrolled approximations are employed to ensure the positivity of the weights. For example, the fixed node approximation is used in DMC [83] while the phaseless or constrained path approximation is used in AFQMC [100–102]. While the biases introduced using these constraints are typically thought to be small [103], the lack of a systematic way to release the constraints and converge to the correct answer can represent a barrier when very high accuracy results are sought. This is particularly a problem when multiple uncontrolled approximations are made as is the case when using pseudopotentials.

In 2009, a new projector approach, known as full configuration interaction QMC (FCIQMC) [104], was developed which can treat correlated systems larger than previously possible without uncontrolled approximations. Since we cannot evaluate the short time propagators appearing in Eq. (2.29), here we instead use the first order approximation

$$|\Psi(\tau)\rangle = (\hat{\mathbb{1}} - \Delta\tau\hat{H})^P |\Psi(0)\rangle, \tag{2.30}$$

and expand the wave function in a discrete basis of Slater determinants as

$$|\Psi(0)\rangle = \sum_{\mathbf{i}} c_{\mathbf{i}}(0) |D_{\mathbf{i}}\rangle. \tag{2.31}$$

This allows us to rewrite Eq. (2.30) as an iterative equation for the expansion coefficients, $c_{\mathbf{i}}$, as

$$c_{\mathbf{i}}(\tau + \Delta\tau) = c_{\mathbf{i}}(\tau) - \Delta\tau \sum_{\mathbf{j}} H_{\mathbf{ij}} c_{\mathbf{j}}(\tau), \tag{2.32}$$

where $H_{\mathbf{ij}} = \langle D_{\mathbf{i}}|\hat{H}|D_{\mathbf{j}}\rangle$. Solving Eq. (2.32) leaves us with the problem that the normalisation of $|E_0\rangle$ constantly changes with τ due to the final decaying exponential in Eq. (2.28). We could remove this by shifting the our energy scale so that $E_0 = 0$. However, since we do not know E_0 in advance we instead iterate

$$c_{\mathbf{i}}(\tau + \Delta\tau) = c_{\mathbf{i}}(\tau) + \Delta\tau \sum_{\mathbf{j}} T_{\mathbf{ij}} c_{\mathbf{j}}(\tau), \tag{2.33}$$

where $T_{\mathbf{ij}} = -(H_{\mathbf{ij}} - \delta_{\mathbf{ij}} S)$, and S is a variable shift which is adjusted during the simulation to control the population of walkers and approximately fix the normalisation of the wavefunction. Another concern with Eq. (2.33) is that in using the short-time propagator we would introduce a time step error. However, it turns out that this is not an issue provided $\Delta\tau < 2/(E_{\max} - E_0)$, where $E_{\max}$ is the maximal eigenvalue of $\hat{H}$ in our basis [105].

Of course, Eq. (2.33) is not very useful since we usually can store neither $|\Psi\rangle$ nor $\hat{H}$. The genius of FCIQMC is that we can instead use a Monte Carlo method to sample the application of the projector and also the wavefunction, thus removing this memory bottleneck. To see this, note that we can interpret Eq. (2.33) as a rate equation for the coefficients of the wave function expansion. We can solve this by stochastically by evolving a set of signed walkers[2] according to Eq. (2.33). A simulation normally begins by specifying some initial distribution of walkers (typically one walker is placed on the Hartree–Fock determinant) after which each walker performs the following set of moves for each time step [104]:

1. A walker $w_{\mathbf{i}}$ attempts to spawn from a determinant $|D_{\mathbf{i}}\rangle$ to another $|D_{\mathbf{j}}\rangle$ with probability $p_s(\mathbf{i} \to \mathbf{j}) = \frac{\Delta\tau}{p_{\text{gen}}(\mathbf{ij})}|T_{\mathbf{ji}}|$. The sign of the spawned walker is $\text{sign}(T_{\mathbf{ij}}) \times \text{sign}(w_{\mathbf{i}})$. To be clear, 'with probability' means that if $p_s > r$, where $r \in [0, 1)$ is a uniformly chosen random number, then we spawn a walker onto $|D_{\mathbf{i}}\rangle$. If $p_s \geq 1$ then we spawn $\lfloor p_s \rfloor$ walkers[3] with probability 1 and another walker with probability $p_s - \lfloor p_s \rfloor$.[4] We also have introduced the generation probability, $p_{\text{gen}}(\mathbf{ij})$, which is the probability of generating the move $\mathbf{i} \to \mathbf{j}$ in the first place. This accounts for the fact that it is far more efficient to only attempt one spawning move per walker rather than trying to attempt spawn to all possible connected determinants. The choice of how to generate excited determinants can greatly affect the efficiency of the algorithm and we refer the reader to Ref. [104] for more details. This spawning step simulates the terms like $c_{\mathbf{j}}(\tau + \Delta\tau) = \Delta\tau \sum_{\mathbf{j}\neq\mathbf{i}} T_{\mathbf{ji}} c_{\mathbf{i}}(\tau)$ in Eq. (2.33).
2. A walker $w_{\mathbf{i}}$ may die and disappear from the simulation or clone a new walker with

[2] Since strictly speaking the walkers do not move, a better name for the walkers might be psi-particles or 'psips' [106]. However, walker is now standard in the literature.
[3] Here the floor operation means, e.g., $\lfloor 2.3999 \rfloor = 2$
[4] We are assuming we represent the walkers using integers. One could just as well use real numbers, and this choice often leads to better statistical error bars [107, 108].

probability $p_d(\mathbf{i}) = \Delta\tau|T_{\mathbf{ii}}|$. If $\mathrm{sign}(T_{\mathbf{ii}}) < 0$ then we remove the walker from the simulation if $\mathrm{sign}(T_{\mathbf{ii}}) > 0$ we clone that walker. We account for multiple cloning events as in step 1. The death/cloning step simulates $c_{\mathbf{i}}(\tau+\Delta\tau) = (1+\Delta\tau T_{\mathbf{ii}})c_{\mathbf{i}}(\tau)$.

3. Walkers of the opposite sign on the same determinant are annihilated and removed from the simulation so that all walkers remaining on a particular determinant have the same sign.

Thus, walkers undergo birth and death processes stochastically, with low weight contributions flitting in and out of existence so that only a snapshot of the wavefunction is stored at each time step. It is only in the long-time limit that the average number of walkers appearing on a given determinant, $\bar{n}^w_{\mathbf{k}}$, is proportional to the corresponding ground-state wavefunction expansion coefficient $c^{\mathrm{gs}}_{\mathbf{k}}$.

During a simulation, a very distinctive walker dynamics emerges, as shown in the top panel of Fig. 2.1. First, there is an initial exponential increase in the walker population followed by a plateau region where the population stabilises. After a certain amount of time the population will begin to increase exponentially again, although at a slower rate. It is after the plateau stage that the sign structure of the ground state wavefunction emerges and correct results for the energy can be found (see the bottom panel of Fig. 2.1). To combat this further rise in population, which only serves to improve statistics, Booth *et al.* introduced a variable shift, adapted from DMC [109], which is designed to dampen this population rise as

$$S(\tau + \Delta\tau) = S(\tau) + \frac{\zeta}{A\Delta\tau}\frac{N_w(\tau + A\Delta\tau)}{N_w(\tau)}, \tag{2.34}$$

where ζ is a damping parameter, A is the number of time steps between which the shift is held constant, and $N_w(\tau)$ is the total number of walkers at imaginary time τ. Typical values of ζ and A are 0.1 and 20 respectively. The form of Eq. (2.34) is chosen so that the shift approaches E_0 once the ground state has been projected out. Both the shift Eq. (2.34) and the projected energy[5],

$$E_{\mathrm{proj}}(\tau) \equiv \frac{\langle D_{\mathbf{0}}|\hat{H}|\Psi(\tau)\rangle}{\langle D_{\mathbf{0}}|\Psi(\tau)\rangle} \tag{2.35}$$

$$= \frac{\sum_{\mathbf{k}} H_{\mathbf{0k}} n^w_{\mathbf{k}}(\tau)}{n^w_{\mathbf{0}}(\tau)}, \tag{2.36}$$

can then be used as estimators for the energy although care must be taken to account for serial correlations between neighbouring data points [110]. Here $n^w_{\mathbf{i}}(\tau)$ is the total number of walkers on $|D_{\mathbf{i}}\rangle$ at time τ. The behaviour of both of these estimators is also shown in Fig. 2.1.

[5] If $\lim_{\tau\to\infty}|\Psi(\tau)\rangle = |E_0\rangle = \sum_{\mathbf{i}} c^{\mathrm{gs}}_{\mathbf{i}}|D_{\mathbf{i}}\rangle$, then we have $\lim_{\tau\to\infty} E_{\mathrm{proj}}(\tau) = \frac{E_0 c^{\mathrm{gs}}_{\mathbf{0}}}{c^{\mathrm{gs}}_{\mathbf{0}}}$.

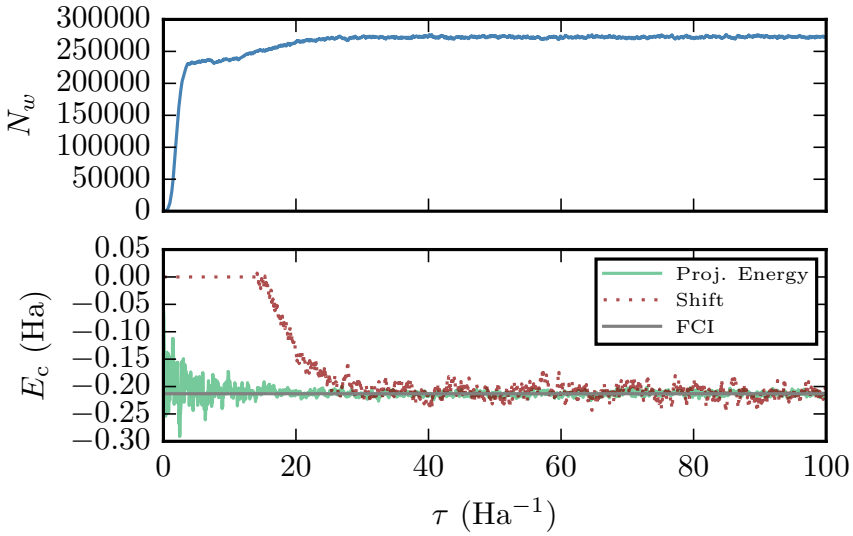

Figure 2.1: Example of the observed population dynamics in an FCIQMC calculation for Ne in an aug-cc-pVDZ basis with 8 electrons correlated in 22 spin orbitals. Note the plateau emerging for $N_w \approx 2.3 \times 10^5$ between $\tau \approx 2 - 10$ Ha^{-1}. Once we exit the plateau phase the shift is allowed to vary and we see, from the bottom panel, that both the shift estimator, and the projected energy estimator converge to the exact FCI correlation energy [111]. This is a reproduction of a similar calculation carried out in Ref. [104], here using the HANDE QMC package [112, 113].

2.2.1 The Sign Problem

FCIQMC naturally suffers from the sign problem since walkers of both positive and negative weight can be spawned. In contrast to PIMC methods, in FCIQMC the severity of the sign problem is measured by the critical number of walkers required to reach the plateau phase of the simulation [104]. To understand the link between the walker dynamics and the sign problem, we follow Spencer *et al.* [114] and write the transition matrix $\mathbf{T} = \mathbf{T}^+ - \mathbf{T}^-$, where $\mathbf{T}^+$ and $\mathbf{T}^-$ contain the absolute values of the positive and negative elements of $\mathbf{T}$ respectively. The decoupled equations of motion for the positive ($n_\mathbf{i}^+$) and negative ($n_\mathbf{i}^-$) walkers without annihilation can then be shown to read

$$\frac{d(n_\mathbf{i}^+ + n_\mathbf{i}^-)}{d\tau} = \sum_\mathbf{j} (T_\mathbf{ij}^+ + T_\mathbf{ij}^-)(n_\mathbf{j}^+ + n_\mathbf{j}^-), \tag{2.37}$$

$$\frac{d(n_\mathbf{i}^+ - n_\mathbf{i}^-)}{d\tau} = \sum_\mathbf{j} (T_\mathbf{ij}^+ - T_\mathbf{ij}^-)(n_\mathbf{j}^+ - n_\mathbf{j}^-). \tag{2.38}$$

One can show that the maximum eigenvalue of $\mathbf{T}^+ + \mathbf{T}^-$ is always greater than the maximum eigenvalue value of $\mathbf{T}^+ - \mathbf{T}^-$ [114]. Thus we see that, in the long time limit, the first solution dominates over the desired second solution, leading to an exponential increase in noise. The inclusion of annihilation serves to dampen the unphysical first

solution without affecting the desired second solution. Thus, once a critical concentration of walkers is reached, the physical solution will win out. This crucial annihilation process is enhanced in a discrete Hilbert space [115] as walkers are more likely to meet and cancel than in the continuous real space. Provided the critical population is small enough to allow for storage on a computer, the FCIQMC algorithm finds the ground state without requiring any uncontrolled approximations.

2.3 Density Matrix Quantum Monte Carlo

In the last section we saw how working in a discrete Hilbert space allows for a more efficient treatment of the sign problem than is usually possible with real space methods. FCIQMC has since been applied to a wide variety of ground state problems in quantum chemistry [116, 117] and physics [118–120] and the ideas have even been adapted to alternative quantum chemistry methods [121–124]. Significant progress has also been made in terms of algorithmic development, such as the use of semi-stochastic projection [107, 108] and more optimal selection probabilities [125], improving the parallel efficiency [126][6] and also understanding the efficiency of the algorithm itself [127, 128]. Alongside these developments, breakthroughs have been made in calculating properties other than the total energy, which has historically been a drawback of many QMC methods. Examples of these include the sampling of reduced density matrices [129], the calculation of unbiased nuclear forces [130] and also spectral properties [131, 132]. The problem of carrying these ideas over to $T > 0$ was addressed by Blunt and coworkers with the development of the density matrix QMC (DMQMC) method [133][7].

To begin, note that the unnormalized density matrix, $\hat{\rho} = e^{-\beta \hat{H}}$, obeys the symmetrized Bloch equation[8]

$$\frac{d\hat{\rho}}{d\beta} = -\frac{1}{2}\left(\hat{H}\hat{\rho} + \hat{\rho}\hat{H}\right), \tag{2.39}$$

which can be solved using a simple Euler update scheme:

$$\hat{\rho}(\beta + \Delta\beta) = \hat{\rho}(\beta) - \frac{\Delta\beta}{2}\left(\hat{H}\hat{\rho}(\beta) + \hat{\rho}(\beta)\hat{H}\right) + \mathcal{O}(\Delta\beta^2). \tag{2.40}$$

As in FCIQMC we can attempt to solve Eq. (2.40) stochastically by evolving a population of signed walkers, where now the walkers live in a discrete operator space made

[6] For more information on the parallel implementation of FCIQMC and some potential optimisations see Appendix A.

[7] We note an alternative way of calculating thermal properties from FCIQMC was presented in Ref. [131].

[8] Here we symmetrise the Bloch equation in anticipation of arriving at a FCIQMC-like algorithm. The advantage of this symmetric form is that walkers will be allowed to spawn along both rows and columns of the density matrix. This allows for better sampling at lower temperature. See Ref. [133] for further discussion. The density matrix trivially obeys the Bloch equation $\frac{d}{d\beta}e^{-\beta\hat{H}} = \hat{H}e^{-\beta\hat{H}}$; to symmetrise the equation we just use the fact that $\hat{H}$ commutes with $\hat{\rho}$.

up of outer products of Slater determinants $|D_\mathbf{i}\rangle\langle D_\mathbf{j}|$. The density matrix can then be written as

$$\hat{\rho}(\beta) = \sum_{\mathbf{ij}} \rho_{\mathbf{ij}}(\beta)|D_\mathbf{i}\rangle\langle D_\mathbf{j}|, \tag{2.41}$$

and the walkers sample the expansion coefficients $\rho_{\mathbf{ij}}(\beta)$. The iterative equation for the expansion coefficients is then

$$\begin{aligned}\rho_{\mathbf{ij}}(\beta+\Delta\beta) &= \rho_{\mathbf{ij}}(\beta) - \frac{\Delta\beta}{2}\sum_{\mathbf{k}} [(H_{\mathbf{ik}} - S\delta_{\mathbf{ik}})\rho_{\mathbf{kj}} - \\ &\qquad \rho_{\mathbf{ik}}(H_{\mathbf{kj}} - S\delta_{\mathbf{kj}})] \qquad (2.42)\\ &= \rho_{\mathbf{ij}}(\beta) + \frac{\Delta\beta}{2}\sum_{\mathbf{k}}(T_{\mathbf{ik}}\rho_{\mathbf{kj}} + \rho_{\mathbf{ik}}T_{\mathbf{kj}}), \qquad (2.43)\end{aligned}$$

where again we have introduced the variable shift S to control the total population of walkers. The rules for evolving the walkers are entirely analogous to those from FCIQMC and follow from Eq. (2.43):

1. A walker can spawn from a density matrix element $\rho_{\mathbf{ik}}$ to $\rho_{\mathbf{ij}}$ with probability $p_s(\mathbf{ik} \to \mathbf{ij}) = \frac{\Delta\beta|T_{\mathbf{kj}}|}{2p_{\text{gen}}(\mathbf{kj})}$, with $\text{sign}(\rho_{\mathbf{ij}}) = \text{sign}(\rho_{\mathbf{ik}}) \times \text{sign}(T_{\mathbf{kj}})$; a similar spawning process takes place from $\rho_{\mathbf{kj}}$ to $\rho_{\mathbf{ij}}$. The definitions of p_{gen} and excitaition generation is exactly the same as in the case of FCIQMC.
2. A walker on the density matrix element $\rho_{\mathbf{ij}}$ clones or dies with probability $p_d(\mathbf{ij}) = \frac{\Delta\beta}{2}|T_{\mathbf{ii}} + T_{\mathbf{jj}}|$. Cloning occurs if $\text{sign}(T_{\mathbf{ii}} + T_{\mathbf{jj}}) \times \text{sign}(\rho_{\mathbf{ij}}) > 0$ and the walker is removed from the simulation otherwise.
3. Walkers of opposite sign on the same density matrix element are annihilated and removed from the simulation.

Unlike FCIQMC where we are only interested in the long time solution, here the intermediate dynamics of the walkers is important to get the correct physical results.

The Bloch equation is a first order linear differential equation whose solution is uniquely defined by the initial condition. The simplest place to begin a simulation is at $\beta = 0$, where the density matrix is the identity matrix and can be sampled by occupying diagonal density matrix elements with uniform probability. A simulation then consists of propagating the initial distribution of walkers with the rules described above to a desired value of β. Estimates for thermodynamic quantities can be found by averaging over many such simulations. Some care needs to be taken when evaluating expectation values as we need to estimate the average of a ratio of two quantities. One way to do this is to average the numerator and denominator separately before taking the ratio:

$$\langle \hat{O} \rangle \approx \frac{\frac{1}{N_s}\sum_s \sum_{\mathbf{ij}} n_{\mathbf{ij}}^{(s)} O_{\mathbf{ji}}}{\frac{1}{N_s}\sum_s \sum_{\mathbf{i}} n_{\mathbf{ii}}^{(s)}}, \tag{2.44}$$

and here $n_{\mathbf{ij}}^{(s)}$ is the number of walkers on density matrix element $|D_\mathbf{i}\rangle\langle D_\mathbf{j}|$ at a specific

temperature β in simulation s from a total of N_s. This estimator still suffers from a population bias although it is typically small and vanishes like N_s^{-1}. As the walker distributions are completely independent of each other from simulation to simulation the variance is simply given by standard error propagation. Writing

$$\bar{O} = \frac{\bar{A}}{\bar{B}} = \frac{\frac{1}{N_s}\sum_s A^{(s)}}{\frac{1}{N_s}\sum_s B^{(s)}} \tag{2.45}$$

and with $\sigma_{\bar{A}}$ calculated as in Eq. (2.3), then we can calculate the standard error in $\bar{O}$ using

$$\sigma_{\bar{O}} \approx \bar{O}\sqrt{\left(\left(\frac{\sigma_{\bar{A}}}{\bar{A}}\right)^2 + \left(\frac{\sigma_{\bar{B}}}{\bar{B}}\right)^2 - 2\frac{\mathrm{Cov}(A,B)}{N_s\bar{A}\bar{B}}\right)} \tag{2.46}$$

where the sample covariance is given as

$$\mathrm{Cov}(A,B) = \frac{1}{(N_s-1)}\sum_s (A^{(s)} - \bar{A})(B^{(s)} - \bar{B}). \tag{2.47}$$

An issue DMQMC runs into in large Hilbert spaces is that (for of a fixed population of walkers) the number of walkers on the diagonal of the density matrix can decay to zero rendering Eq. (2.44) a poor estimator. This is perhaps unsurprising given that the number of off-diagonal elements relative to the number of diagonal elements increases exponentially with system size. Additionally, since most of the weight contributing to the evaluation of a given estimator is determined by walkers on the diagonal, or at most a few excitation levels[9] away, a method to improve the sampling of these important parts of Hilbert space seems like a good idea. Problems such as these are usually overcome using importance sampling [82], which was adapted to DMQMC in Ref. [133]. To begin we write the importance sampled density matrix as

$$\tilde{\rho}_{\mathbf{ij}} = \rho^T_{\mathbf{ij}}\rho_{\mathbf{ij}}, \tag{2.48}$$

where element-wise multiplication is assumed and ρ^T is some trial density matrix. Inserting Eq. (2.48) into Eq. (2.39) we find

$$\frac{d\tilde{\rho}_{\mathbf{ij}}}{d\beta} = \frac{d\rho^T_{\mathbf{ij}}}{d\beta}\rho_{\mathbf{ij}} + \rho^T_{\mathbf{ij}}\frac{d\rho_{\mathbf{ij}}}{d\beta} \tag{2.49}$$

$$= \frac{d\rho^T_{\mathbf{ij}}}{d\beta}\rho_{\mathbf{ij}} + \frac{1}{2}\rho^T_{\mathbf{ij}}\sum_{\mathbf{k}}\left(T_{\mathbf{ik}}\rho_{\mathbf{kj}} + \rho_{\mathbf{ik}}T_{\mathbf{kj}}\right) \tag{2.50}$$

$$= \frac{d\log\rho^T_{\mathbf{ij}}}{d\beta}\tilde{\rho}_{\mathbf{ij}} + \frac{1}{2}\sum_{\mathbf{k}}\left(\left(\rho^T_{\mathbf{ij}}T_{\mathbf{ik}}\frac{1}{\rho^T_{\mathbf{kj}}}\right)\left(\rho^T_{\mathbf{kj}}\rho_{\mathbf{kj}}\right) + \left(\rho^T_{\mathbf{ik}}\rho_{\mathbf{ik}}\right)\left(\rho^T_{\mathbf{ij}}T_{\mathbf{kj}}\frac{1}{\rho^T_{\mathbf{ik}}}\right)\right) \tag{2.51}$$

which is nearly identical to the usual equation we solve with DMQMC except now

[9] The excitation level (n_{ex}) is defined such that the bra and ket of a density matrix element differ by n_{ex} particle-hole pairs.

the transition rates between states are modified by the importance function. Thus, by choosing ρ^T appropriately, we can encourage walkers to stay near the diagonal (or wherever we want them to be). These re-weighting factors can be determined dynamically through the simulation and are usually arbitrarily chosen so that the weight of walkers on each excitation level is roughly uniform [133]. In larger systems the weights necessary for constructing ρ^T can become quite large. This can restrict spawning and introduce potential population biases if the weights are used from the beginning of the simulation [133]. To overcome this, the weights can instead be gradually introduced as β increases. The first term on the right hand side of Eq. (2.51) accounts for this temperature dependence.

Finally, given that the space we sample using DMQMC is the square of that in FCIQMC one might think that the number of walkers required for the ground-state to emerge may be square of that required for an FCIQMC simulation of the same system. This was indeed found to be the case in the first application of DMQMC to the strongly-correlated Heisenberg model [133], but it remained to be seen how bad this would be for other systems.

2.4 Summary

In this chapter we introduced some quantum Monte Carlo methods which could in principle be applied to studying warm dense matter. The primary limitation of all of these methods is the sign problem, which arises due to Fermionic statistics. DMQMC is a promising new method that should be able to build on the successes of FCIQMC in treating systems with sign problems in an unbiased fashion. In the next chapter we will see to what extent this hope is fulfilled.

Chapter 3

The Uniform Electron Gas

Solving the many-electron Hamiltonian (Eq. (1.3)) directly using any of the methods described so far is complicated significantly by the presence of the nuclear potential felt by the electrons. A simpler approach to understanding the importance of electron-electron interactions is to first neglect the presence of the nuclei and smear out the charge so that the electrons only interact with themselves and a positive uniform background. In this way one arrives at the uniform electron gas (UEG) model[1] the first and simplest attempts to describe a realistic metallic system. Despite the simplification of homogeneity, the UEG has been of fundamental importance in the development of our modern understanding of correlated electron systems. Concepts such as Fermi liquid theory [32], collective excitations and quasiparticles [134, 135], the BCS theory of superconductivity [136], screening [28], and Hohenberg-Kohn-Sham DFT [39, 40], were all built on our understanding of the low temperature properties of the UEG.

The UEG is also of central importance for the quantitative accuracy of DFT. Most DFT exchange-correlation functionals are built, in some way, on the local density approximation (see Section 1.6), which relies on an accurate fit of the exchange-correlation energy of the UEG as a function of density [137–140]. As DFT results are sensitive to the approximate functional used, these fits are in turn based on Ceperley and Alder's highly accurate QMC results for the UEG [141]. For thermal DFT we likewise need an accurate fit for the exchange-correlation free energy of the UEG across the whole density-temperature plane, which should again be based on accurate QMC results.

The first step towards providing this much needed data was taken by Brown *et al.*, who they applied restricted path-integral Monte Carlo (RPIMC) to the warm dense UEG [11]. Unfortunately, subsequent configuration path-integral Monte Carlo (CPIMC) simulations disagreed substantially with the RPIMC results in the high density regime [12]. As these QMC data are already being incorporated into exchange-correlation free energy functionals [50, 51, 142] it is important to have an alternative method capable of resolving this disagreement.

In this chapter we will outline how we can adapt DMQMC to treat the warm dense UEG and discuss some of the complications which arise in the process. In overcoming

[1] Also known as the homogeneous electron gas or one-component plasma.

these difficulties we present an alternative formalism analogous to working in the interaction picture (see Eq. (2.24)). This interaction-picture DMQMC algorithm enables us to perform some initial proof-of-principle benchmark studies against CPIMC and also provides a direct route to evaluating the free energy.

3.1 The Hamiltonian

The UEG describes a system N electrons in a cubic box of volume $V = L^3$ where the charge distribution of the nuclei is now uniformly distributed throughout space as $n(\mathbf{r}) = n = N/V$ so that the whole system is charge neutral. Writing down the many-electron Hamiltonian as in Eq. (1.28) is complicated by the long-range Coulomb interaction. Here we will first discuss the thermodynamic limit result which is a bit simpler [14], before discussing the problem of simulating a finite number of electrons in Section 3.1.1.

Let us first assume we are in a very large box so that the error resulting from replacing summations with integrals is small, and we can further extend the domain of integration to all space. One finds immediately that the electron-background, electron-electron and background-background contributions all individually diverge. Since the system is physical, all of these divergences must cancel and it is usual to first regularise the Coulomb interaction by including factors of $e^{-\lambda r}$ so that intermediate results remain finite. The limit of $\lambda \to 0$ will then be taken at the end of the calculation.

With these points in mind we can write

$$\hat{V}_{en} \to \hat{V}_{\mathrm{e-bg}} = -\sum_i \int d\mathbf{R} \frac{n e^{-\lambda|\mathbf{r}_i - \mathbf{R}|}}{|\mathbf{r}_i - \mathbf{R}|} \tag{3.1}$$

$$= -Nn\frac{4\pi}{\lambda^2}, \tag{3.2}$$

and

$$\hat{V}_{nn} \to \hat{V}_{\mathrm{bg-bg}} = \frac{1}{2}\int d\mathbf{R}' \int d\mathbf{R} \frac{n^2 e^{-\lambda|\mathbf{R}' - \mathbf{R}|}}{|\mathbf{R}' - \mathbf{R}|} \tag{3.3}$$

$$= \frac{1}{2} Nn\frac{4\pi}{\lambda^2}. \tag{3.4}$$

All that now remains of the many-body Hamiltonian are terms including only the electrons. To make further progress we need to evaluate the matrix elements of the kinetic and electron-electron interaction. The most natural single-particle basis we can use for the UEG are plane waves, where we have

$$\varphi_i(\mathbf{r}) \to \langle \mathbf{r}|\mathbf{k}_i\sigma_i\rangle = \varphi_{\mathbf{k}_i\sigma_i}(\mathbf{r}) = \frac{1}{\sqrt{V}} e^{i\mathbf{k}_i\cdot\mathbf{r}}\chi_{\sigma_i}, \tag{3.5}$$

with $\mathbf{k} \in \left\{\frac{2\pi}{L}\mathbf{n}\right\}$ where $\mathbf{n}$ is a vector of integers and the spin functions are given as

$$\chi_\uparrow = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \tag{3.6}$$

$$\chi_\downarrow = \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \tag{3.7}$$

The matrix elements necessary to construct the Hamiltonian can now be calculated as (see Ref. [14] for the details)

$$\begin{aligned} t_{ij} &\to \langle \mathbf{k}_i \sigma_i | -\frac{1}{2}\nabla^2 | \mathbf{k}_j \sigma_j \rangle && (3.8)\\ &= \frac{1}{V}\delta_{\sigma_i\sigma_j} \int d\mathbf{r} e^{-i\mathbf{k}_i\cdot\mathbf{r}} \left(-\frac{1}{2}\nabla^2 e^{i\mathbf{k}_j\cdot\mathbf{r}}\right) && (3.9)\\ &= \frac{1}{2}\mathbf{k}_j^2 \delta_{\sigma_i\sigma_j}\delta_{\mathbf{k}_i\mathbf{k}_j}, && (3.10) \end{aligned}$$

and

$$\begin{aligned} v_{ijkl} &\to \langle \mathbf{k}_i\sigma_i \mathbf{k}_j\sigma_j | \frac{e^{-\lambda|\mathbf{r}-\mathbf{r}'|}}{|\mathbf{r}-\mathbf{r}'|} | \mathbf{k}_k\sigma_k \mathbf{k}_l\sigma_l \rangle && (3.11)\\ &= \frac{1}{V^2}\delta_{\sigma_i\sigma_k}\delta_{\sigma_j\sigma_l} \int d\mathbf{r} \int d\mathbf{r}' e^{-i\mathbf{k}_i\cdot\mathbf{r}} e^{-\mathbf{k}_j\cdot\mathbf{r}'} \frac{e^{-\lambda|\mathbf{r}-\mathbf{r}'|}}{|\mathbf{r}-\mathbf{r}'|} e^{i\mathbf{k}_k\cdot\mathbf{r}} e^{i\mathbf{k}_l\cdot\mathbf{r}'} && (3.12)\\ &= \frac{1}{V}\frac{4\pi}{(\mathbf{k}_i-\mathbf{k}_k)^2+\lambda^2}\delta_{\sigma_i\sigma_k}\delta_{\sigma_j\sigma_l}\delta_{\mathbf{k}_i+\mathbf{k}_j,\mathbf{k}_k+\mathbf{k}_l}. && (3.13) \end{aligned}$$

We can use the momentum conservation Kronecker delta above to simplify the potential term by writing $\mathbf{k}_i - \mathbf{k}_k = \mathbf{q} = \mathbf{k}_l - \mathbf{k}_j$. Collecting all the terms together we have

$$\begin{aligned} \hat{H} = &\sum_{\mathbf{k}\sigma} \frac{1}{2}\mathbf{k}^2 \hat{c}^\dagger_{\mathbf{k}\sigma}\hat{c}_{\mathbf{k}\sigma} + \frac{1}{2V}\sum_{\mathbf{q}\neq\mathbf{0}}\sum_{\mathbf{k}\mathbf{p}\sigma\sigma'} \frac{4\pi}{\mathbf{q}^2+\lambda^2} \hat{c}^\dagger_{\mathbf{k}+\mathbf{q}\sigma}\hat{c}^\dagger_{\mathbf{p}-\mathbf{q}\sigma'}\hat{c}_{\mathbf{p}\sigma'}\hat{c}_{\mathbf{k}\sigma} \\ &+ \frac{1}{2V}\frac{4\pi}{\lambda^2}(\hat{N}^2-\hat{N}) + \frac{1}{2}Nn\frac{4\pi}{\lambda^2} - Nn\frac{4\pi}{\lambda^2}, \end{aligned} \tag{3.14}$$

where we have separated the $\mathbf{q} = \mathbf{0}$ contribution from the electron-electron interaction and, to tidy things up, we dropped the extra Latin index of the k-points. The remaining divergent terms in Eq. (3.14) vanish in the thermodynamic limit since

$$\begin{aligned} &\lim_{N,V\to\infty} N^{-1}\left\langle \frac{1}{2V}\frac{4\pi}{\lambda^2}(\hat{N}^2-\hat{N}) + \frac{1}{2}Nn\frac{4\pi}{\lambda^2} - Nn\frac{4\pi}{\lambda^2}\right\rangle && (3.15)\\ &\quad = \lim_{N,V\to\infty} -\frac{1}{2}\frac{4\pi}{\lambda^2}\frac{n}{N} = 0. && (3.16) \end{aligned}$$

Only now can we safely set $\lambda = 0$ and hence

$$\hat{H} = \sum_{\mathbf{k}\sigma}\frac{1}{2}\mathbf{k}^2\hat{c}^\dagger_{\mathbf{k}\sigma}\hat{c}_{\mathbf{k}\sigma} + \frac{1}{2V}\sum_{\mathbf{q}\neq\mathbf{0}}\sum_{\mathbf{k}\mathbf{p}\sigma\sigma'} v_{\mathbf{q}}\hat{c}^\dagger_{\mathbf{k}+\mathbf{q}\sigma}\hat{c}^\dagger_{\mathbf{p}-\mathbf{q}\sigma'}\hat{c}_{\mathbf{p}\sigma'}\hat{c}_{\mathbf{k}\sigma}, \tag{3.17}$$

where $v_{\mathbf{q}} = \frac{4\pi}{\mathbf{q}^2}$.

Inspecting the powers of L^{-1} appearing in Eq. (3.17) we see that the kinetic energy goes like r_s^{-2}, while the potential term goes like r_s^{-1} (see Eq. (1.55)). Thus, at $T = 0$, and in the thermodynamic limit[2], the only free parameter for the UEG is r_s, which measures the relative strength of these two interactions. For $r_s \ll 1$ (high density) the kinetic term dominates and the system behaves nearly like an ideal gas. For $r_s \gg 1$ (low density) the potential term dominates and it is thought the system should form a Wigner crystal [28]. Most metals lie in between these regimes, and since there is no small coupling parameter, we usually need to resort to accurate quantum Monte Carlo simulations for quantitatively accurate results.

3.1.1 Finite System Sizes

Of course, most QMC methods work far away from the thermodynamic limit, only explicitly simulating a handful of electrons. This complicates the discussion of the electron-electron interaction considerably[3]. From a classical perspective, the total Coulombic energy in a periodic box is

$$V_{ee} = \frac{1}{2} \sum_{i,j \neq i} \sum_{\mathbf{R}} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j - \mathbf{R}|}, \tag{3.18}$$

where $\mathbf{R} = \sum_i n_i \mathbf{a}_i$ are a set of translation vectors, and $\mathbf{a}_i$ are the primitive lattice vectors of the (here) cubic system. Unfortunately, Eq. (3.19) is an ill-defined sum which converges only conditionally. The conventional solution to this problem is to instead use the Ewald summation technique which leads to a well defined result [143, 144].

We first note that another way to calculate the electrostatic energy is from

$$V_{ee} = \frac{1}{2} \sum_i \bar{\phi}(\mathbf{r}_i), \tag{3.19}$$

where

$$\bar{\phi}(\mathbf{r}_i) = \lim_{\mathbf{r} \to \mathbf{r}_i} \left(\phi(\mathbf{r}) - \frac{1}{|\mathbf{r} - \mathbf{r}_i|} \right), \tag{3.20}$$

is the electrostatic potential at $\mathbf{r}_i$ due to all the other electrons except itself. $\phi(\mathbf{r})$ is the full electrostatic potential due to all the electrons, including $\mathbf{r}_i$, and can be found solving by Poisson's equation:

$$\nabla^2 \phi(\mathbf{r}) = -4\pi \rho(\mathbf{r}). \tag{3.21}$$

[2] For a finite system we usually specify r_s and N which then determines the box size L. At $T > 0$ we also require $\Theta = T/T_F$, where the warm dense regime corresponds to $r_s \approx \Theta \approx 1$.

[3] The kinetic energy remains unaffected since we have the identity $\delta_{\mathbf{k}\mathbf{k}'} = \frac{1}{V} \int_V d\mathbf{r} e^{i(\mathbf{k}-\mathbf{k}')\cdot\mathbf{r}}$ for a periodic system.

For a collection of periodically repeated point-charges in a cancelling uniform background we have

$$\rho(\mathbf{r}) = \sum_{\mathbf{R}} \delta(\mathbf{r} - \mathbf{r}_i - \mathbf{R}) - \rho_{\text{bg}}, \tag{3.22}$$

where the background density, ρ_{bg}, is chosen so that the cell is charge neutral. The idea behind Ewald's method is to first add and subtract a periodic array of Gaussians to Eq. (3.22) so that we can write

$$\rho(\mathbf{r}) = \rho_1(\mathbf{r}) + \rho_2(\mathbf{r}), \tag{3.23}$$

where

$$\rho_1(\mathbf{r}) = \sum_{\mathbf{R}} \left(\delta(\mathbf{r} - \mathbf{r}_i - \mathbf{R}) - \frac{\kappa}{\sqrt{\pi}} e^{-\kappa^2 |\mathbf{r} - \mathbf{r}_i - \mathbf{R}|^2} \right), \tag{3.24}$$

and

$$\rho_2(\mathbf{r}) = \frac{\kappa}{\sqrt{\pi}} \sum_{\mathbf{R}} e^{-\kappa^2 |\mathbf{r} - \mathbf{r}_i - \mathbf{R}|^2} - \rho_{\text{bg}}. \tag{3.25}$$

We can then solve Poisson's equation twice, once with $\rho_1(\mathbf{r})$ and once with $\rho_2(\mathbf{r})$, before adding the two potentials together at the end. Solving for the potential generated by $\rho_1(\mathbf{r})$ is most easily done in real space, where we can use the fact that the potential associated with a Gaussian charge distribution is related to the error function. The potential associated with $\rho_2(\mathbf{r})$ is most easily solved for in reciprocal space as the Fourier transform of a Gaussian is again a Gaussian.

Solving these equations and adding the two solutions together we can now express $\phi(\mathbf{r})$ in terms of two rapidly converging sums:

$$\begin{aligned} \phi(\mathbf{r} - \mathbf{r}') = & \frac{4\pi}{V} \sum_{\mathbf{k} \neq \mathbf{0}} \frac{\exp\left(-\mathbf{k}^2/4\kappa^2 - i\mathbf{k} \cdot (\mathbf{r} - \mathbf{r}')\right)}{\mathbf{k}^2} \\ & + \sum_{\mathbf{R}} \frac{\operatorname{erfc}(\kappa |\mathbf{r} - \mathbf{r}' - \mathbf{R}|)}{|\mathbf{r} - \mathbf{r}' - \mathbf{R}|} - \frac{\pi}{\kappa^2 V}, \end{aligned} \tag{3.26}$$

where $\mathbf{k}$ is a reciprocal lattice vector. The value of κ is arbitrary but can be chosen to speed up the convergence of the real or reciprocal space sum [145].

The total electrostatic energy of the electrons can then be written as

$$V_{ee} = \frac{1}{2} \sum_{i=1}^{N} \sum_{\substack{j=1 \\ j \neq i}}^{N} \phi(\mathbf{r}_i - \mathbf{r}_j) + \frac{N\xi}{2}, \tag{3.27}$$

where

$$\xi = \lim_{\mathbf{r}_i \to \mathbf{r}_j} \left(\phi(\mathbf{r}_i - \mathbf{r}_j) - \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} \right) \tag{3.28}$$

$$= \frac{4\pi}{V}\sum_{\mathbf{k}\neq\mathbf{0}}\frac{\exp\left(-\mathbf{k}^2/4\kappa^2\right)}{\mathbf{k}^2} - \frac{\pi}{\kappa^2 V} + \sum_{\mathbf{R}}\frac{\mathrm{erfc}(\kappa|\mathbf{R}|)}{|\mathbf{R}|} - \frac{2\kappa}{\sqrt{\pi}}, \tag{3.29}$$

accounts for the interaction of a point charge at $\mathbf{r}_i$ with all its periodic images.

We can now calculate the matrix elements necessary for Eq. (3.17). Replacing $\frac{1}{|\mathbf{r}_i-\mathbf{r}_j|}$ by $\phi(\mathbf{r}_i-\mathbf{r}_j)$ in Eq. (3.13) and ignoring spin indices we have

$$\langle \mathbf{k}_i\mathbf{k}_j|\phi(\mathbf{r}-\mathbf{r}')|\mathbf{k}_k\mathbf{k}_l\rangle = \delta_{\mathbf{k}_i+\mathbf{k}_j,\mathbf{k}_k+\mathbf{k}_l}\frac{1}{V}\int_V d\mathbf{x}\phi(\mathbf{x})e^{i(\mathbf{k}_i-\mathbf{k}_k)\cdot\mathbf{x}} \tag{3.30}$$

$$= -\delta_{\mathbf{k}_i+\mathbf{k}_j,\mathbf{k}_k+\mathbf{k}_l}\frac{1}{V}\int_V d\mathbf{x}\nabla^2\phi(\mathbf{x})\frac{1}{(\mathbf{k}_i-\mathbf{k}_k)^2}e^{i(\mathbf{k}_i-\mathbf{k}_k)\cdot\mathbf{x}} \tag{3.31}$$

$$= \delta_{\mathbf{k}_i+\mathbf{k}_j,\mathbf{k}_k+\mathbf{k}_l}\frac{1}{V}\frac{4\pi}{(\mathbf{k}_i-\mathbf{k}_k)^2}\left(1-\rho_{\mathrm{bg}}\delta_{\mathbf{k}_i-\mathbf{k}_k,\mathbf{0}}\right). \tag{3.32}$$

In going from the first to the second line we integrated by parts twice; the boundary terms vanish because $\phi(\mathbf{r})$ and $e^{i\mathbf{k}\cdot\mathbf{r}}$ are periodic functions. The Fourier components necessary for Eq. (3.17) are then

$$v_{\mathbf{q}} = \begin{cases} \frac{4\pi}{\mathbf{q}^2} & \text{if } \mathbf{q}\neq\mathbf{0} \\ 0 & \text{if } \mathbf{q}=\mathbf{0} \end{cases}, \tag{3.33}$$

where the $\mathbf{q}=\mathbf{0}$ term can be set to zero because we chose ρ_{bg} so that $\phi(\mathbf{r})$ averages to zero over our simulation cell (the average of a function over a volume corresponds to the $\mathbf{k}=\mathbf{0}$ Fourier component of that function). Thus we find, practically speaking, the only modification to the electron-electron term in Eq. (3.17) is that the constant Madelung contribution of $\frac{1}{2}N\xi$ needs to be added.

3.2 Applying DMQMC

In principle, all that is required to simulate a model Hamiltonian using DMQMC is a knowledge of the Hamiltonian matrix elements and some way to sample the initial condition[4]. If we define $|D^{\mathbf{k}_a}_{\mathbf{k}_i}\rangle = \hat{c}^\dagger_{\mathbf{k}_a}\hat{c}_{\mathbf{k}_i}|D\rangle$, and $|D^{\mathbf{k}_a\mathbf{k}_b}_{\mathbf{k}_i\mathbf{k}_j}\rangle = \hat{c}^\dagger_{\mathbf{k}_a}\hat{c}^\dagger_{\mathbf{k}_b}\hat{c}_{\mathbf{k}_j}\hat{c}_{\mathbf{k}_i}|D\rangle$, and to simplify the notation we have written $\mathbf{k}_i=(\mathbf{k}_i\sigma_i)$, then from Section 1.3 we have

$$\langle D|\hat{H}|D\rangle = \sum_{\mathbf{k}_i}\frac{1}{2}\mathbf{k}_i^2 - \frac{1}{2V}\sum_{\mathbf{k}_i}\sum_{\mathbf{k}_j\neq\mathbf{k}_i}\frac{4\pi}{(\mathbf{k}_i-\mathbf{k}_j)^2}\delta_{\sigma_i\sigma_j}, \tag{3.34}$$

$$\langle D|\hat{H}|D^{\mathbf{k}_a}_{\mathbf{k}_i}\rangle = 0, \tag{3.35}$$

[4] The implementation of the Slater-Condon rules and excitation generators for the UEG was done originally in HANDE [112, 113] by Dr. James Spencer. The original DMQMC algorithm was implemented by Dr. Nick Blunt and Tom Rogers [133].

and

$$\langle D|\hat{H}|D^{\mathbf{k}_a\mathbf{k}_b}_{\mathbf{k}_i\mathbf{k}_j}\rangle = \begin{cases} \frac{4\pi}{V}\left(\frac{1}{(\mathbf{k}_i-\mathbf{k}_a)^2}\delta_{\sigma_i\sigma_a}\delta_{\sigma_j\sigma_b} - \frac{1}{(\mathbf{k}_i-\mathbf{k}_b)^2}\delta_{\sigma_i\sigma_b}\delta_{\sigma_j\sigma_a}\right) & \text{if } \mathbf{k}_i+\mathbf{k}_j = \mathbf{k}_a+\mathbf{k}_b \\ 0 \text{ otherwise} \end{cases}, \tag{3.36}$$

where the summations are over occupied states $\mathbf{k}_i \in |D\rangle$, we have neglected the constant terms as they do not effect the dynamics and the box length is given as $L = \left(\frac{4\pi N}{3}\right)^{1/3} r_s$. Note again that Eq. (3.36) is for a very particular choice of excited determinant. In general we need to work out an additional phase factor which can arise when reordering the creation operators [112]. Since the Hamiltonian only couples Slater determinants of the same total momentum (momentum is conserved), $\mathbf{K} = \sum_{\mathbf{k}_{\text{occ}}} \mathbf{k}$, we see that there are no single-particle excitations allowed.

At $T = 0$ we can use momentum conservation to write the Hamiltonian matrix in block diagonal form and only solve for the ground state of the total momentum symmetry sector which yields the lowest energy. For a closed shell, finite-particle, UEG this will be the $\mathbf{K} = \mathbf{0}$ sector of the Hamiltonian. We can no longer use this saving at $T > 0$ and we instead need to average results over all symmetry sectors at once.

For the UEG, we can write the density matrix explicitly as

$$\hat{\rho} = \sum_{\zeta}\sum_{\mathbf{K}}\sum_{\mathbf{i},\mathbf{j}} \rho_{\mathbf{ij}} |D_{\mathbf{i}}(\zeta,\mathbf{K})\rangle\langle D_{\mathbf{j}}(\zeta,\mathbf{K})|, \tag{3.37}$$

where we used the fact that the Hamiltonian conserves total momentum and spin, and the configurational index $\mathbf{i}$ is understood to be for a given $\mathbf{K}$ and ζ. It is typical to restrict the sum over spin polarisations in Eq. (3.37) to either the unpolarised $\zeta = 0$ or fully polarised $\zeta = 1$ case. Thus, to sample the initial condition of $\beta = 0$ we distribute N_w walkers along the diagonal by randomly occupying $N_\uparrow$ and $N_\downarrow$ of the M plane waves states which lie within a sphere of cutoff radius k_c as shown in Fig. 3.1. If we wish to average over spin polarisation we allocate a fraction of the walkers to each spin polarisation sector of the Hilbert space, where the fraction is determined by the spin sector's size relative to the total size of the Hilbert space given as[5]

$$\text{Dim}(\mathcal{H}) = \sum_{\substack{N_\uparrow,N_\downarrow \\ \{N_\uparrow+N_\downarrow=N\}}} \binom{M}{N_\uparrow} \times \binom{M}{N_\downarrow}. \tag{3.38}$$

A simulation then proceeds using the steps outlined in Section 2.3. Unfortunately, we find that a direct application of the DMQMC algorithm as outlined in Section 2.3

[5] One can use Monte Carlo to estimate the relative sizes of the Hilbert spaces when this formula becomes unstable [112].

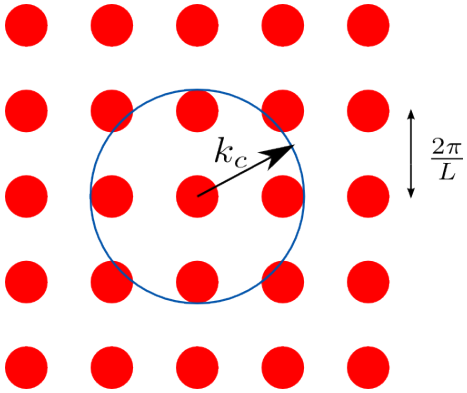

Figure 3.1: Allowed single particle states in a cubic simulation cell. Slater determinants are constructed by occupying any combination of N states inside the spherical cutoff of radius k_c.

to the UEG fails catastrophically. In Fig. 3.2 we show the internal energy per particle[6] evaluated using DMQMC with various walker numbers in a seven electron system in a small basis set of 81 plane waves. We see an unphysical knee region in the graph in the intermediate temperature regime followed by a slow decay to the ground state FCIQMC result.

To investigate this further we can try to reproduce this behaviour in an even smaller system where exact results are possible. In Fig. 3.3(a) we reduce M to 19 and restrict ourselves to the total momentum $\mathbf{K} = \mathbf{0}$ symmetry sector of the Hilbert space. Even here DMQMC can result in estimates for the internal that are too high in the intermediate temperature range. The source of this can be understood by considering the initial condition of infinite temperature. Here, we uniformly occupy any of the possible $\binom{M}{N}$ determinants so that the probability of selecting the Hartree–Fock determinant rapidly approaches zero. On those rare occasions when the Hartree-Fock determinant or another low-energy determinant is sampled at $\beta = 0$, the population of walkers arising from that low-energy determinant will dominate the simulation, but most simulations miss the low-energy part of the Hilbert space altogether. As shown in Fig. 3.3(b), this sampling problem reduces as the number of simulations (N_s) (or the population of walkers per simulation) increases, thus increasing the chance of sampling the low-energy space; however, this brute force sampling approach soon becomes impractical. For instance, even with the modest basis set size used to obtain Fig. 3.2,the computational cost to get the *wrong* answer increases to roughly 24 hours on 48 cores. Given that the accurate results we discuss later in this thesis require 1000s of plane waves to get accurate results, this brute force approach is doomed to fail.

[6] Throughout this thesis we use the convention that quantities per electron use a lower case letter. Thus, the internal energy per particle is $u = U/N$, the kinetic energy per particle is $t = T/N$ etc.

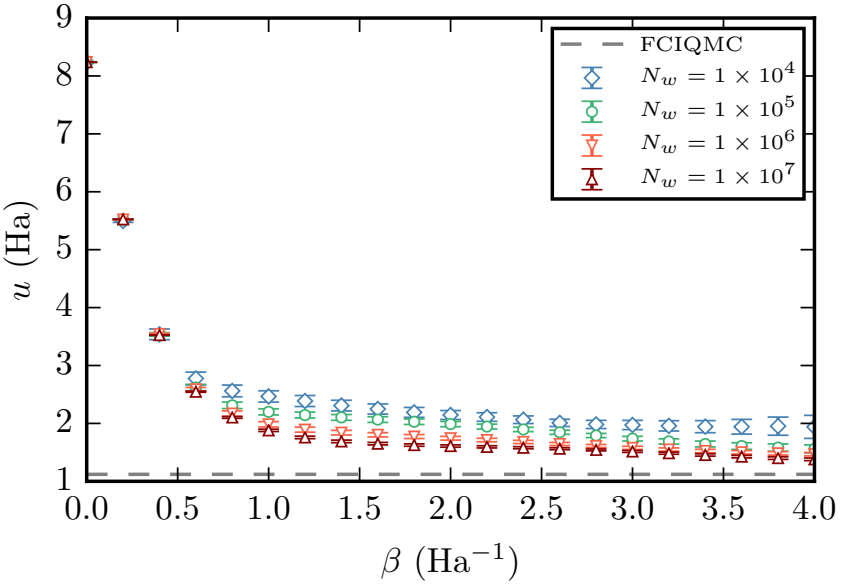

Figure 3.2: Internal energy per particle calculated using the conventional DMQMC algorithm for a UEG with $N = 7$, $\zeta = 1$ UEG with $M = 81$. The knee occurs roughly at 2.1 Ha, which is close to the first excited state for this system, suggesting the DMQMC settles down to the wrong ground state before further projection slowly brings it to the ground state.

3.2.1 Moving to the Interaction Picture

There are two sampling issues present in the original DMQMC algorithm; the distribution of weight in the density matrix changes rapidly as a function of β and important determinants are rarely present in our initial configurations. Feynman points out in Ref. [16] that if we write $\hat{H} = \hat{H}^0 + \hat{H}'$, where $\hat{H}'$ is small compared to $\hat{H}^0$, then the quantity $e^{\beta \hat{H}^0} \hat{\rho}$ will be a slowly-varying function of β. Building on this idea, we can define a new operator

$$\hat{f}(\tau) = e^{-(\beta-\tau)\hat{H}^0} e^{-\tau \hat{H}}, \tag{3.39}$$

with the properties

$$\hat{f}(\tau = 0) = e^{-\beta \hat{H}^0}, \tag{3.40}$$

$$\hat{f}(\tau = \beta) = e^{-\beta \hat{H}} = \hat{\rho}(\beta). \tag{3.41}$$

From Eq. (3.40) above we see that, by working with the operator $\hat{f}$, we can start the simulation from $e^{-\beta \hat{H}^0}$ instead of the identity. For weakly-correlated systems this should provide a good first approximation to the distribution of weight in the fully interacting density matrix.

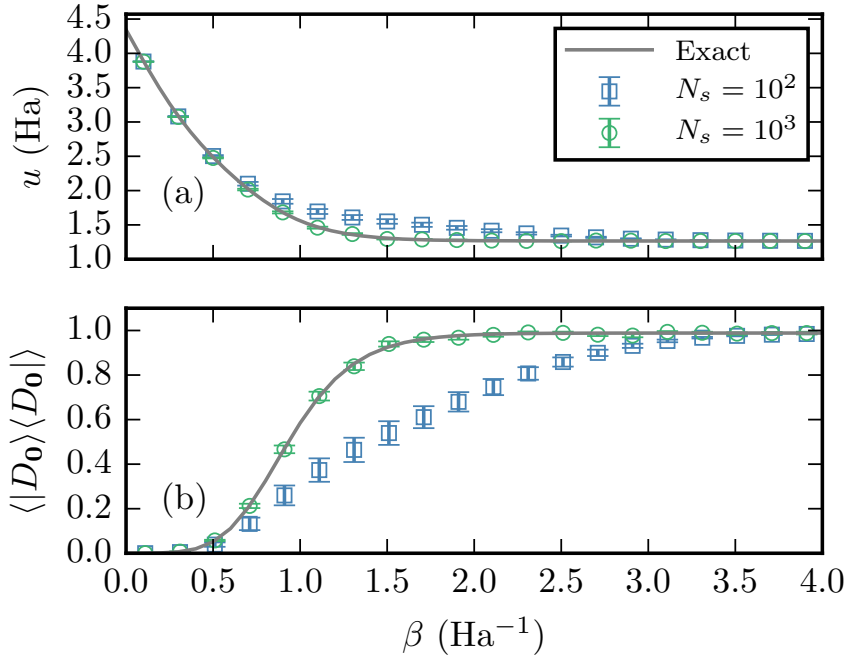

Figure 3.3: Panel (a) shows the energy calculated for a seven-electron spin-polarized electron gas at $r_s = 1$ with $M = 33$ plane waves in the total momentum $\mathbf{K} = \mathbf{0}$ subspace using DMQMC and an initial walker population $N_w = 10^4$. Increasing the number, N_s, of simulations over which averages are taken from 10^2 (squares) to 10^3 (circles) results in a more accurate answer being reproduced. We also see that the error bars do not reflect the true errors for $N_s = 10^2$ in the intermediate β regime. Panel (b) shows that the average occupation on the Hartree-Fock density matrix element ($|D_\mathbf{0}\rangle\langle D_\mathbf{0}|$) is too small in the intermediate temperature range. Reproduced with permission from Ref. [146].

Differentiating Eq. (3.39) with respect to τ we find

$$\frac{d\hat{f}}{d\tau} = \hat{H}^0 e^{-(\beta-\tau)\hat{H}^0} e^{-\tau\hat{H}} - e^{-(\beta-\tau)\hat{H}^0}\hat{H}e^{-\tau\hat{H}} \tag{3.42}$$

$$= \hat{H}^0\hat{f} - \hat{f}\hat{H} \tag{3.43}$$

We can write the derivative in a different way starting from Eq. (3.42) and inserting the identity operator $e^{-(\beta-\tau)\hat{H}^0}e^{(\beta-\tau)\hat{H}^0}$ after $\hat{H}$ and $\hat{H}^0$ to find

$$\frac{d\hat{f}}{d\tau} = \hat{H}^0 e^{-(\beta-\tau)\hat{H}^0} e^{(\beta-\tau)\hat{H}^0}\hat{f} - e^{-(\beta-\tau)\hat{H}^0}\hat{H}e^{(\beta-\tau)\hat{H}^0} e^{-(\beta-\tau)\hat{H}^0} e^{-\tau\hat{H}} \tag{3.44}$$

$$= e^{-(\beta-\tau)\hat{H}^0}\left(\hat{H}^0 - \hat{H}\right)e^{(\beta-\tau)\hat{H}^0}\hat{f} \tag{3.45}$$

$$= -\hat{H}'_I(-(\beta-\tau))\hat{f}, \tag{3.46}$$

where we have used the definition

$$\hat{O}_I(\alpha) = e^{\alpha\hat{H}^0}\hat{O}e^{-\alpha\hat{H}^0}, \tag{3.47}$$

for $\alpha \geq 0$ and noted that $\hat{H}^0$ commutes with $e^{-(\beta-\tau)\hat{H}^0}$. We call this the interaction-picture version of DMQMC (IP-DMQMC) and, as it is used exclusively throughout the rest of this thesis, we will drop the IP- prefix. In practice the exponential factors appearing in Eq. (3.46) are time-consuming to evaluate and so Eq. (3.43) is easier to work with. Since we choose $\hat{H}^0$ to be diagonal in our basis set, the only modifications to the original DMQMC algorithm are that $p_d(\mathbf{ij}) = \Delta\tau|H^0_{\mathbf{ii}} - H_{\mathbf{jj}}|$ and that now walkers only spawn along columns of the density matrix. We find this simple modification drastically increases the range of applicability of DMQMC in the case of the UEG as will be demonstrated in Section 3.4. Note, however, that by Eq. (3.41), only one temperature value specified by β is now available in a given simulation in contrast to the conventional DMQMC algorithm where *all* temperatures up to a given $\beta_{\max}$ are available. In principal it is possible to reweight estimators using the new technique to remove the factors of $e^{-(\beta-\tau)\hat{H}^0}$; however this is usually only possible for short $\beta - \tau$ as the required factors of $e^{(\beta-\tau)\hat{H}^0}$ can grow quite large resulting in noisy estimators.

3.2.2 Symmetric Form

Note that Eq. (3.46) is in an asymmetrical form as walkers can only spawn along columns of the density matrix. While formally there is nothing wrong with this algorithm, we have found that it can lead to sampling problems at low temperatures when attempting to evaluate properties which do not commute with the Hamiltonian (see Fig. 3.4). To understand this note that at low temperatures and in small systems most of the walkers at $\tau = 0$ are placed on the Hartree–Fock determinant, which are then propagated only along columns of the density matrix. Generally speaking, the true density matrix will be smeared out along the diagonal in our Hartree–Fock basis, but walkers can never reach these states due to being restricted initially. With increased sampling this problem should go away, but the convergence is slow as we need to resolve $e^{-\beta(E^0_{\mathrm{HF}} - E^0_{\mathbf{i}})}$ where $\hat{H}^0|D_{\mathbf{i}}\rangle = E^0_{\mathbf{i}}|D_{\mathbf{i}}\rangle$ for large β. An algorithm which allows walkers to move along rows or columns (similar to the original DMQMC algorithm) therefore seems like a good idea.

We can write such a symmetrised form as

$$\hat{f} = e^{-\alpha\hat{H}^0} e^{-\tau\hat{H}} e^{-\alpha\hat{H}^0},, \tag{3.48}$$

for $\alpha = \frac{1}{2}(\beta - \tau)$, so that

$$\frac{\partial \hat{f}}{\partial \tau} = \frac{1}{2}\{\hat{H}^0, \hat{f}\} - \frac{1}{2}(\hat{H}_I(-\alpha)\hat{f} + \hat{f}\hat{H}_I(\alpha)). \tag{3.49}$$

The spawning probabilities are now given as

$$p_s(\mathbf{ik} \rightarrow \mathbf{ij}) = \frac{\Delta\tau}{2p_{\mathrm{gen}}}|H_{\mathbf{kj}}|e^{-\alpha(E^0_{\mathbf{j}} - E^0_{\mathbf{k}})} \tag{3.50}$$

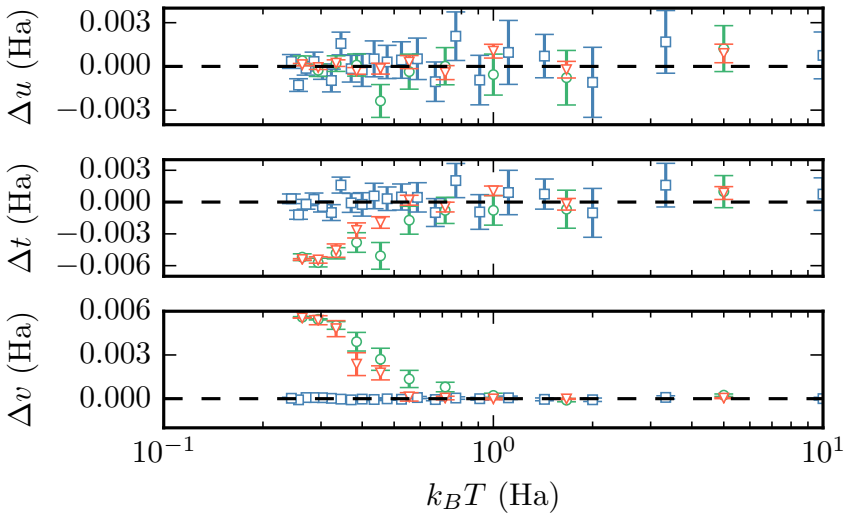

Figure 3.4: Deviation of the internal, kinetic and potential energy from exact (FCI) values for the symmetric (blue squares) and asymmetric (orange triangles and green circles) equations of motion. For the asymmetric equations we fixed the number of simulations and increased the walker population from $N_w = 10^3$ (green circles) to $N_w = 10^4$ (orange triangles). We see that better sampling slowly yields more accurate results for the potential and kinetic energy calculated using the asymmetric form. Note that both methods produce exact results for observables which commute with the Hamiltonian. The system shown here is for $N = 7$, $M = 19$ and $\zeta = 1$.

$$p_s(\mathbf{kj} \to \mathbf{ij}) = \frac{\Delta\tau}{2p_{\text{gen}}}|H_{\mathbf{ki}}|e^{-\alpha(E^0_{\mathbf{i}}-E^0_{\mathbf{k}})} \tag{3.51}$$

We see above that spawning from low energy to high energy states ($E_{\mathbf{k}} - E_{\mathbf{j}} > 0$) is suppressed while spawning from high to low energy configurations ($E_{\mathbf{j}} - E_{\mathbf{k}} < 0$) is greatly enhanced. This enhancement and suppression is diminished as $\alpha \to 0$. These large exponential factors can result in large blooming events (multiple walkers spawned onto a single determinant) near the beginning of the simulation, but do not affect results when $\tau = \beta$. In principle it should be possible to generate moves with probability

$$p_{\text{gen}}(\mathbf{ik} \to \mathbf{ij}) \propto e^{-\alpha(E^0_{\mathbf{j}}-E^0_{\mathbf{k}})}, \tag{3.52}$$

so that the large exponential factors do not need to be explicitly evaluated, but this is quite complicated to achieve exactly and we did not attempt it. In practice, since we are normally only interested in results at $\tau = \beta$, this turns out to be not such a big problem. We see from Fig. 3.4 that using the symmetric algorithm vastly improves estimates of observables other than the internal energy, although at an added cost in terms of simulation time (roughly a factor of 2). Often, if it is only the internal energy we want to calculate, the asymmetrical form is preferred, particularly at low temperatures where the time required to complete a single run from $\tau = 0$ to $\tau = \beta$ can be significant.

3.2.3 Sampling the initial condition

The choice of $\hat{H}^0$ is somewhat arbitrary, but it should allow for an efficient sampling of $\hat{f}(\tau = 0)$ and this is most easily achieved if $\hat{H}^0$ is diagonal in our many-particle basis. In principle, any initial density matrix can be sampled using the Metropolis algorithm [88], but due to the sensitivity of our method to the initial conditions, we have found this approach problematic. This is particularly an issue at lower temperatures and larger basis sets due to the longer equilibration times required for the Metropolis algorithm. A different route, which is free from such issues, is to sample the grand canonical density matrix corresponding to $\hat{H}^0$ in such a way that the desired, canonical, distribution is reached.

The probability of selecting a particular determinant $|D_\mathbf{i}\rangle$ in the grand canonical ensemble is

$$P_{GC}(|D_\mathbf{i}\rangle) = \frac{1}{Z_{GC}} \prod_{i \in |D_\mathbf{i}\rangle} e^{-\beta(\varepsilon_i - \mu)}, \tag{3.53}$$

where we are assuming that we can write

$$\hat{H}^0 = \sum_i \varepsilon_i \hat{c}_i^\dagger \hat{c}_i, \tag{3.54}$$

and μ can be determined from Eq. (1.52). However, we wish to generate determinants in the canonical ensemble where the correct probability is

$$P_C(|D_\mathbf{i}\rangle) = \frac{1}{Z_C} \prod_{i \in |D_\mathbf{i}\rangle} e^{-\beta(\varepsilon_i - \mu)}. \tag{3.55}$$

The constant factor of $e^{\beta\mu N}$ appearing in the canonical probabilities is accounted for in Z_C.

We see that $P_C(|D_\mathbf{i}\rangle) \propto P_{GC}(|D_\mathbf{i}\rangle)$ and so, by independently occupying orbitals with probability given by the Fermi factors f_i (Eq. (1.53)) and then discarding those configurations with $\langle \hat{N} \rangle \neq N$, we attain the correct proportionality factor Z_{GC}/Z_C. Averaging over spin sectors using this technique is also simpler since we can simply relax the constraint of $N_\uparrow = N_\downarrow$. This method is slightly inefficient in the sense that only about one in $\sqrt{N}$ of the configurations sampled has the right value of N. However, this is really never an issue given the system sizes we can realistically sample using QMC. The chemical potential can be obtained by numerically inverting Eq. (1.52) in the appropriate finite basis set using, for example, the bisection or Newton-Raphson method [147]. A demonstration of the whole procedure is given in Fig. 3.5, where we see that $\langle \hat{H} \rangle$ is indeed a slowly varying function of τ and that the correct estimate is reproduced at $\tau = \beta$.

Finally, we note that any diagonal density matrix can be obtained by reweighting the

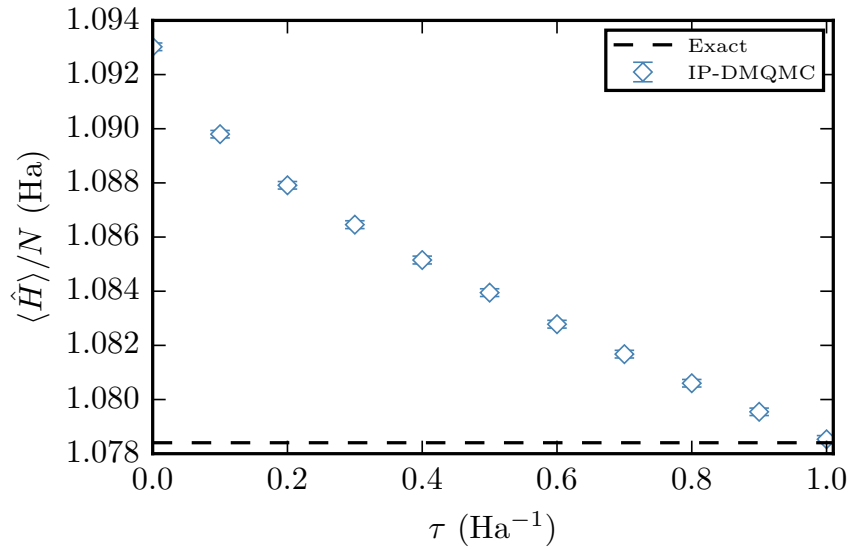

Figure 3.5: Variation of $\langle \hat{H} \rangle$ with τ using $\hat{H}^0 = \hat{T}$. The grand canonical procedure described in Section 3.2.3 was used. The system shown is a four-electron, spin polarized gas at $r_s = 1$ with $M = 81$ and $\beta = 1$. To obtain these results we used approximately 10^3 walkers and averaged over 100 simulations. The dashed line represents the exact FCI result, which DMQMC reproduces at $\tau = \beta$ as expected. For comparison, at $\beta = 0$, $\langle \hat{H} \rangle / N = 12.687(1)$ Ha from the original DMQMC algorithm. Reproduced with permission from Ref. [146].

configurations which result from the above sampling procedure as

$$P_{\text{new}}(|D_{\mathbf{i}}\rangle) = P_{\text{old}}(|D_{\mathbf{i}}\rangle) e^{-\beta(E^0_{\text{new}} - E^0_{\text{old}})}, \tag{3.56}$$

where E_{new} and E_{old} are the new and old total energies of a given configuration $|D_{\mathbf{i}}\rangle$, respectively.

3.2.4 Free Energies

Directly calculating free energies in QMC is a difficult task as one needs to evaluate the entropy $S = -k_B \text{Tr}[(\hat{\rho}/Z) \log(\hat{\rho}/Z)]$, which requires all the eigenstates and eigenvalues of $\hat{H}$ to be known. One typically resorts to coupling constant integration, which increases the computational effort and leads to potential additional sources of error. It turns out that the interaction picture formalism allows the exchange-correlation free energy[7] $F_{\text{xc}} = F - F_0$ to be calculated nearly for free.

To see this, we define

$$\tilde{F}(\tau) = -k_B T \log Z(\tau) \tag{3.57}$$

$$= -k_B T \log \text{Tr}\left[\hat{f}(\tau)\right], \tag{3.58}$$

[7] This is the appropriate definition for the UEG since there is no Hartree term.

for $\hat{f}$ defined as in Eq. (3.48) with $\tilde{F}(\tau = 0) = F_0$ and $\tilde{F}(\tau = \beta) = F$. Differentiating Eq. (3.58) with respect to τ we find

$$\frac{\partial \tilde{F}}{\partial \tau} = -k_B T Z(\tau)^{-1} \mathrm{Tr}\left[\frac{\partial \hat{f}}{\partial \tau}\right], \tag{3.59}$$

$$= -k_B T Z(\tau)^{-1} \mathrm{Tr}\left[(\hat{H}^0 - \hat{H}_I(-\alpha))\hat{f}\right], \tag{3.60}$$

$$= k_B T \langle \hat{V}_I(-\alpha)\rangle_\tau, \tag{3.61}$$

where we have inserted Eq. (3.49) for $\frac{\partial \hat{f}}{\partial \tau}$ and used the cyclicity of the trace and the definition of $\hat{V}_I(-\alpha) = e^{-\alpha \hat{H}^0}\hat{V}e^{\alpha \hat{H}^0} = e^{-\alpha \hat{H}^0}(\hat{H} - \hat{H}^0)e^{\alpha \hat{H}^0} = \hat{H}_I(-\alpha) - \hat{H}^0$ in arriving at Eq. (3.61). Finally, integrating Eq. (3.61) we find

$$F_{\mathrm{xc}} = k_B T \int_0^\beta \langle \hat{V}_I(-\alpha)\rangle_\tau \, d\tau, \tag{3.62}$$

which is the integral of a simple estimator over the course of the simulation.

Of course, nothing in life is free and there are some added complications. Consider evaluating the integral Eq. (3.62) using the simplest possible integration rule[8]. We can then write

$$I = \beta^{-1}\Delta\tau \sum_{\tau_i} \bar{f}(\tau_i), \tag{3.63}$$

$$= \frac{1}{N_\tau} \sum_{\tau_i} \frac{\frac{1}{N_s}\sum_s a^s(\tau_i)}{\frac{1}{N_s}\sum_s b^s(\tau_i)}, \tag{3.64}$$

where N_τ is the number of points at which the integrand is evaluated and in taking the average of the integrand we have again noted that $\overline{\left(\frac{a}{b}\right)} \neq \frac{\bar{a}}{\bar{b}}$ and have therefore refrained from interchanging the order of the summation over time steps and simulations. This form of Eq. (3.64) renders the estimation of the standard error in Eq. (3.62) quite tricky. Since points at neighbouring points in imaginary time are serially correlated we can not simply sum the variances for $\bar{f}(\tau_i)$, but instead we need to evaluate

$$\sigma_{\bar{I}}^2 \approx \sum_{\tau_i} \sigma_{\bar{f}(\tau_i)}^2 + \frac{2}{N_s} \sum_{i<j} \mathrm{Cov}[f(\tau_i), f(\tau_j)]. \tag{3.65}$$

We can use standard error propagation [148] to evaluate $\mathrm{Cov}[f(\tau_i), f(\tau_j)]$ to find a (biased) estimate as

$$\begin{aligned}\mathrm{Cov}[f(\tau_i), f(\tau_j)] \approx \frac{\bar{a}(\tau_i)\bar{a}(\tau_j)}{\bar{b}(\tau_i)\bar{b}(\tau_j)} \Bigg[&\frac{\mathrm{Cov}[a(\tau_i), a(\tau_j)]}{\bar{a}(\tau_i)\bar{a}(\tau_j)} + \frac{\mathrm{Cov}[b(\tau_i), b(\tau_j)]}{\bar{b}(\tau_i)\bar{b}(\tau_j)} \\ &- \left(\frac{\bar{a}(\tau_i)\bar{b}(\tau_i)\mathrm{Cov}[a(\tau_i), b(\tau_j)] + \bar{a}(\tau_j)\bar{b}(\tau_j)\mathrm{Cov}[b(\tau_i), a(\tau_j)]}{\bar{a}(\tau_i)\bar{a}(\tau_j)\bar{b}(\tau_i)\bar{b}(\tau_j)}\right)\Bigg],\end{aligned} \tag{3.66}$$

[8] The error arising from numerical integration is typically a few orders of magnitude smaller than that of the statistical error so that this is usually a safe thing to do.

but this estimator is not numerically stable and is rather complicated. A simpler, albeit approximate, estimate for the error can be found by assuming that the value of the integral for a particular simulation is not too bad an estimate for the mean and instead evaluate the error as

$$\sigma_{\bar{I}}^2 = \frac{1}{N_s(N_s-1)} \sum_s (I_s - \bar{I}')^2, \tag{3.67}$$

where

$$\bar{I}' = \frac{1}{N_s} \sum_s I_s, \tag{3.68}$$

and I_s is the estimate for the integral from a single simulation which are completely independent from one another. In other words, we have interchanged the order of summations which is something we said we should not do. As a rough measure of how good this error bar is, we can check if the difference between the true mean and the biased mean Eq. (3.68) is significant (usually within an order of magnitude of the error) and flag this point as not being reliable if it falls outside this threshold. To summarise, we evaluate the mean using Eq. (3.64) and the variance using Eq. (3.67) under the restriction that the biased mean is not too biased.

The second, more serious, issue is that in using the symmetric form of the density matrix given in Eq. (3.49) to overcome the sampling issue present in evaluating $\langle \hat{V}_I(-\alpha) \rangle_\tau$ we have, ironically, introduced another sampling problem. Ignoring normalisation, the integrand is

$$\langle \hat{V}_I(-\alpha) \rangle_\tau = \rho_{\mathbf{ij}}(\tau) V_{\mathbf{ji}} e^{-\alpha(E_{\mathbf{j}}^0 - E_{\mathbf{i}}^0)}, \tag{3.69}$$

which can be very large if $(E_{\mathbf{j}}^0 - E_{\mathbf{i}}^0) < 0$ and lead to noisy estimators for large α (small τ) (see Fig. 3.6). Using the asymmetric formulation leads to a more stable estimator for the free energy, but leaves us with the original sampling problem discussed in Section 3.2.2[9].

In Fig. 3.7 we investigate these two issues by comparing to exact digitalisation results. We see that the approximate calculation of the error works very well at most temperatures, giving us confidence that this is a good approach. At lower temperatures we see a systematic bias that can be understood in light of the preceding discussion and Fig. 3.6. Unfortunately we have not found a good solution to the latter problem but, since the free energy approaches the internal energy at lower temperatures, it is not that significant and can be monitored by investigating the quality of the integrand as in Fig. 3.6.

[9] In other words, there is a 'conservation of pain' - W.M.C. Foulkes.

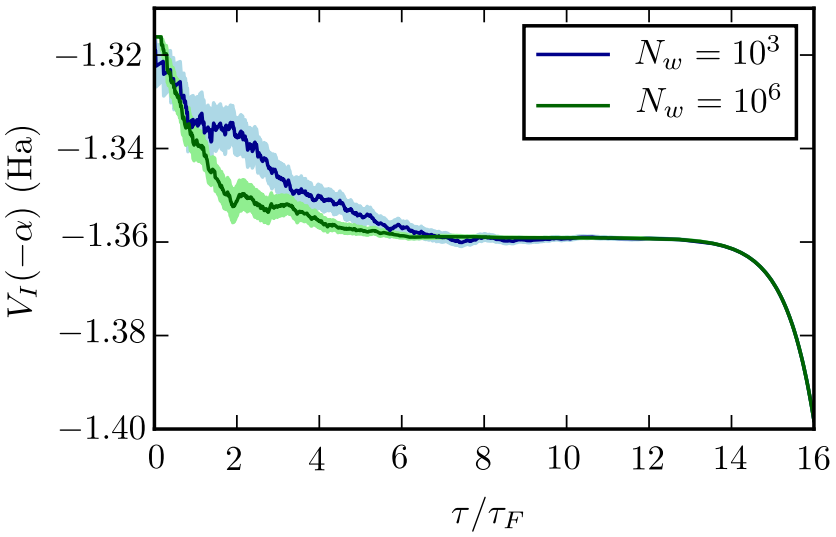

Figure 3.6: The noise in estimator for $\langle \hat{V}(-\alpha) \rangle$ at low τ. Here $\tau_F = E_F^{-1}$. Note that improving the sampling generally leads to better convergence. System shown is $N = 7$, $\zeta = 1$, $r_s = 1$ and $M = 19$ at $\Theta = 0.0625$.

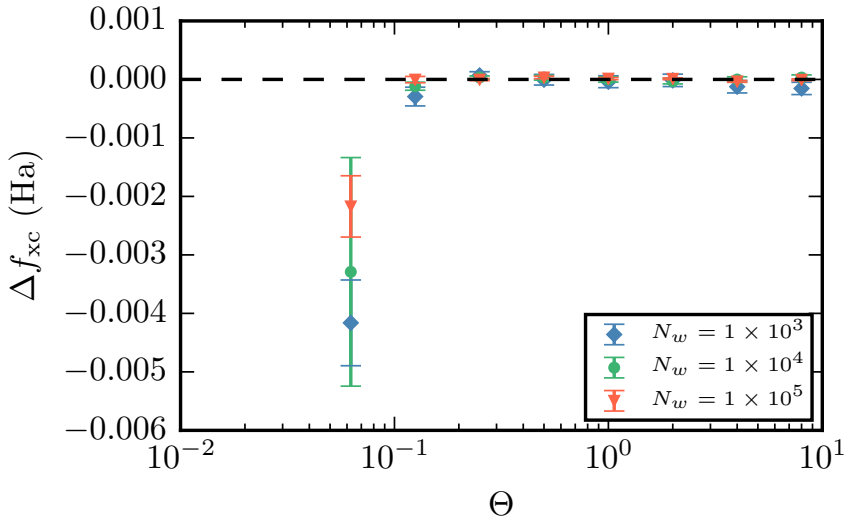

Figure 3.7: Deviation of the exchange correlation free energy from the exact FCI result for $N = 7$, $\zeta = 1$, $r_s = 1$ and $M = 19$ with various walker numbers N_w.

Ideal Free Energy

The evaluation of the total free energy, $F = F_0 + F_{\rm xc}$, requires the ideal contribution evaluated in the canonical ensemble, which is again a non-trivial task[10]. Thankfully, Monte Carlo saves us again and we can reuse the ideas developed in Section 3.2.3.

Let $N_{\rm accept}$ be the number of N-particle states generated during a grand-canonical Monte Carlo run of $N_{\rm attempts}$, then define

$$\delta \equiv \frac{N_{\rm accept}}{N_{\rm attempt}} = \frac{e^{\beta\mu_0 N} Z_C^0(N)}{Z_{GC}^0}. \tag{3.70}$$

If follows then that

$$-k_B T \log \delta = -\mu_0 N + F_0 - \Omega_0, \tag{3.71}$$

$$F_0 = \Omega_0 + \mu_0 N - k_B T \log \delta, \tag{3.72}$$

which corrects the usual Legendre transform of Eq. (1.50) for the case of finite system sizes. Thus, we can calculate F_0 simply using

$$\Omega_0 = -k_B T \sum_{\mathbf{k}\sigma}^{\mathbf{k}_c} \log\left(1 + e^{-\beta(\varepsilon_{\mathbf{k}\sigma} - \mu)}\right), \tag{3.73}$$

and we can determine δ directly from the Monte Carlo run. The validity of this approach is demonstrated in Fig. 3.8.

[10] See, e.g., Ref. [149] for recent semi-analytic work.

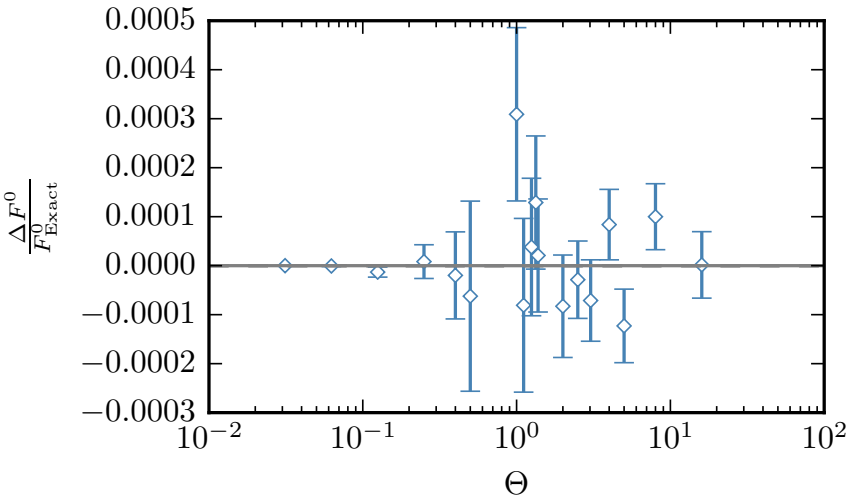

Figure 3.8: Relative deviation in the ideal free energy evaluated using Eq. (3.72) from exact results found by enumerating all ideal total energies and directly constructing Z_C. The system shown is for $N = 7$, $\zeta = 1$, $r_s = 1$ and $M = 19$. Note that the ideal free energy passes through zero between $\Theta = 0.5$ and $\Theta = 1$ which causes the error bars to grow in this region.

3.3 Basis Sets

The discussion so far has assumed that we are working in a finite basis set of M plane waves; to remove basis set incompleteness errors we need to extrapolate to the $M \to \infty$ limit. It is useful to understand how different parts of the internal energy converge with M so that reliable extrapolations can be performed. For the UEG at $T > 0$ there are essentially two contributions. The first is the convergence of the interacting kinetic energy, which can be mostly understood on a single-particle level. The second is the convergence of the potential energy where one needs to resolve cusps in the many-electron wavefunction. These cusps arise when electrons coalesce and are a many-body effect.

We will begin by trying to understand the convergence of the cusp, which is most easily done at $T = 0$. Previous work has shown that the correlation energy, $E_c = E - E_{\text{HF}}$, converges like M^{-1} for the unpolarised ($\zeta = 0$) case [103, 150, 151]. The fully polarised case ($\zeta = 1$) has not had as much attention, but is important as some of the largest deviations between CPIMC and RPIMC occurred for the polarised gas [12] and some early CPIMC results [152] assumed a M^{-1} convergence when extrapolating[11]. Results from the RPA suggest that the cusp error should converge like $M^{-5/3}$ [153, 154], and from a physical point of view, this faster convergence for the polarised gas can be understood in terms of the Pauli exclusion principle. Like spins are repelled and cannot exist in the same point in space, which reduces the strength of the Coulomb interaction

[11] Practically speaking this is of little importance if the error from extrapolation is small.

felt by any pair of electrons. A relatively straightforward way to arrive at the limiting behaviour is to use second order perturbation theory [151][12,13], which is the simplest non-trivial theory capturing the relevant physics.

Consider the conventional second order expression for the correlation energy of the UEG

$$E_c^{(2)} = -\sum_{\mathbf{i}\neq\mathbf{0}} \frac{|\langle D_\mathbf{i}|\hat{V}|D_\mathbf{0}\rangle|^2}{\langle D_\mathbf{i}|\hat{T}|D_\mathbf{i}\rangle - \langle D_\mathbf{0}|\hat{T}|D_\mathbf{0}\rangle} \tag{3.74}$$

$$= -\frac{1}{4}\sum_{ij\in|D_\mathbf{0}\rangle}\sum_{ab}\frac{|v_{ijab}-v_{ijba}|^2}{\varepsilon_a+\varepsilon_b-\varepsilon_i-\varepsilon_j}. \tag{3.75}$$

Let us pick a particular $ij \to ab$ pair and investigate Eq. (3.76) further. We can use the Slater-Condon rules (Eq. (3.36)) to find

$$\frac{|v_{ijab}-v_{ijba}|^2}{\varepsilon_a+\varepsilon_b-\varepsilon_i-\varepsilon_j} \propto \frac{\left(\frac{1}{(\mathbf{k}_i-\mathbf{k}_a)^2}\delta_{\sigma_i\sigma_a}\delta_{\sigma_j\sigma_b} - \frac{1}{(\mathbf{k}_i-\mathbf{k}_b)^2}\delta_{\sigma_i\sigma_b}\delta_{\sigma_j\sigma_a}\right)^2}{\mathbf{k}_a^2+\mathbf{k}_b^2-\mathbf{k}_i^2-\mathbf{k}_j^2}. \tag{3.76}$$

Using momentum conservation and writing the momentum transfer $\mathbf{q} = \mathbf{k}_i - \mathbf{k}_a$, we find that the term of the right hand side of Eq. (3.76) is

$$\frac{1}{\Delta}\times\left(\frac{1}{\mathbf{q}^2}\delta_{\sigma_i\sigma_a}\delta_{\sigma_j\sigma_b} - \frac{1}{(\mathbf{q}+\mathbf{k}_j-\mathbf{k}_i)^2}\delta_{\sigma_i\sigma_b}\delta_{\sigma_j,\sigma_a}\right)^2, \tag{3.77}$$

where $\Delta = \mathbf{k}_a^2+\mathbf{k}_b^2-\mathbf{k}_i^2-\mathbf{k}_j^2$. Splitting the terms into $\uparrow\uparrow$ and $\uparrow\downarrow$ interactions we have

$$E_c^{(2)}(ij\to ab) \propto \frac{1}{\Delta}\times\begin{cases}\left(\frac{2\mathbf{q}\cdot(\mathbf{k}_j-\mathbf{k}_i)+(\mathbf{k}_j-\mathbf{k}_i)^2}{\mathbf{q}^2(\mathbf{q}+\mathbf{k}_j-\mathbf{k}_i)^2}\right)^2 & \text{for } \uparrow\uparrow \\ \left(\frac{1}{\mathbf{q}^2}\right)^2 \text{ or } \left(\frac{1}{(\mathbf{q}+\mathbf{k}_j-\mathbf{k}_i)^2}\right)^2 & \text{for } \uparrow\downarrow\end{cases}. \tag{3.78}$$

Next we rewrite the summation over a and b in Eq. (3.75) as one over $\mathbf{q}$ and look at the residual q contributions beyond the cutoff wavevector $\mathbf{k}_c$. If we then assume the large $\mathbf{q}$ limit, so that $|\mathbf{q}| \gg |\mathbf{k}_i|$, $|\mathbf{q}| \gg |\mathbf{k}_j|$ and $\Delta \approx \mathbf{q}^2$, then we find the basis set error, $\Delta E_c^{(2)} = E_c^{(2)}(M=\infty) - E_c^{(2)}(M)$, goes like

$$\begin{aligned}\Delta E_c^{(2)} &= \sum_{\mathbf{q}>\mathbf{k}_c}^{\infty} E_c^{(2)} \approx \int_{\mathbf{k}_c}^{\infty} d\mathbf{q} E_c^{(2)} \\ &\propto \begin{cases}\int dq \frac{q^2}{q^8} = k_c^{-5} \propto E_c^{-5/2} \propto \left(\frac{M}{V}\right)^{-5/3} & \text{for } \uparrow\uparrow \\ \int dq \frac{q^2}{q^6} = k_c^{-3} \propto E_c^{-3/2} \propto \left(\frac{M}{V}\right)^{-1} & \text{for } \uparrow\downarrow\end{cases},\end{aligned} \tag{3.79}$$

where $M = \frac{\frac{4\pi}{3}k_c^3}{\rho_\mathbf{k}} \propto k_c^3 V$ and $\rho_\mathbf{k} = \frac{(2\pi)^3}{V}$ is the density of states. Thus, we arrive at the

[12] This argument for $T=0$, particularly the $\zeta=1$ case, is due to Dr. James Shepherd.

[13] Of course, second-order perturbation theory diverges for the thermodynamic limit UEG, but provides well defined results in a finite box. In this case the UEG is effectively an insulator with a small band gap at the Fermi level.

proposed limiting behaviour for the fully polarised gas where only ↑↑ interactions are allowed. For the unpolarised case the ↑↓ interactions dominate.

At $T > 0$ there is a competition between the convergence of the kinetic and potential energies with the basis set size. Strictly speaking $\langle \hat{V} \rangle$ also converges differently at $T > 0$, often in a non-monotonic fashion. This might seem surprising at first as we might expect the potential energy to always become more negative with increasing M. To understand why this happens, consider the expression for the exchange (free) energy in the grand canonical ensemble

$$F_{\mathrm{x}} = \langle \hat{V} \rangle_0^{\mathrm{GC}} = -\frac{1}{2V} \sum_{ij}^{\mathbf{k}_c} \frac{4\pi}{(\mathbf{k}_i - \mathbf{k}_j)^2} f_{\mathbf{k}_i} f_{\mathbf{k}_j} \delta_{\sigma_i \sigma_j}, \tag{3.80}$$

where the sum runs over all i and j in our basis set. As the basis set size increases (at a fixed temperature) higher energy states are more likely to become occupied. Since larger momentum transfers yield smaller contributions to exchange energy the exchange energy can therefore become less negative. This is demonstrated in Fig. 3.9 where we plot the evolution of the Fermi factor with increasing M. We see that the distribution of weight in Eq. (3.80) shifts from low $\mathbf{k}$ to higher $\mathbf{k}$ making the exchange energy less negative. The discussion of the convergence of the correlation energy at $T > 0$ is more complicated (not least because there is still some debate about what form it should take [74, 77]), but we expect the limiting behaviour to be the same as at $T = 0$.

To investigate this further we plot in Fig. 3.10 the total energy of a two-electron spin-polarized system, which can be solved using exactly diagonalization in large basis sets. We see that the total energy initially increases rapidly with basis-set size before appearing to saturate. From the inset of Fig. 3.10 we see that as the size of the basis set is further increased, a slight reduction in the total energy is observed, with the residual error apparently proportional to $M^{-5/3}$.

The initial increase of the total energy with respect to M at non-zero temperatures can be understood largely by looking at the non-interacting kinetic energy as a function of basis-set size, which is most easily analyzed in the grand canonical ensemble. The non-interacting basis-set error is purely kinetic and given as

$$\Delta T_0(M) = T_0(\infty) - T_0(M) \tag{3.81}$$

$$= \sum_{\mathbf{k} > \mathbf{k}_c} \varepsilon_{\mathbf{k}} f_{\mathbf{k}}. \tag{3.82}$$

For $\varepsilon_c = \frac{1}{2} k_c^2 \gg 1$, this can be approximated as

$$\Delta T_0 \approx \int_{\varepsilon_c}^{\infty} \varepsilon^{3/2} e^{-\beta(\varepsilon - \mu)} d\varepsilon \tag{3.83}$$

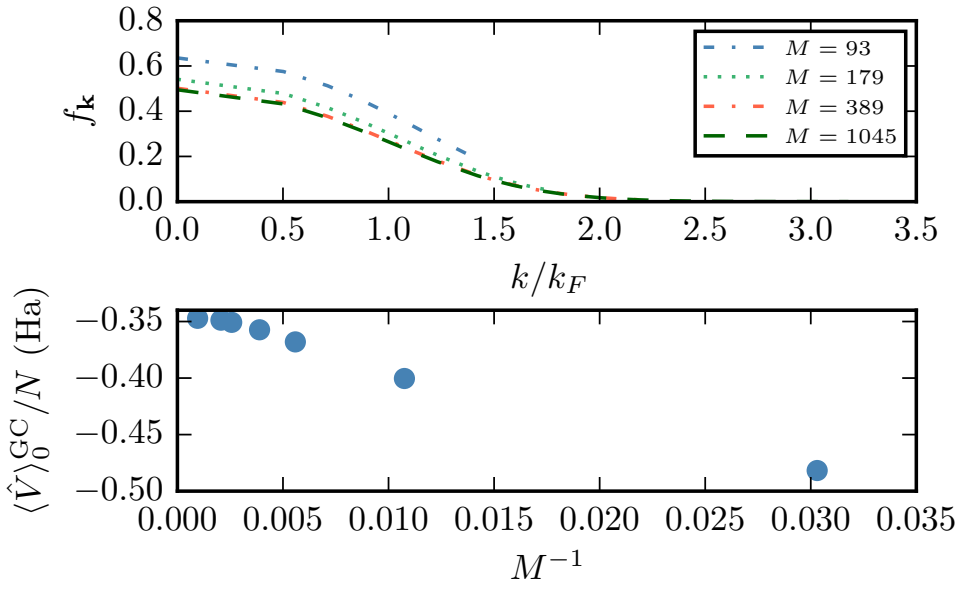

Figure 3.9: Evolution of the single-particle occupations and exchange free energy in the grand canonical ensemble with increasing basis set size. The system here is a $N = 66$, $r_s = 1$, $\zeta = 0$ spin-polarised UEG at $\Theta = 1$. The top panel shows the evolution of the Fermi factor $f_{\mathbf{k}}$ with increasing basis size M. Note that for a finite system only certain discrete values of k are allowed, which explains the non-smooth nature of the trend lines. The bottom panel shows the convergence of the exchange energy with basis set size, demonstrating that it does become less negative.

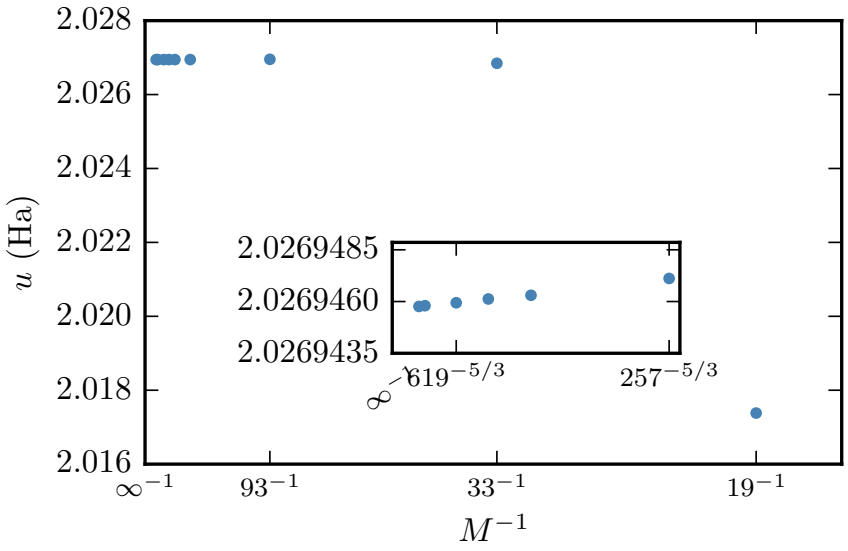

Figure 3.10: Behavior of the FCI internal energy per-particle with basis-set size for an $N = 2$, $r_s = 1$ spin-polarized system at $\Theta = 0.5$. Here we see the competition between the exponential convergence of the total energy at low M (main plot) and the $M^{-5/3}$ behaviour for high M (inset). Reproduced with permission from Ref. [146]

where we have used $f_{\mathbf{k}} \approx e^{-\beta(\varepsilon_{\mathbf{k}}-\mu)}$ for $\varepsilon_{\mathbf{k}} \gg \mu$. Hence

$$\Delta T_0(M) \approx \int_0^\infty (\varepsilon_c + x)^{3/2} e^{-\beta(\varepsilon_c+x-\mu)} dx \tag{3.84}$$

$$\approx \varepsilon_c^{3/2} e^{-\beta(\varepsilon_c-\mu)} \int_0^\infty e^{-\beta x} dx, \tag{3.85}$$

if $\varepsilon_c \gg \beta^{-1}$ so that $(\varepsilon_c + x)^{3/2} \approx \varepsilon_c^{3/2}$ everywhere $e^{-\beta x}$ is significant. It then follows that the leading-order correction is

$$\Delta T_0(M) \approx \beta^{-1} \varepsilon_c^{3/2} e^{-\beta(\varepsilon_c-\mu)}. \tag{3.86}$$

From Eq. (3.86) we see, unsurprisingly, that the kinetic energy begins to converge exponentially once $\varepsilon_c \approx k_B T$ or

$$M \approx V \left(\frac{\Theta}{r_s^2} \right)^{3/2} \approx N \Theta^{3/2}. \tag{3.87}$$

Thus, we see that for large Θ the kinetic energy and hence the total energy converge quite slowly as functions of M. This is an issue for DMQMC simulations as the cost of a calculation currently increases roughly linearly with basis-set size.

We can overcome this problem by instead extrapolating the temperature-dependent correlation energy, $U_c(\beta, M) = U(\beta, M) - U_{\mathrm{HF}}(\beta, M)$. Here the 'Hartree–Fock' energy can be calculated from $U_{\mathrm{HF}} = \langle \hat{H} \rangle_{\mathrm{HF}}$, using the density matrix[14],

$$\hat{\rho}_{\mathrm{HF}} = \sum_{\mathbf{i}} e^{-\beta E_{\mathbf{i}}^{\mathrm{HF}}} |D_{\mathbf{i}}\rangle\langle D_{\mathbf{i}}|, \tag{3.88}$$

where $E_{\mathbf{i}}^{\mathrm{HF}} = \langle D_{\mathbf{i}} | \hat{H} | D_{\mathbf{i}} \rangle$ and the sum runs over all determinants in the basis set. The infinite basis-set total energy can then be reconstructed as

$$U(\beta, M = \infty) = U_{\mathrm{HF}}(\beta, M = \infty) + U_c(\beta, M = \infty), \tag{3.89}$$

with the hope that $U_c(\beta, M)$ is small and converges more rapdily as is often the case with energy differences. The problem remains to calculate $U_{\mathrm{HF}}(\beta)$ in the canonical ensemble. Fortunately we can again use the sampling procedure outlined in Section 3.2.3, i.e.,

$$U_{\mathrm{HF}} = \frac{1}{Z_{\mathrm{HF}}} \sum_{\mathbf{i}} E_{\mathbf{i}}^{\mathrm{HF}} e^{-\beta E_{\mathbf{i}}^{\mathrm{HF}}} \tag{3.90}$$

$$= \frac{1}{Z_{\mathrm{HF}}} \sum_{\mathbf{i}} E_{\mathbf{i}}^{\mathrm{HF}} e^{-\beta(E_{\mathbf{i}}^{\mathrm{HF}} - E_{\mathbf{i}}^0)} e^{\beta E_{\mathbf{i}}^0} \tag{3.91}$$

$$= \frac{\sum_{\mathbf{i}} E_{\mathbf{i}}^{\mathrm{HF}} w(\mathbf{i}) p(\mathbf{i})}{\sum_{\mathbf{i}} w(\mathbf{i}) p(\mathbf{i})}, \tag{3.92}$$

[14] This is not the usual thermal Hartree–Fock approach discussed in Section 1.5.

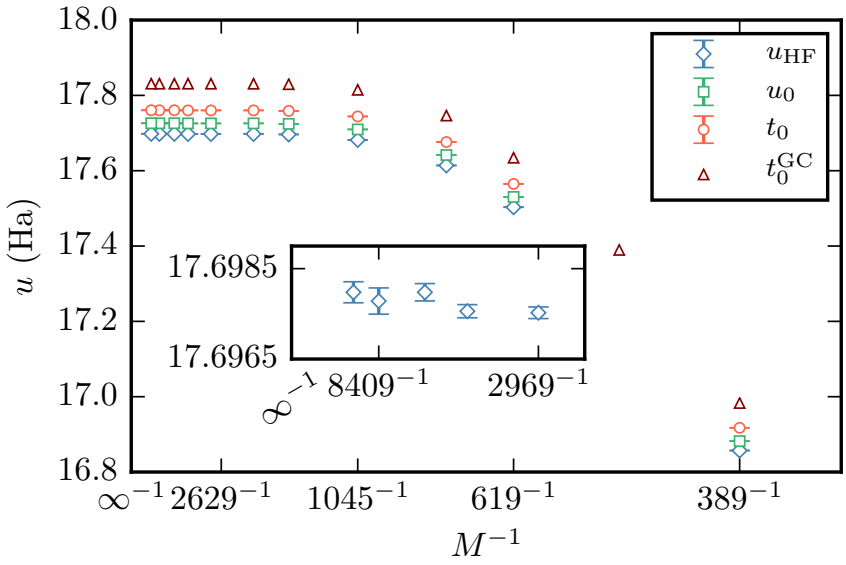

Figure 3.11: Exponential convergence of different estimates for the internal energy per particle with basis-set size for $N = 4$, $r_s = 1$ and $\Theta = 4$. See main text for definitions. Note that no Madelung contribution is included for the Hartree-Fock estimates in this figure. The infinite basis-set limit for U_{HF} is estimated as 17.6980(2) Ha. Reproduced with permission from Ref. [146].

where $w(\mathbf{i}) = e^{-\beta(E_{\mathbf{i}}^{\mathrm{HF}} - E_{\mathbf{i}}^0)}$ and $p(\mathbf{i}) = Z_0^{-1} e^{-\beta E_{\mathbf{i}}^0}$. Thus, by generating determinants as described in Section 3.2.3 and reweighting them using $w(\mathbf{i})$, we can instead sample $\hat{\rho}_{\mathrm{HF}}$ and, as a result, estimate U_{HF} as desired. In Fig. 3.11 we show the convergence of $U_{\mathrm{HF}}(\beta, M)$ as a function of basis set for a four-electron, spin-polarized system at $r_s = 1$ and $\Theta = 4$. Note the large basis-set sizes required to converge the total energy to within statistical error bars. Fig. 3.11 also shows the kinetic ($t_0 = \langle \hat{T} \rangle_0 / N$) and ideal exchange plus kinetic energy ($u_0 = \langle \hat{H} \rangle_0 / N$) calculated in the ideal canonical ensemble as functions of M. Any of these could in principle be subtracted from $U(\beta, M)$ to define a (exchange-)correlation energy, but the quantity defined by subtracting $U_{\mathrm{HF}}(\beta, M)$ extrapolates most smoothly to the infinite M limit. The non-interacting grand canonical energy ($t_0^{\mathrm{GC}} = \langle \hat{T} \rangle_0^{\mathrm{GC}} / N$, see Eq. (3.82)) is significantly larger than the canonical estimates.

Fig. 3.12 shows how $U_c(\beta, M)$ depends on M at a number of different temperatures. For small basis sets U_c shows a power-law decay with M, but this ceases for large enough M and the energies begin to increase again. The increase is due to kinetic effects that are not captured in the non-interacting expression we subtract and depends on the definition of the correlation energy. Given that it is the free energy which is variational at $T > 0$, it is not surprising that the internal energy converges non-monotonically with M.

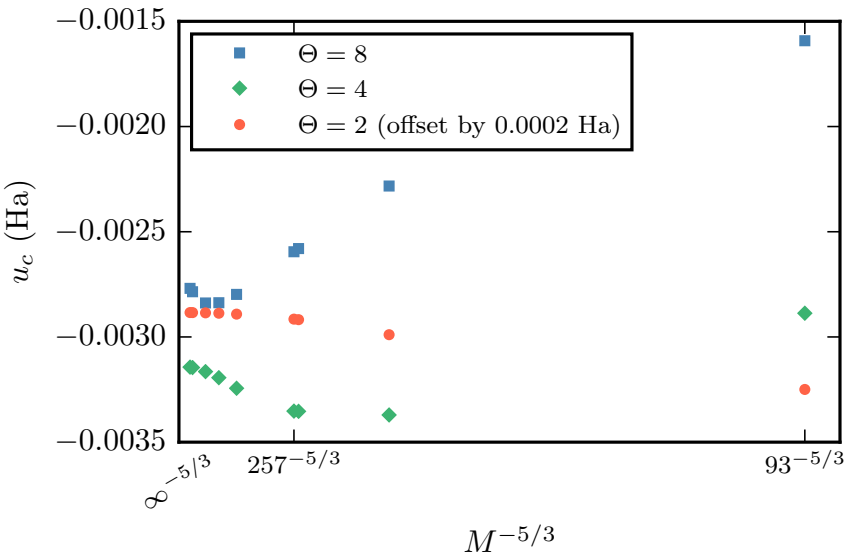

Figure 3.12: Comparison of the convergence of u_c with basis-set size calculated using exact diagonalization for a two-electron spin-polarized system at $r_s = 1$ for different values of Θ. The $\Theta = 2$ data has been shifted down by 0.0002 Ha for visibility. Reproduced with permission from Ref. [146].

We can estimate the infinite basis set result from the largest basis set $M_{\max}$ as

$$U(\beta, M = \infty) = U(\beta, M_{\max}) + U_{\mathrm{HF}}(\beta, M = \infty) - U_{\mathrm{HF}}(\beta, M = M_{\max}) \tag{3.93}$$

$$= U(\beta, M_{\max}) + \Delta U_{\mathrm{HF}}(\beta, M = \infty, M_{\max}), \tag{3.94}$$

where in practice we replace $M = \infty$ by a finite value which can be estimated with the help of Eq. (3.87). This will in general over estimate U_c (it will be too negative as we may miss slight upturn in the correlation energy at very high temperature as in Fig. 3.12), but the remaining discrepancy is typically smaller than the statistical error bar.

3.4 Results

Following all this development we are now in a position to carry out some preliminary benchmarking studies to convince ourselves of our implementation. We will focus on the simple case of four spin-polarized electrons, a small but non-trivial system and one for which there already exist benchmark CPIMC calculations [152]. We will only consider evaluating the internal energy as that was the only estimator presented in [152]. As a first step we compare our four-electron DMQMC results to FCI results in small basis sets and see perfect agreement across the whole temperature range (Fig. 3.13); and also that basis set errors are still significant.

Next we extend these results to basis sets far beyond the reach of conventional FCI; the

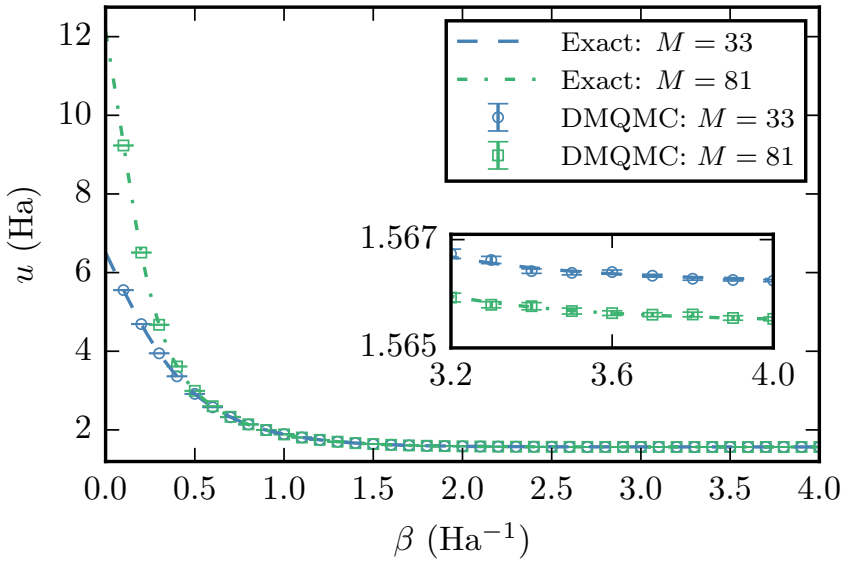

Figure 3.13: Comparison of DMQMC results (markers) with FCI results (dashed-lines) for the internal energy per-particle of the UEG with $N = 4$ and $r_s = 1$ in two different basis-sets. The inset shows the low T behavior where we see increasing the basis-set size serves to decrease the total energy in contrast to the high T behavior, where the opposite occurs. Reproduced with permission from Ref. [146].

largest density matrix sampled contains approximately 10^{22} matrix elements. We used the asymmetric formulation and the initialization procedure outlined in Section 3.2.3 and the free-electron Hamiltonian for $\hat{H}^0$ for $r_s \leq 1$; for $r_s > 1$ we found it advantageous to use the Hartree-Fock density matrix defined in Eq. (3.88). The calculations were initialized with 10^3–10^7 walkers and the results averaged over 100–5000 simulations, each using a different random number seed. No population control was used, although since population control these effects are typically small (since we use such large populations of walkers), population control should generally be used as it allows for a more efficient simulation. Time steps $\Delta\tau$ ranging from $0.01/E_F$ to $0.001/E_F$ were used, with a smaller time-step required at lower r_s; the values chosen were small enough that we could resolve no time-step error within the statistical errors. Each (r_s, Θ, M) calculation was typically run for 2 hours on 48 cores with a total computational cost of approximately 80000 core hours. The separate calculations of U_{HF} required 9000 core hours.

In Figs. 3.14 and 3.15 we show the convergence of the DMQMC results with basis set at low and high temperatures, respectively. We find that, for this system, a direct extrapolation of the total energy with respect to M is best for $\Theta \leq 0.25$; the non-interacting kinetic energy contributions are minimal in this regime and there is a clear cusp-related trend in the total energy. The basis set correction procedure outlined in Section 3.3 is best suited for temperatures above this, becoming increasingly useful above $\Theta = 2$ as more highly excited states become accessible. In between these too regimes both methods produce statistically identical results. Fig. 3.16 summarizes our

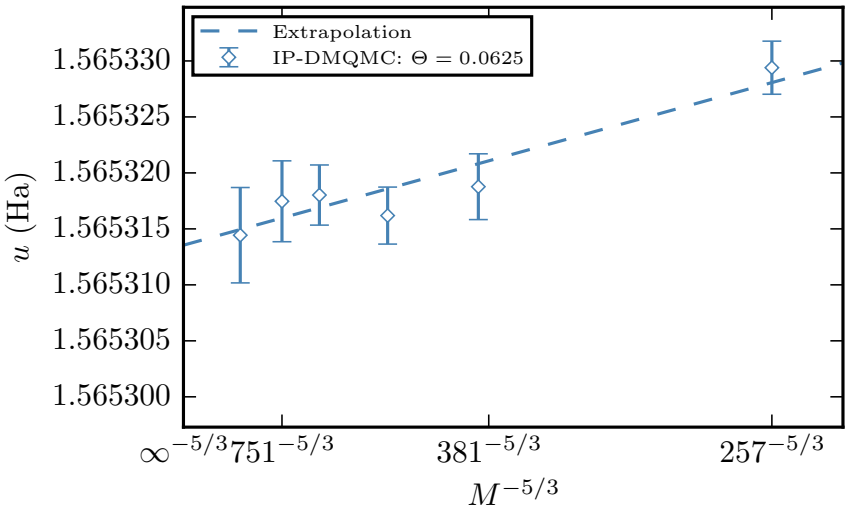

Figure 3.14: Total energy of a four-electron spin-polarized system at $r_s = 1$ and $\Theta = 0.0625$, showing a convergence with $M^{-5/3}$. The dashed line represents an extrapolation to the infinite-basis-set limit carried out using a weighted least-squares fit as implemented in Scipy [155]. Reproduced with permission from Ref. [146].

results and shows perfect agreement with the available CPIMC data from Ref. 152. Further results at higher temperatures and other r_s values are available in tabular form in the supplementary material of Ref. [146] and again agree with the available CPIMC results.

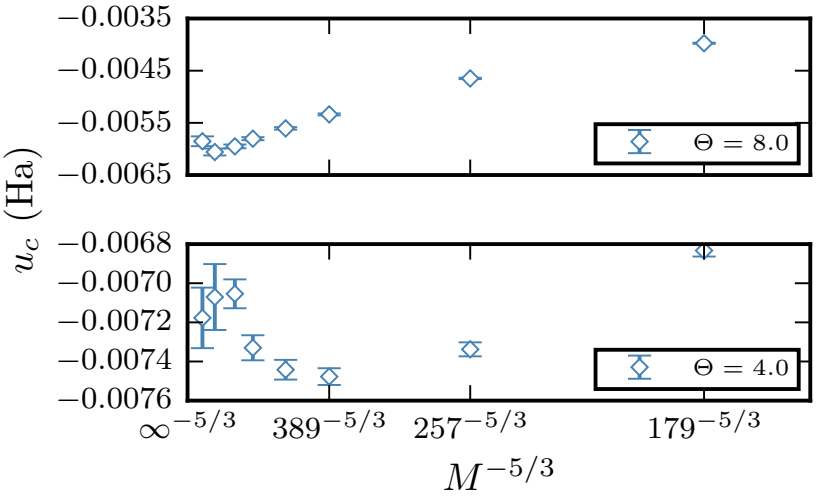

Figure 3.15: Convergence of u_c with basis-set size for the four-electron $r_s = 5$ UEG at $\Theta = 8$ (top panel) and $\Theta = 4$ (bottom panel) calculated using DMQMC. Note the similar behaviour to that found in the two-electron case in Fig. 3.12. Reproduced with permission from Ref. [146].

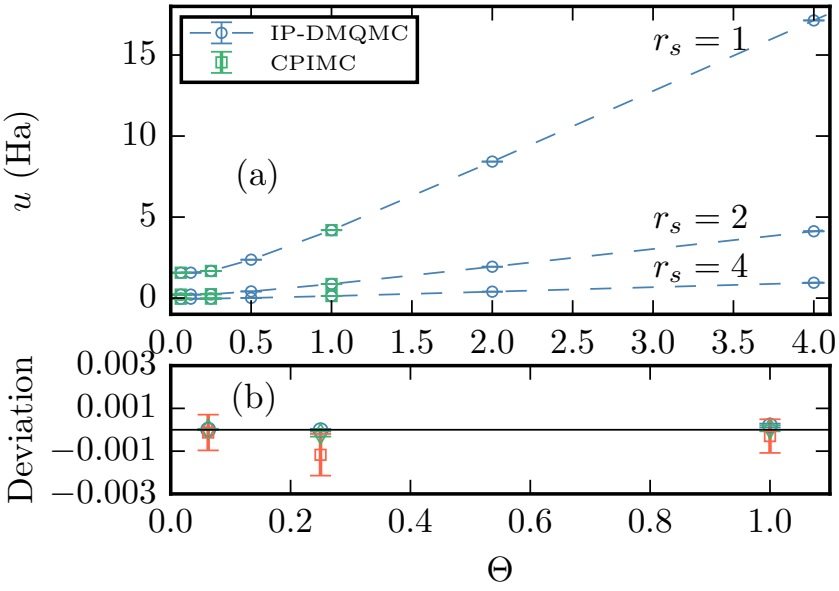

Figure 3.16: (a) Extrapolated DMQMC (blue circles) total energies per particle for the four-electron system at $r_s = 1, 2, 4$ showing exact agreement with the CPIMC results (green squares) of Ref. **??**. Dashed lines are meant as guides to the eye. (b) Relative deviation, $(E_{\mathrm{DMQMC}} - E_{\mathrm{CPIMC}})/E_{\mathrm{DMQMC}}$, as a function of temperature showing statistically identical results for $r_s = 1$ (circles), $r_s = 2$ (diamonds) and $r_s = 4$ (squares). Reproduced with permission from Ref. [146].

3.5 Summary

In this chapter we have shown how DMQMC can be adapted to treat the UEG. In overcoming sampling problems present in the original algorithm we have vastly extended the range of applicability of the method. This new formulation has the added benefit of granting direct access to the exchange-correlation free energy. By developing numerical basis-set corrections we have further reduced the computational cost in reaching the infinite basis set limit. These corrections are particularly useful at higher temperatures due to the significant thermal occupation of highly energetic one-electron states. Using these developments we have found exact agreement with CPIMC results for a small benchmark system across the whole temperature and density range, suggesting that both our approaches are sound. Of course, the systems treated here are too small to draw any conclusions regarding the disagreement between RPIMC and CPIMC for larger systems. To accomplish this we require the 'initiator' approximation investigated in the next chapter.

Chapter 4

The Initiator Approximation

The fundamental limitation on the system sizes and temperatures accessible to any QMC method is the sign problem. For systems without a sign problem, nearly arbitrarily large systems can be treated [156], whereas this number is reduced to about 10 spin polarised electrons for the case of the electron gas at $T = 0$ and $r_s = 1$ in FCIQMC. In the previous chapter we saw that, given a sufficient population of walkers, statistically exact results for nearly any estimator can in principle be attained using DMQMC for very small numbers of electrons. For DMQMC to be useful as something more than a benchmarking tool, much larger systems need to be treated. In this chapter we will test the application of the 'initiator' approximation (adapted from FCIQMC [157]) to DMQMC. In the process we will show that even given the limitations of the initiator approximation, useful and sufficiently accurate results can be attained for surprisingly large systems. Using this approximation we resolve the large disagreement between internal energies calculated using RPIMC and CPIMC [11, 12] at high densities and low temperatures and provide additional data in the intermediate density regime where none existed before. At higher temperatures we again find exact agreement with CPIMC and provide some of the first free energy QMC data for the UEG for finite particle numbers.

4.1 Origins

On the face of it, attempting to solve the many-electron Schrödinger equation using the linear expansion of the many-electron wavefunction in the space of Slater determinants is a hopeless thing to do. The memory required to store the expansion coefficients will never be available and even if it were the computational cost of doing anything with it does not bear thinking. For relatively weakly-correlated systems where mean-field theory is not such a bad approximation, the linear expansion can however become dominated by a handful of determinants with vanishingly small weight on highly excited determinants. We might be tempted to simply discard these low weight determinants and truncate the FCI expansions. However, doing this leads to results which are not size extensive and thus we cannot reach the macroscopic limit. FCIQMC is a very clever

algorithm as it provides a way to exploit the sparsity of the Hilbert space and instead sample these low-weight coefficients; and yet, since every determinant remains accessible, the results obtained remain size consistent. Unfortunately, for larger systems, it becomes harder to maintain the concentration of walkers required to maintain a healthy annihilation rate to overcome the sign problem and sample the wavefunction exactly [114]. To overcome this, Cleland *et al.* [157] introduced the initiator approximation for FCIQMC, which imposes some survival of the fittest criteria on the walkers and thus accelerates convergence with walker number.

The initiator approximation to FCIQMC (i-FCIQMC) amounts to imposing the following heuristic spawning criteria:

- spawning to unoccupied determinants is only permitted from a set of (initiator) determinants whose total walker population is above a user defined threshold called n_{add};
- spawning events from non-initiator determinants to unoccupied determinants are only permitted if the event is 'sign coherent', i.e., if multiple non-initiator determinants attempt to spawn onto the same unoccupied determinant with the same sign at the same time step;
- spawning to already occupied determinants is unaffected.

The origin of these rules can be understood by noting that that determinants that contribute most to the eventual solution originate from spawning events where the spawning probabilities $p_s(\mathbf{j} \to \mathbf{i}) \propto \Delta\tau |H_{\mathbf{ij}}| c_{\mathbf{j}}$ are largest so these should be favoured. Due to the vastness of the Hilbert space, walkers tend to proliferate in low weight areas and essentially get 'lost' which prevents the definite sign structure of the wavefunction from emerging. Somehow biasing walkers to stay in important areas of the Hilbert space should therefore help convergence. Restricting the spawning in this way naturally introduces a bias but this can be monitored by increasing the number of walkers, N_w, with the original algorithm recovered in the $N_w \to \infty$ limit (see Fig. 4.1). This approximation has met with considerable success and has been used in a variety of chemical [116, 117, 119], model [103, 150] and solid-state FCIQMC simulations [118]; it has also been applied to UEG systems of up to 54 unpolarised electrons at zero temperature [103, 150].

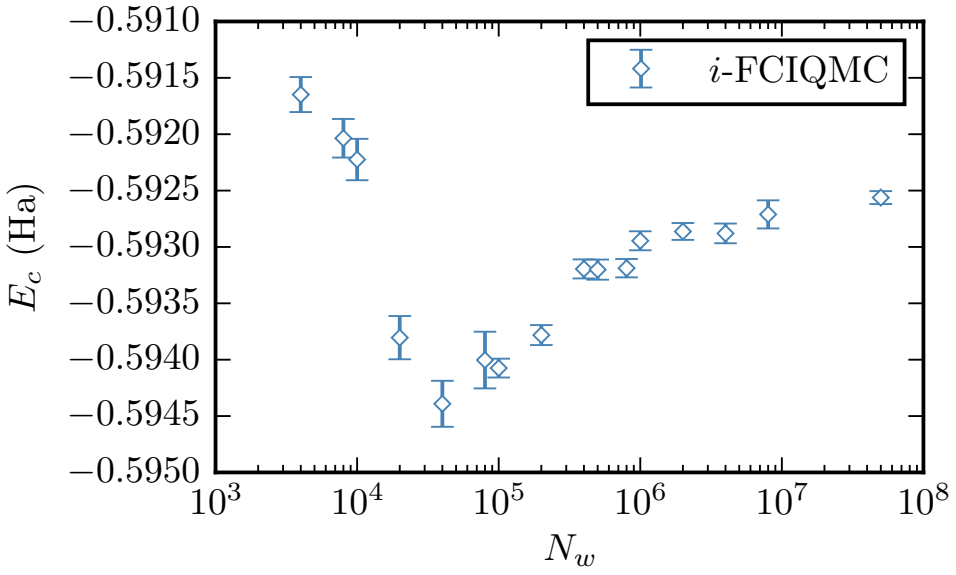

Figure 4.1: Demonstration of the convergence of the correlation energy (E_c) with walker number N_w of a UEG with $N = 14$, $\zeta = 0$ UEG at $r_s = 0.5$. Note that the projected energy estimator used to obtain this data is not a variational, so some care needs to be taken to ensure convergence. This calculation would not have been possible using the regular FCIQMC algorithm.

4.2 Adaptation to DMQMC

Given the success of i-FCIQMC and the promise from the initial DMQMC results for small electron gas systems, it is not unreasonable to hope that an analogous implementation if i-DMQMC would be equally as successful. Unfortunately, this turns out to be only partially true.

In DMQMC, a direct application of the initiator approximation can lead to incorrect averages at higher temperature as shown in Fig. 4.2. We can understand these biases by considering the evolution of the 'shape' of the density matrix as a function of temperature (see Fig. 4.3). Note we assume we are using the interaction-picture DMQMC algorithm throughout. At very high temperatures, the density matrix at $\tau = 0$ is close to the identity matrix, so that initially the likelihood of any given determinant having a population greater than n_{add} is very small. This significantly biases results as seen inFig. 4.2 (a) since effectively no spawning can occur. The exception here is for very low temperatures, where we approach an algorithm very similar to i-FCIQMC, and most of the walker distribution is initially located on the reference Hartree–Fock determinant. Thus, an obvious step to make is to set all diagonal elements of the density matrix to be automatically, and permanently, to be initiator determinants, free to spawn in an unrestricted way regardless of their occupation. We see from Fig. 4.2 (b) that this improves results at high temperatures, while at low temperatures results remain unaffected. There are, however, still significant deviations in the intermediate temperature

regime, which turns out to be the most difficult to sample using i-DMQMC[1]. Here, thermally excited states are important, so the walker distribution spreads up the diagonal, while off-diagonal elements become relatively more important due to the decrease in temperature. Allowing diagonal elements to spawn helps initially, but there is insufficient imaginary time available for the walkers to 'find' the important parts of Hilbert space dynamically, which leads to a much slower convergence with N_w. To overcome this issue, we have found that setting all walkers at excitation level less than or equal to 2 (from the diagonal) yields much better results (see Fig. 4.2 (c)). Here the excitation level, n_{ex}, is defined as the number of particle-hole excitations the bra and ket of a density matrix element differ by. Note that the use of a larger initiator space often leads to larger statistical errors as now more walkers end up away from the diagonal so that, for a fixed population, the weight on the diagonal is reduced.

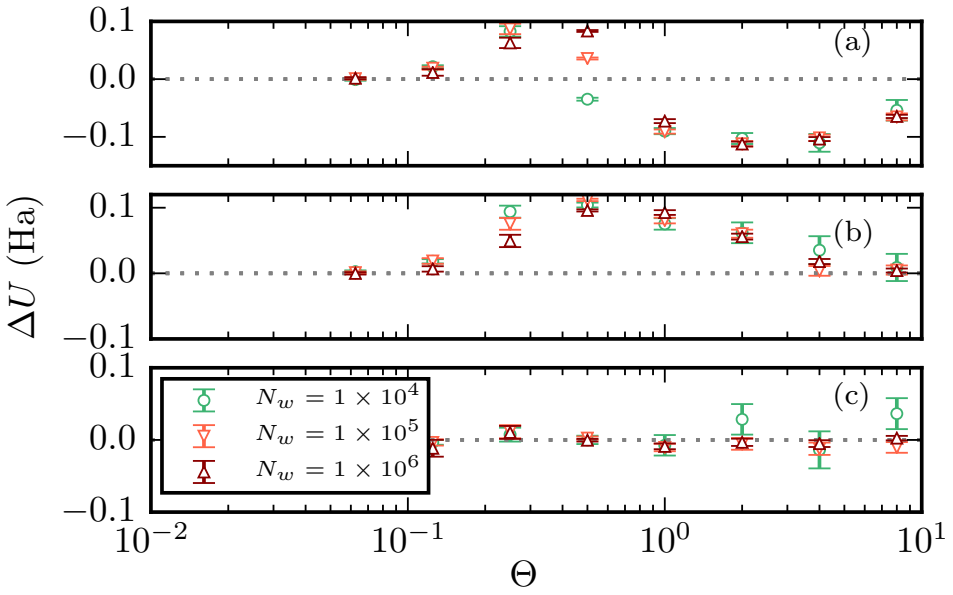

Figure 4.2: Deviation of i-DMQMC internal energies from exact diagonalisation for $N = 7$, $r_s = 1$, $\zeta = 1$ and $M = 19$. Here we used the asymmetric formulation of DMQMC, but similar behaviour is found for the symmetric case. We see that imposing either no initiator space (panel (a)), or just the diagonal (panel (b)) leads to biases in the intermediate temperature regime by effectively truncating the density matrix expansion. Panel (c) shows that setting all determinants at $n_{\text{ex}} \leq 2$ leads to much better convergence.

4.2.1 Other Properties

So far we have largely been concerned with determining the internal energy using DMQMC as this was usually the quantity which was compared in previous previous PIMC studies [11, 12]. One of the most appealing features of DMQMC is that we can

[1] Which is perhaps unsurprising given that it corresponds to the warm dense regime.

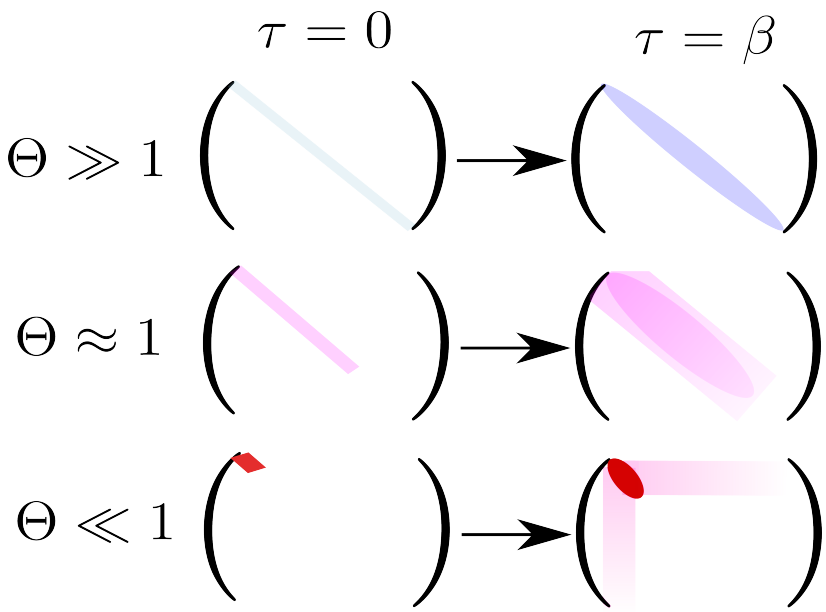

Figure 4.3: Schematic of the evolution of the distribution of weight in the density matrix with imaginary time in the interaction-picture formalism. The intermediate temperature regime is difficult to sample due to the relative weight on and off the diagonal being comparable.

determine the expectation value of nearly arbitrarily complicated operators and even the free energy in a simple and unbiased fashion. Unfortunately, we have found that the initiator approximation tends to bias expectation values of operators which do not commute with the Hamiltonian. Evidence of this is shown in Fig. 4.4, where we see that the convergence with walker number is much slower than for the case of the internal energy, and impractically large walker numbers are required, even for modest system sizes. For smaller systems where exact diagonalisation is possible, we do generally find convergence.

The origin of this slower convergence can be understood, in part, on variational grounds. At zero temperature there is a variational principle for the total energy. If we try to estimate the energy using some trial wave function ψ^T, then the errors in the energy are of the order $\delta\psi^2 = (\psi_0 - \psi^T)^2$, where ψ_0 is the true ground state wavefunction. Expectation values of operators which do not commute with the Hamiltonian suffer from errors on the order of $\delta\psi$ leading to slower convergence, so Fig. 4.4 may not come as such a surprise

At $T > 0$ the variational principle is for the Helmholtz free energy as calculated using Eq. (1.16). Note, however, that in DMQMC we calculate the free energy from $F = -k_B T \log Z$ via thermodynamic integration, which in turns depends on the expectation value of the potential energy. Thus we cannot hope (and we do not find) that our estimate for f_{xc} fares any better. If we could evaluate F using the functional form of Eq. (1.16) it is likely we would find a better estimate but this is hampered by the virtual

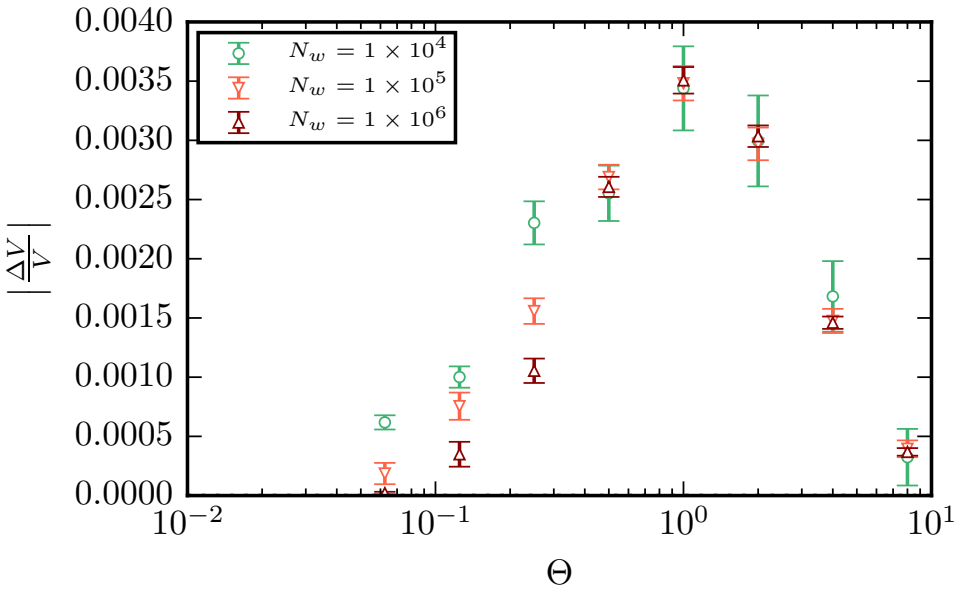

Figure 4.4: Absolute relative deviation of the i-DMQMC potential energy from the DMQMC values for $N = 7$, $r_s = 1$, $\zeta = 1$ and $M = 256$ calculated using the symmetric algorithm.

impossibility of determining the entropy.

The initiator approximation is, in some sense, a very strange approximation to make. We can interpret it as replacing the actual Hamiltonian with a modified operator which depends on the distribution of walkers at the same point in simulation time, such that the equations we now iterate are

$$\rho'_{\mathbf{ij}}(\tau + \Delta\tau) = \rho'_{\mathbf{ij}}(\tau) - \Delta\tau \sum_{\mathbf{k}} H_{\text{init}}(\rho'; \tau)_{\mathbf{ik}} \rho'_{\mathbf{kj}}(\tau), \tag{4.1}$$

where we have used the original, asymmetric, DMQMC algorithm for simplicity. With this interpretation, $\hat{\rho}$ is not necessarily of the form of $e^{-\beta \hat{H}'}$, and moreover, the differential equation obeyed by this modified density matrix is not simply the Bloch equation. Perhaps an alternative perturbative approach along using the expansion of Eq. (2.23) might lead to a more accurate estimator for the free energy, or at to a better understanding of the limitations of the current method. However, going beyond the initiator approximation is difficult to do and we have not found any useful solution to this problem thus far.

4.3 Results for larger system sizes

From these initial test cases we see that i-DMQMC can potentially allow us to resolve the disagreement between RPIMC and CPIMC, at least for the internal energy. To reach

these system sizes requires pushing i-DMQMC to its limits and carefully monitoring basis set, initiator and potential time step errors.

4.3.1 Convergence checks

Converging all the parameters in a DMQMC calculation is a computationally demanding challenge. Basis set and initiator errors may change with r_s and certainly change with temperature. Extensive investigation at $T = 0$ [150] has shown that typically the initiator error does not increase dramatically with basis set. Furthermore, while converging the wavefunction gets harder as r_s increases, this is not such a problem at the values of r_s accessible to DMQMC. Thus, in principle, we just need to converge the energy with respect to basis set size at a single r_s and then monitor convergence with N_w and r_s.

To check some of these assumptions we pick the relatively easy case of $r_s = 0.6$, $N = 33$ and $\zeta = 1$. In Fig. 4.5 we show the convergence of the i-DMQMC internal energy with walker number for different basis sizes. We see that the convergence is in general non-monotonic and that care needs to be taken to ensure convergence is achieved. For this low temperature of $\Theta = 0.0625$, we find that a basis set size of $M \approx 1045$ is sufficiently accurate for comparison to the CPIMC data. From the CPIMC results of Ref. [12] and additional DMQMC simulations we find this basis-set size is sufficient for all temperatures $\Theta < 0.5$ at this particle number.

In Fig. 4.6 we compare the internal energy calculated using the symmetric and asymmetric formulations. We see that in both cases very similar results are achieved, but the symmetric form leads to larger error bars for the same amount of computational effort. We thus stick to the asymmetric formulation at lower temperatures.

We find that the initiator error converges at a comparable rate in the intermediate temperature regime as shown in Fig. 4.7, which is consistent with Fig. 4.2. At higher temperatures, any obvious trend is masked by the naturally larger error bars as the temperature is increased, so care must be taken.

In Table 4.1 we test the convergence of the initiator error with r_s at the lowest temperatue of $\Theta = 0.0625$, where one expects the sign problem to be most severe. As expected, the initiator error grows with r_s but any bias is essentially negligible (estimated at less than 1%) for $N_w > 10^6$ and much smaller than $k_B T$. Pushing DMQMC to even lower densities is a challenge. We have found that, in certain cases, we can reach these regimes by extrapolating results obtained from a lower basis set, where the simulations are less demanding. This approach is illustrated in Fig. 4.8 for the case of $r_s = 4$ and $\Theta = 0.0625$. We see that the results are beginning to converge for the larger walker numbers and we can estimate the remaining systematic error as the difference between the extrapolated value for the two largest walker numbers.

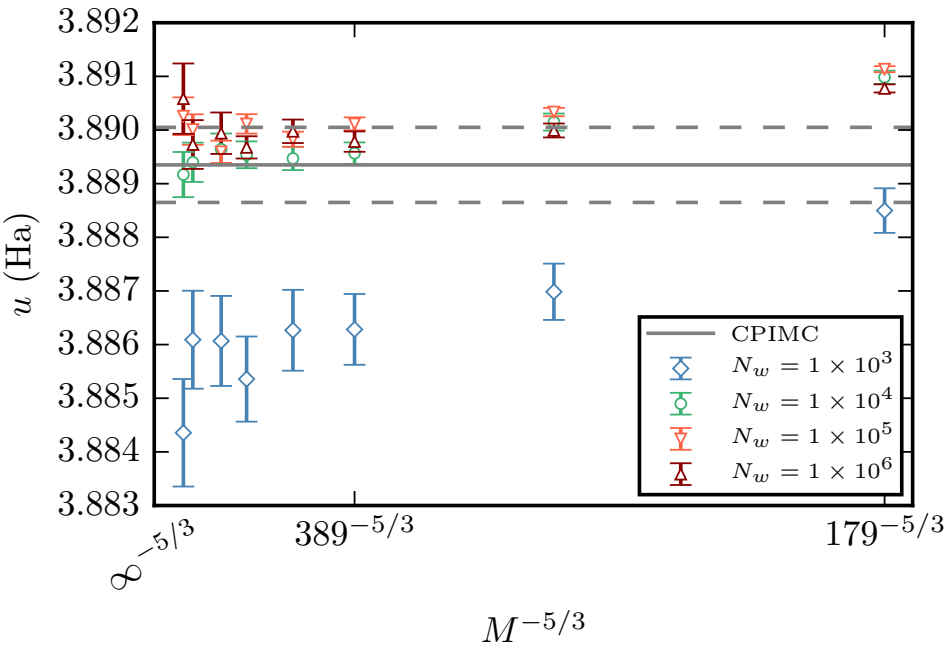

Figure 4.5: Convergence of the i-DMQMC internal energy with basis set size (M) for various different target walker populations (N_w) for $N = 33$, $r_s = 0.6$, $\Theta = 0.0625$. Here the asymmetric equations of motion were used and we ran with a fixed population. Also plotted is the CPIMC energy (solid line) and 1-σ error bars (dashed lines).

	r_s		
N_w	0.6	1	2
1×10^4	3.8894(2)	1.1458(2)	0.12491(5)
1×10^5	3.8899(1)	1.1460(2)	0.12423(9)
1×10^6	3.8893(2)	1.1455(2)	0.12223(7)
5×10^6	3.8893(2)	1.1452(2)	0.1211(5)

Table 4.1: Convergence of the i-DMQMC internal energy per particle with the target walker number, N_w, for $\Theta = 0.0625$, $N = 33$ and $M = 1045$ at a variety of r_s values with $n_{\text{add}} = 3$. For $r_s = 2$ we used $n_{\text{ex}} = 0$ and grew the population past $N_w = 10^5$, whilst for $r_s = 0.6, 1$ we used $n_{\text{ex}} = 2$ and kept the population fixed at N_w. The initiator error is smaller than 1 milli-Hartree per particle at $N_w > 10^6$.

Finally, since we are using a simple Euler update scheme propagate the density matrix, we can run into time-step errors. Fortunately, these tend to be small as shown in Fig. 4.9 and not distinguishable from statistical error for typical values of the time step $\Delta\tau$ we use in this work.

4.3.2 Higher Temperatures

For $\Theta \geq 1$ and $r_s \leq 1$ for $N = 33$ we find that no initiator approximation is required so that in principle exact results can be found straightforwardly. At the highest tem-

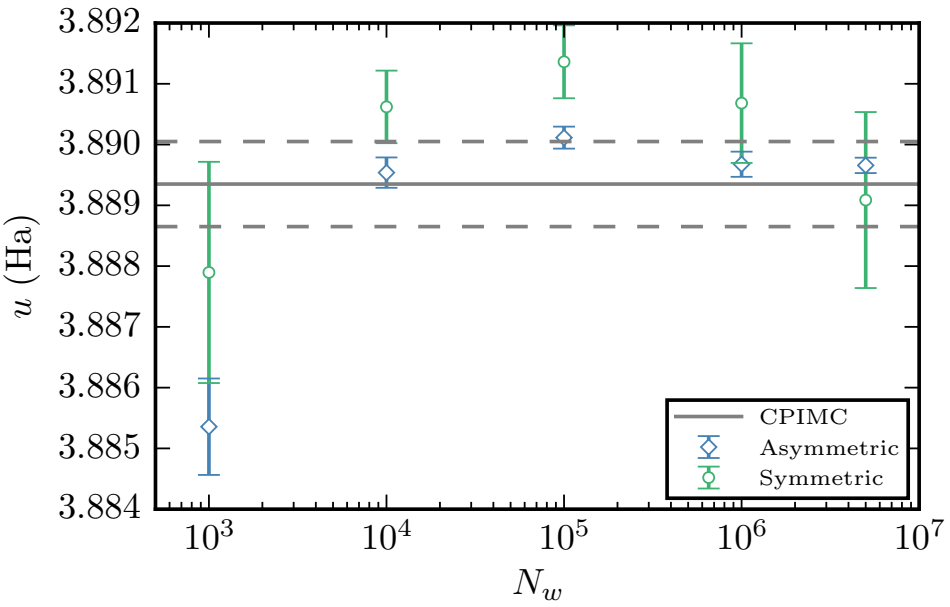

Figure 4.6: Convergence of the i-DMQMC internal energy with target population (N_w) for $N = 33$, $r_s = 0.6$, $\Theta = 0.0625$. We see that the convergence of initiator error differs when using the symmetric (squares) and asymmetric (circles) equations of motion. However, both methods produce identical results in the large N_w limit. The shift was varied throughout. Also plotted is the CPIMC energy (solid line) and 1-σ error bars (dashed lines) from Ref [12].

peratures we can use the basis set corrections developed in Section 3.3 to reduce the computational effort. In Fig. 4.10 we demonstrate the usefulness of these corrections for $r_s = 0.6$ and the hardest case of $\Theta = 8$. We see that applying the Hartree–Fock-like correction $\Delta u_{\mathrm{HF}}(M, \Theta)$ calculated using Eq. (3.88) dramatically speeds up convergence so that relatively modest basis sets can be used. For comparison, the direct convergence carried out in Ref. [12] required $M \approx 40000$ plane waves, which would represent a significant overhead for a DMQMC simulation, whose computational cost rises linearly with M for a fixed number of walkers.

In Figs. 4.11 and 4.12 we investigate the performance of the basis corrections for the free energy. First, in Fig. 4.11, we plot the total free energy for the interacting and ideal system, which exhibit a very similar behaviour. This suggests that an ideal contribution of $\Delta F^0(M, \Theta) = F^0(\infty, \Theta) - F^0(M, \Theta)$ based on the free energy of the ideal system should accelerate convergence, which, as we see, is indeed the case. Of course, this is just a statement that the exchange-correlation free energy converges more rapidly than the total free energy. However, we see in Fig. 4.12 that the residual basis set error for f_{xc} can be significant. Using the same logic as before, we can attempt to accelerate the convergence by adding a first-order (in the potential) correction to f_{xc}. This amounts to removing the contribution from f_{x}, i.e., we add $\Delta f_{\mathrm{x}}^0 = f_{\mathrm{x}}^0(\infty) - f_{\mathrm{x}}^0(M)$ where $f_{\mathrm{x}}^0 = N^{-1}\langle \hat{V} \rangle_0$ is the first order exchange contribution to the free energy per

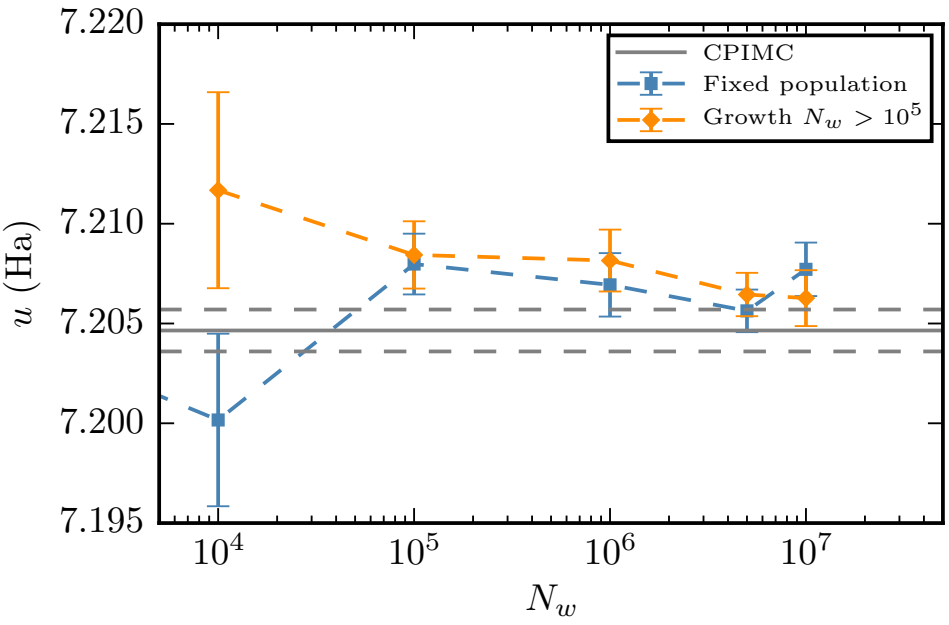

Figure 4.7: Convergence of i-DMQMC internal energy with target walker population (N_w) for $N = 33$, $r_s = 0.6$, $\Theta = 0.5$, $M = 1045$ and $\zeta = 1$. We also demonstrate that the choice of population control does not affect results (fixing the population at N_w and growing the population to N_w). Also plotted is the CPIMC result from Ref. [12] (solid line) and the corresponding error bar (dashed line).

electron. We see in Fig. 4.12 that the residual correlation free energy can be extrapolated more easily and we can estimate the remaining systematic basis set error as the difference between the linear extrapolation and the result in the largest basis set.

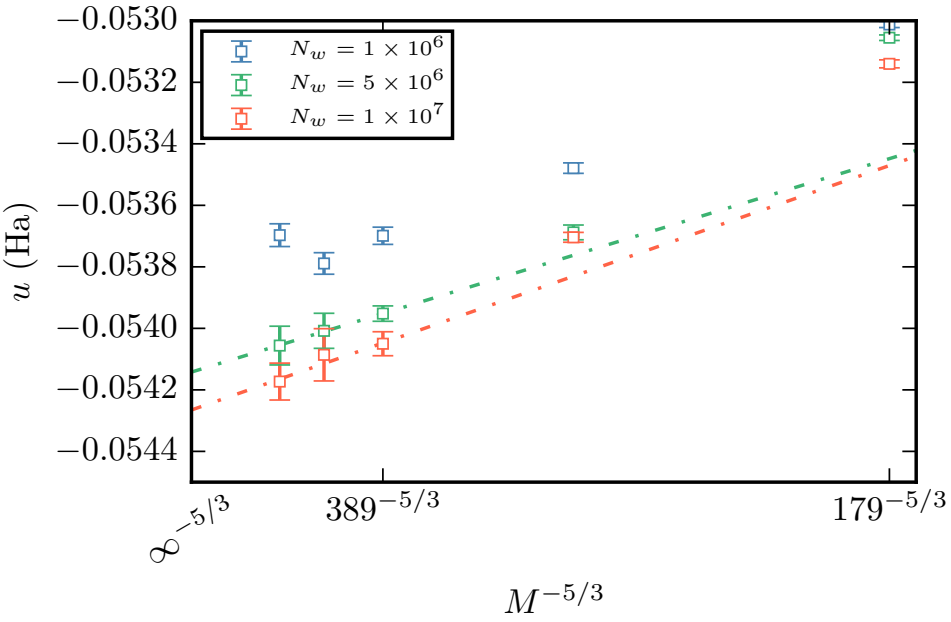

Figure 4.8: Convergence of i-DMQMC internal energy with target walker population (N_w) for $N = 33$, $r_s = 4$, $\Theta = 0.0625$ and $\zeta = 1$. Here, and for the $r_s = 3$ point in Fig. 4.13, the error is estimated as the difference between the extrapolated values for the two largest populations.

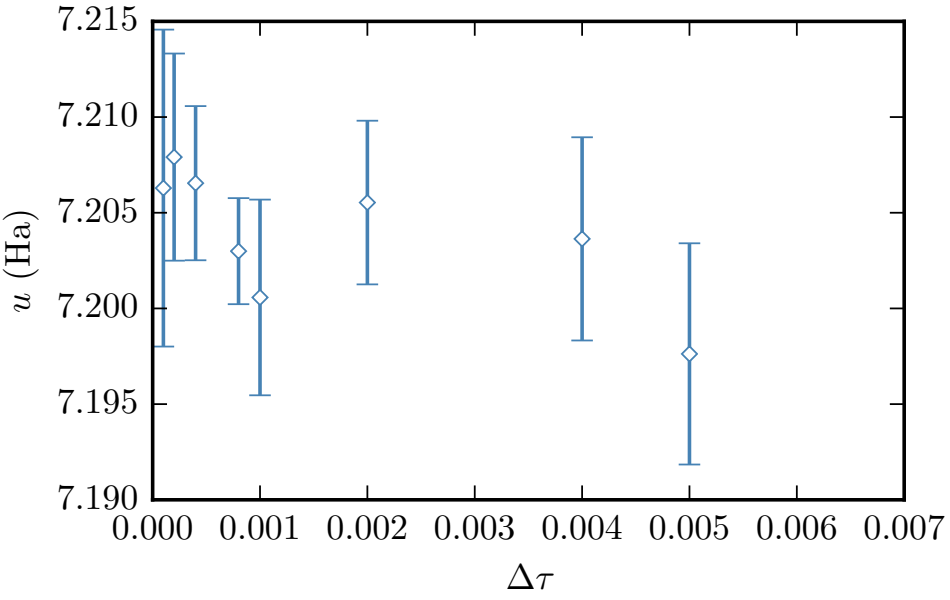

Figure 4.9: Variation of internal energy with time step $\Delta\tau$ for $N = 33, \zeta = 1$ and $r_s = 0.6$ at $\Theta = 0.5$ with $M = 751$.

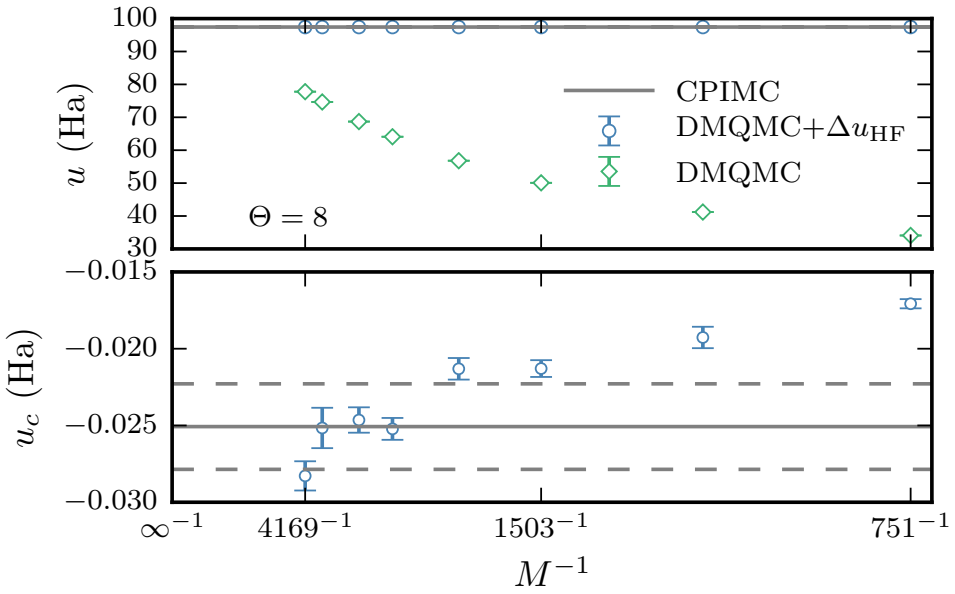

Figure 4.10: Top Panel: Slow convergence of the DMQMC energy as a function of basis-set size for $r_s = 0.6$ and $\Theta = 8$. Adding the Hartree–Fock ($\Delta u_{\mathrm{HF}}(M, \Theta)$) correction leads to a much smaller (relative) basis set error. Bottom Panel: The correlation internal energy ($u_c = u - u_{\mathrm{HF}}$)converges much more rapidly with M and is in good agreement with the CPIMC data from [12]. There is no particular significance to plotting against M^{-1} here.

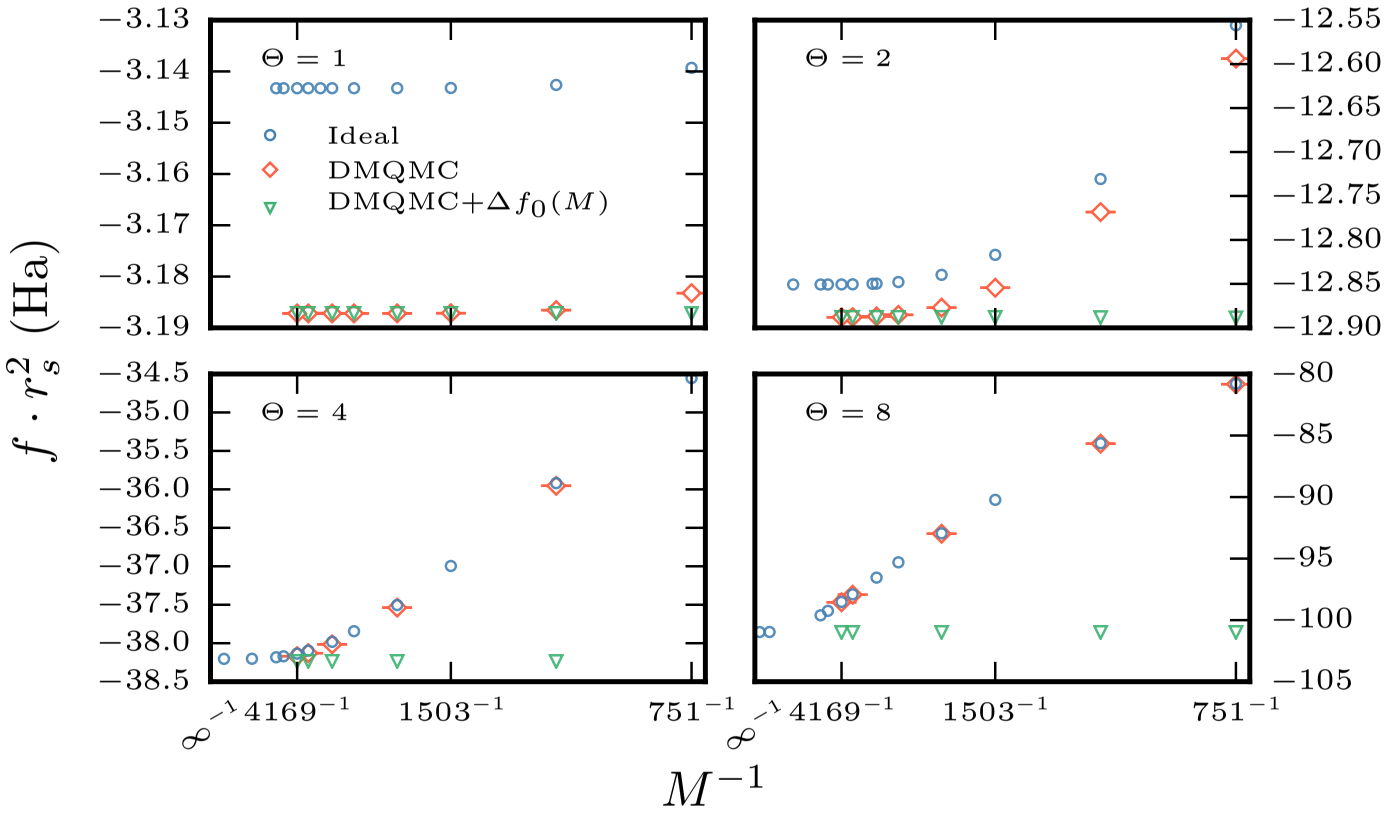

Figure 4.11: Convergence of the free energy calculated using DMQMC with basis set size for $N = 33$, $r_s = 0.1$, $\zeta = 1$ for various temperatures. We see that the free energy is variational with respect to increasing basis-set size which allows for a more systematic convergence. Also plotted is the ideal free energy for the same system; this is larger than the interacting free energy which is expected on variational grounds. The effects of interactions diminishing with increasing temperature.

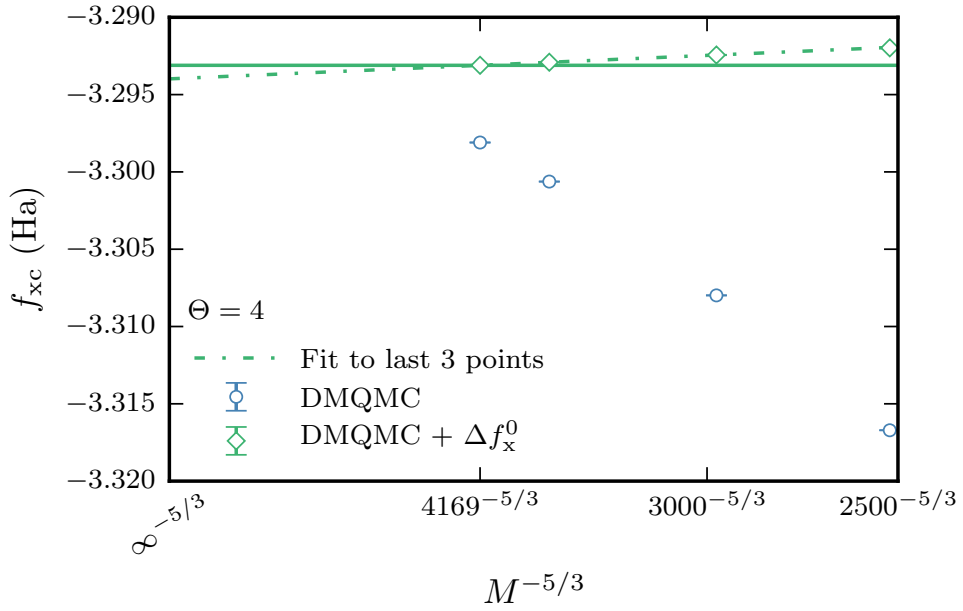

Figure 4.12: Basis set corrections applied to DMQMC data for f_{xc} for $N = 33, \zeta = 1, r_s = 0.1$ and $\Theta = 4$, showing the reduction in overall basis set error when corrections are applied. Here we reach an accuracy of ≈ 2 mHa per electron based on the difference between the maximum DMQMC data point considered at $M = 4169$ plane waves (horizontal line) and the extrapolated value (dot-dashed line). The extrapolation was performed using a weighted least squares fit [155] assuming an $M^{-5/3}$ behaviour to the last 3 points.

4.3.3 Resolving the Disagreement

We are now in a position to provide results for the $N = 33$, $\zeta = 1$ UEG. We first focus on the region $0.6 \leq r_s \leq 2$ and $0.0625 \leq \Theta \leq 0.5$, where the differences between the CPIMC and RPIMC results are largest and no other data are available [12].

All simulations used real amplitudes to represent the walkers in order to improve the stochastic efficiency [107, 108] and fixed $n_{\text{add}} = 3$. In most simulations we initialised the density matrix with N_w walkers and allowed the shift to vary throughout. For higher densities and the largest walker numbers we instead grew the total population to its target value. The results are averages over values from between 20 and a few thousand independent runs, depending on the temperature and target population considered. Typical values for the parameters used in our simulations can be found in Table 4.2.

r_s	N_w	M	$\Delta\tau \times E_F$
< 1	1×10^6	1045	1×10^{-3}
1	5×10^6	1045	0.5×10^{-3}
2	5×10^6	1045	0.25×10^{-3}
> 2	1×10^7	extrap Fig. 4.8	0.25×10^{-3}

Table 4.2: Typical parameters used in our i-DMQMC simulations. For $r_s \geq 2$, $\Theta \leq 0.125$ we set $n_{\text{ex}} = 0$, i.e., only diagonal elements were permanently set to be initiators. For $\Theta \geq 0.125$ we used the same parameters except with $M = 1503$ and $5 \times 10^6 \leq N_w \leq 1 \times 10^7$.

The i-DMQMC results for the exchange-correlation energy per particle presented in Fig. 4.13 are in very good agreement with the CPIMC results at all values of r_s up to the maximum of $r_s = 1$ considered by Schoof *et al.* [12]. (The sign problem prohibited the use of configuration PIMC at higher r_s for the temperatures considered.) The agreement is even better at lower r_s values. In particular, our results confirm that the kink-potential approximation used by Schoof [12] for $r_s \geq 0.6$ is well controlled and that the RPIMC results are significantly too low at $r_s = 1$. Our additional points in the physically important range $1 \leq r_s \leq 2$ ($1 \leq r_s \leq 4$ at low temperatures) further suggest that the RPIMC results are unreliable for all $r_s \leq 4$. We find a slight, apparently systematic, disagreement with CPIMC at $\Theta = 0.5$, although all points remain within error bars. The origin of this discrepancy is likely due to the slower convergence of the initiator error at higher temperature. This is combined with a slight tendency of the extrapolation with the number of kinks in CPIMC to underestimate the total energy. The true value of the total energy probably lies in the upper half of the CPIMC error bar, which is systematic in nature and not statistical.

To further confirm the accuracy of our results, we have carried out independent i-FCIQMC calculations of the internal energy at zero temperature. In Fig. 4.14 we check the convergence of the i-FCIQMC correlation energy with basis set size and walker number, where we see that any remaining systematic error is essentially negligible for

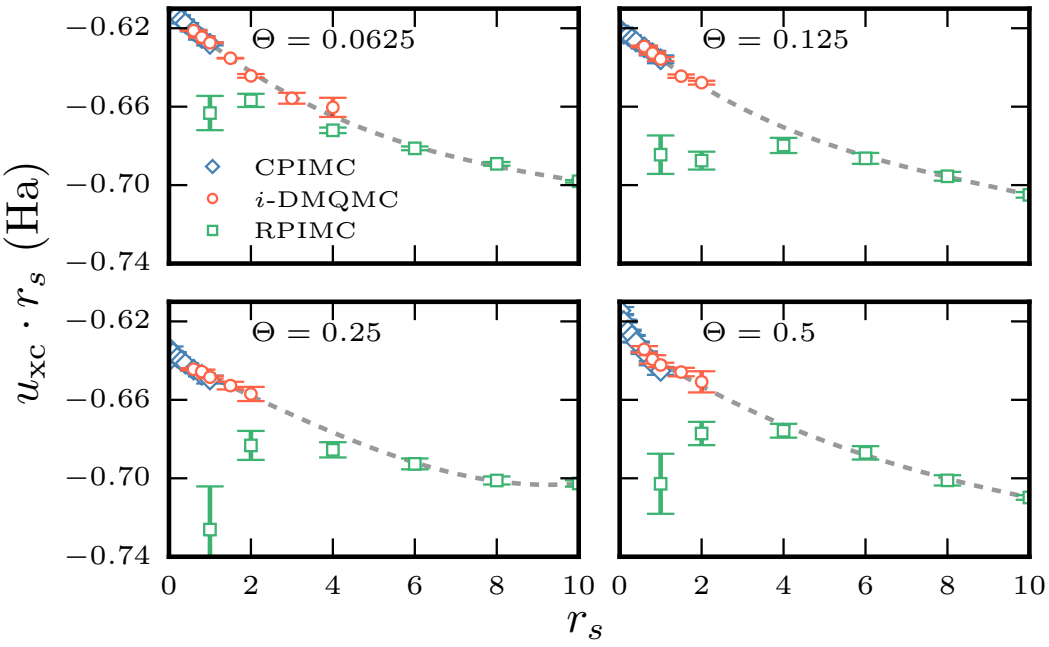

Figure 4.13: Exchange-correlation energy per particle (times r_s) as a function of r_s, showing excellent agreement between i-DMQMC and CPIMC for $r_s \leq 1$ and differences between i-DMQMC and RPIMC for $1 \leq r_s \leq 4$. The dashed lines are weighted third-order polynomial interpolations [158] between the i-DMQMC data and the restricted PIMC data for $r_s > 4$ and are meant as guides to the eye. Reproduced with permission from Ref. [54].

the accuracy considered here. Assuming that the energy varies like T^2 for small T, we can attempt to extrapolate the i-DMQMC and RPIMC results to zero temperature and compare them with the ground state result. Figure 4.15 shows that the extrapolated i-DMQMC energy agrees with the ground state result, but that the extrapolated restricted PIMC energy is too low. This is in contrast to the seemingly reliable extrapolation of the size-corrected restricted PIMC data performed in [11], which agreed well with the Perdew-Zunger parametrization of the $T = 0$ correlation energy of the UEG [138]. Also plotted in Fig. 4.15 is u_{HF} which which is seen to perform relatively well. In fact, it is no more inaccurate than the RPIMC result.

Above $\Theta = 0.5$ we found essentially identical results to CPIMC so we instead only present new data for f_{xc} in Fig. 4.16 where we also plot the exchange-correlation entropy ($s_{\text{xc}} = T^{-1}(u_{\text{xc}} - f_{\text{xc}})$). As expected, interactions tend to lower the entropy for this uniform system ($S \leq S_0$) by an amount that increases with r_s and vanishes in the high T limit. As $\lim_{T\to 0} s_{\text{xc}} = 0$ and given the behaviour of f_{xc}, we expect s_{xc} to reach a minimum in the warm dense regime which tends to counteract a similar minimum found in u_{xc} [11, 53].

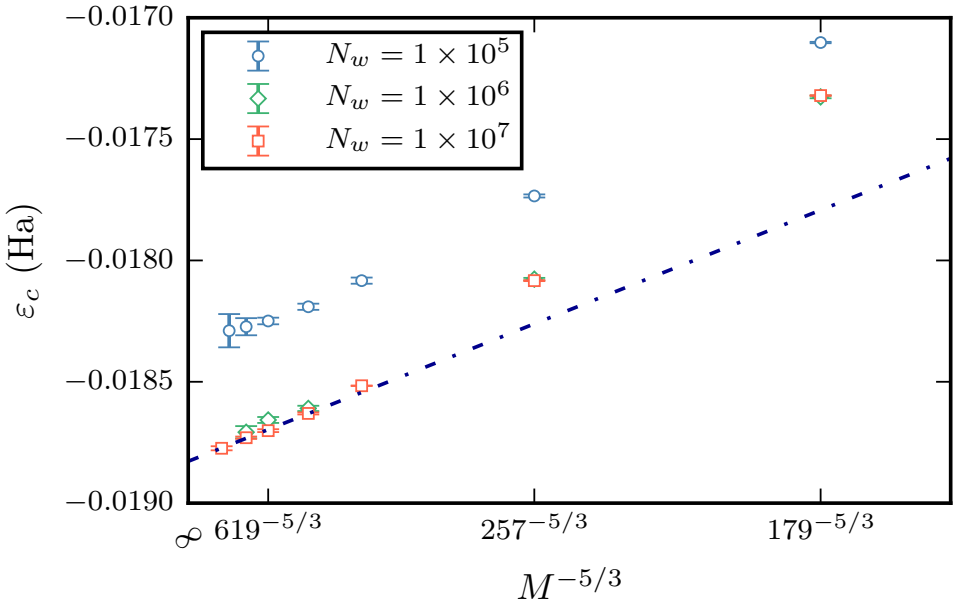

Figure 4.14: Convergence of the *i*-FCIQMC correlation energy (per particle) with basis-set size, M, and walker number, N_w, for $N = 33$, $r_s = 1$ and $\zeta = 1$. The dot-dashed line is a weighted least squares fit [155] of the last four ($N_w = 10^7$) points assuming a $M^{-5/3}$ dependence.

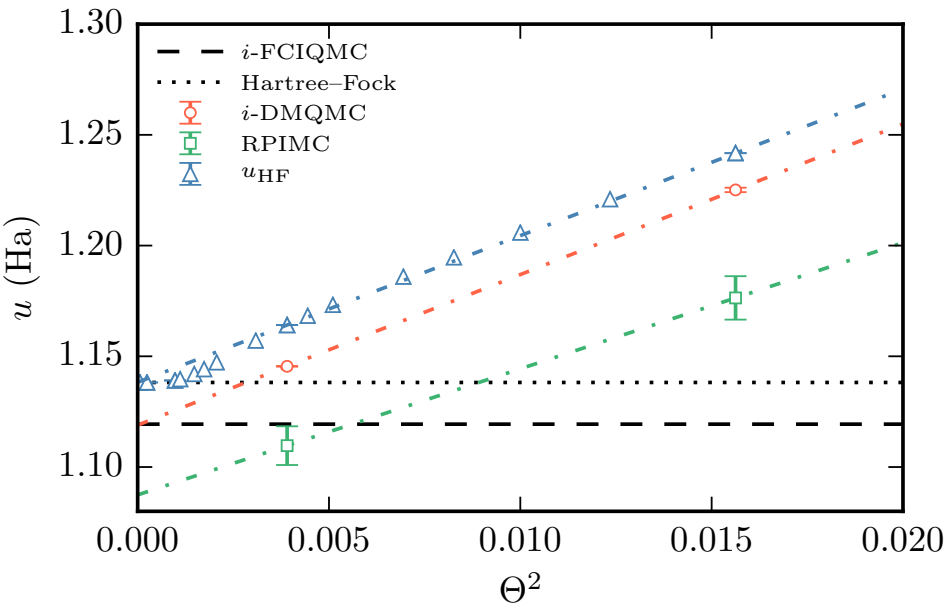

Figure 4.15: Extrapolation of the internal energy to $\Theta = 0$ for the $N = 33$, $\zeta = 1$, $r_s = 1$ system. The RPIMC energies are systematically too low and extrapolate to a value considerably below the i-FCIQMC ground-state energy. The i-DMQMC (and by extension CPIMC) results fare significantly better. This discrepancy cannot be explained by finite-size effects alone. Although two data points is not enough to perform a reliable extrapolation, we see from the Hartree–Fock extrapolation (u_{HF} see Eq. (3.88)) which is also carried out using only the $\Theta = 0.125$ and $\Theta = 0.0625$ points, that a linear fit between these is accurate. Reproduced with permission from Ref. [54].

.

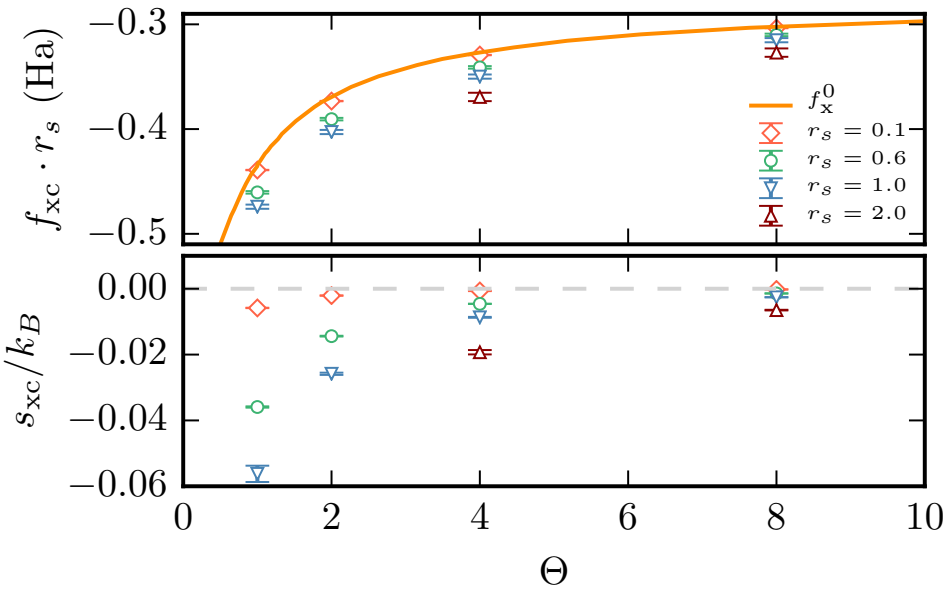

Figure 4.16: Top panel: Exchange-correlation free energies for the $N = 33, \zeta = 1$ electron gas calculated using DMQMC. The sign problem prohibits calculations below $\Theta = 4$ at $r_s = 2$. Also plotted is $f_{\rm x}^0 = N^{-1}\langle \hat{V} \rangle_0$, the first-order exchange contribution to the free energy evaluated in the canonical ensemble. Bottom panel: The exchange-correlation entropy for the same system. Reproduced with permission from Ref. [54].

4.4 Summary

In this chapter we adapted the initiator approximation to DMQMC. This allowed us to resolve the significant disagreement between the published RPIMC and CPIMC results for the exchange correlation energy. The initiator approximation dramatically extends the scope of the DMQMC method and is particularly useful at low temperatures where errors in the internal energy converge more rapidly due to the variational principle. Worryingly, the errors due to the initiator approximation affect expectation values of operators that do not commute with the Hamiltonian much more severely than they affect the energy itself and converge very slowly with walker number. This makes accurate initiator results for the free energy difficult to obtain.

Chapter 5

The Thermodynamic Limit

In the previous chapters we showed how DMQMC could be adapted to treat the UEG and how the initiator approximation could significantly extend the system sizes we could tackle using it. These developments allowed us to attain highly accurate benchmark results for finite UEG systems which are important in their own right, particularly in the development of approximate methods. However, for the purposes of developing a DFT functional, we really need results in the thermodynamic limit and thus need to perform some extrapolation with electron number. In this chapter we will investigate to what extent we can reliably approach the thermodynamic limit using DMQMC. In particular, we will focus on the use of analytic finite size corrections in QMC simulations and suggest an improved scheme of particular use in the high temperature and high density regimes, comparing where possible to existing theories[1].

5.1 Finite Size effects

5.1.1 Single-particle size effects

All of the QMC calculations presented in this work were carried out in a cubic simulation cell of length L with N particles subject to periodic boundary conditions. The use of periodic boundary conditions serves to mimic an infinite solid but leads to two sources of finite size errors. The first is that the allowed k-points that define our single particle basis are spaced evenly (for a cubic box) at intervals of $\Delta k = 2\pi/L$. This means that expectation values, which would otherwise be evaluated as integrals over $\mathbf{k}$ in the thermodynamic limit, are replaced by discrete sums. This leads to an integration error similar to that found in ground state mean field electronic structure methods when integrating over the Brillouin zone. Numerous methods exist to aid in the removal of this error, including the use of special k-points [159] for insulators or cleverly chosen grids of k-points [160] that allow accurate results to be obtained with significantly coarser k-point meshes.

[1] Parts of this work (particularly concerning the analytic corrections to the potential energy) were developed independently by Tobias Dornheim and Simon Groth. Our results were published jointly in Ref. [55].

In many-body simulations, the density of k-points is determined entirely by the number of electrons in the simulation cell. By increasing the number of electrons (and keeping the electron density fixed) we can increase the k-point density and thus reduce these integration errors. Some issues with this approach are that we are usually restricted to simulating a handful of electrons, and also that the QMC energy can display significant 'shell'-effects where the energy exhibits jumps as the electron number changes, rendering a direct extrapolation difficult [161]. In the case of the UEG one can choose so-called 'magic' numbers of electrons so that a complete shell of k-points is occupied at $T = 0$ (for the corresponding ideal system). For example, $N = 7, 19, 33, \ldots$, all represent magic numbers for the spin polarised 3D UEG. The use of closed shell systems generally produces total energies closer to the thermodynamic limit value, but the issue of extrapolating the residual error remains [161]. A more general way to reduce these shell effects is instead to perform twist averaging [161].

Recall that the many-electron Hamiltonian is invariant under the translation of any electron by a simulation cell translation vector $\mathbf{R}_s$. This implies that translating any single electron by $\mathbf{R}_s$ can only change the many-electron wavefunction by a phase factor:

$$\Psi(\mathbf{r}_1, \mathbf{r}_2, \cdots, \mathbf{r}_i + \mathbf{R}_s, \cdots, \mathbf{r}_N) = e^{i\mathbf{k}_s \cdot \mathbf{R}_s} \Psi(\mathbf{r}_1, \mathbf{r}_2, \cdots, \mathbf{r}_i, \cdots, \mathbf{r}_N), \tag{5.1}$$

where $\mathbf{k}_s$ can take any value in the simulation cell Brillouin zone, and $\mathbf{k}_s = \mathbf{0}$ corresponds to the usual case of periodic boundary conditions. Twist averaging exploits this symmetry and averages results over different values of $\mathbf{k}_s$ while keeping the number of electrons fixed[2]. For metallic systems this helps resolve the sharp discontinuity in the electronic occupations at the Fermi energy. For DMQMC, twist averaging amounts to evaluating

$$\langle \hat{O} \rangle_{\mathbf{k}_s} = \frac{1}{(2\pi)^3} \int_{\mathrm{BZ}} d\mathbf{k}_s \; Z(\mathbf{k}_s)^{-1} \mathrm{Tr}\left[\hat{\rho}(\mathbf{k}_s)\hat{O}\right]. \tag{5.2}$$

Eq. (5.2) can be evaluated approximately using Monte Carlo integration where twist vectors are generated randomly so that we can estimate the integral as

$$\langle \hat{O} \rangle_{\mathbf{k}_s} \approx \bar{O} = \frac{1}{N_{\mathbf{k}_s}} \sum_{\mathbf{k_s}} \bar{O}(\mathbf{k}_s), \tag{5.3}$$

where $\bar{O}(\mathbf{k}_s)$ is the QMC result for $\langle \hat{O} \rangle$ at a particular $\mathbf{k}_s$. An estimate for the standard error can be found from

$$\sigma^2_{\langle O \rangle_{\mathbf{k}_s}} \approx \sigma^2_{\mathrm{QMC}} + \sigma^2_{\mathrm{Integration}} \tag{5.4}$$

$$\approx \frac{1}{N^2_{\mathbf{k}_s}} \sum_{\mathbf{k_s}} \sigma^2_{\bar{O}(\mathbf{k}_s)} + \frac{1}{N_{\mathbf{k}_s}(N_{\mathbf{k}_s} - 1)} \sum_{\mathbf{k}_s} (\bar{O}(\mathbf{k}_s) - \bar{O})^2, \tag{5.5}$$

where the first term accounts for the statistical errors in the individual QMC estimates

[2] Twist averaging in the grand canonical ensemble is also possible [161] but we will not discuss this here.

at different $\mathbf{k}_s$ (which are statistically independent), and the second term accounts for the error in the Monte Carlo integration. We see, therefore, that twist averaging should not significantly affect the computational load of a DMQMC simulation as we need to perform many independent simulations anyway. This is of course only true if the error from the Monte Carlo integral converges rapidly, which is not always the case [161, 162].

Twist averaging is ubiquitous for $T = 0$ QMC simulations of UEG systems and in real solid-state applications [163, 164] but its application at $T > 0$ has not been as widely reported. The RPIMC simulations of Ref. [11] did not perform twist averaging at all but used magic numbers of $N = 33$ and 66, whilst Schoof *et al.* [12] found that the use of twist averaging did not significantly affect results for $\Theta \gtrapprox 0.5$. This can be understood by noting that at the high temperatures considered here, the electron occupations no longer exhibit a sharp discontinuity at the Fermi energy but are smoothly spread out to higher energy states, as in the case of the usual Fermi factor in a non-interacting picture.

To investigate this further we plot in Fig. 5.1 twist averaged results for the non-interacting kinetic energies calculated in the canonical ensemble relative to their thermodynamic limit values. We look at the ideal case as the shell effects for the non-interacting and interacting system are roughly proportional to one another [162] and it is much less computationally demanding to obtain results for the ideal system. In the thermodynamic limit we can evaluate the kinetic energy per-particle as

$$t_0 = (2-\zeta)\frac{2^{3/2}}{3\pi} r_s^3 \beta^{-5/2} I_{3/2}(\eta_0), \tag{5.6}$$

$$\tag{5.7}$$

where $\eta_0 = \beta\mu_0$,

$$I_\nu(\eta) = \int_0^\infty dx \frac{x^\nu}{1+e^{x-\eta}}, \tag{5.8}$$

is the usual Fermi integral, and we can determine the chemical potential from

$$n = (2-\zeta)\frac{\sqrt{2}}{2\pi^2}\beta^{-3/2} I_{1/2}(\eta_0). \tag{5.9}$$

Note that here, and for remainder of this thesis, the prefactor of $(2-\zeta)$ is understood to be valid for the fully spin polarised and unpolarised systems only. We see in Fig. 5.1 that, as expected [161, 162], at low temperatures the twist averaged result is generally above that of the infinite system size and that the number of twists needed to converge the average is rather large [12, 162]. The number of twists can be reduced by using better choices of twist vectors, but we have not investigated this here. At higher temperatures, we find that the twist averaged kinetic energy is generally lower that the thermodynamic limit result, with the effect of twist averaging diminishing with increasing Θ. Indeed, for

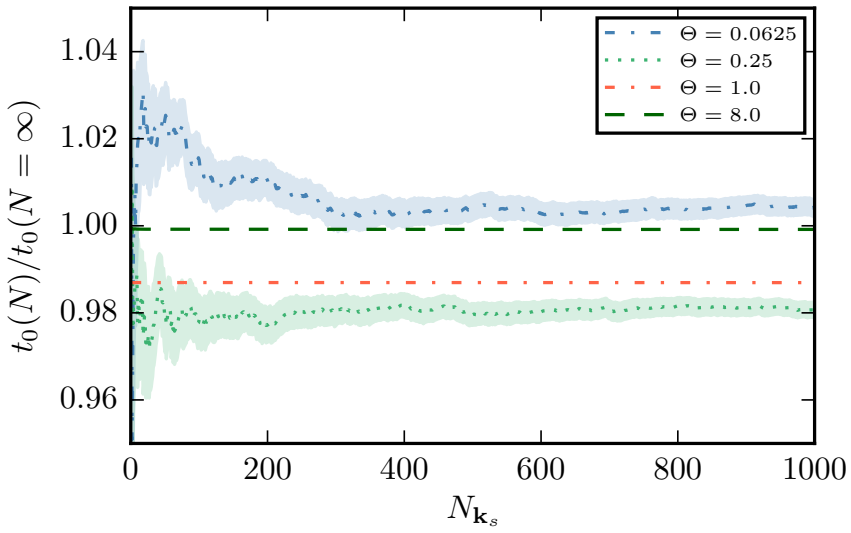

Figure 5.1: Variation of the twist averaged ideal kinetic energy per particle of an $N = 14$, $\zeta = 0$, $r_s = 1$ UEG calculated in the canonical ensemble as a function of the number of twists, $N_{\mathbf{k}_s}$. Energies are relative to the corresponding result in the thermodynamic limit. The shaded region around the lines represents the error bar calculated from Eq. (5.5).

$\Theta \geq 1$, for this system, the twist averaged results are practically indistinguishable from the $\mathbf{k}_s = \mathbf{0}$ result. We also see that residual one-electron effects can remain significant in the canonical ensemble and that some extrapolation is necessary. We find that size corrections to the ideal part of the free energy are typically larger than the internal energy, particularly at higher temperatures [149], and again become less sensitive to twist averaging as the temperature increases.

In Fig. 5.2 we plot the convergence of the twist averaged i-DMQMC internal energy with the number of twists at for a $N = 14$ UEG at $r_s = 1$ and $\Theta = 0.0625$. As was the case for the ideal system, we see that the convergence of the intergral with $N_{\mathbf{k}_s}$ is slow. One way to accelerate the convergence of the twist averaging procedure is to use a control variate [165] and instead average

$$\tilde{U}(\mathbf{k}_s) = U(\mathbf{k}_s) - (T_0(\mathbf{k}_s) - \langle \hat{T} \rangle_{0,\mathbf{k}_s}), \tag{5.10}$$

or

$$\tilde{U}(\mathbf{k}_s) = U(\mathbf{k}_s) - (U_{\mathrm{HF}}(\mathbf{k}_s) - \langle \hat{H} \rangle_{\mathrm{HF},\mathbf{k}_s}), \tag{5.11}$$

where $\langle \hat{T} \rangle_{0,\mathbf{k}_s}$ denotes the twist averaged ideal kinetic energy. Averaging either of these quantities uses the fact that the kinetic energy and QMC energy for a given $\mathbf{k}_s$ are usually highly correlated and that the term in brackets on the right hand side vanishes in the limit of perfect twist averaging. Using this control variate technique we see from Fig. 5.2 that the total energy converges much more smoothly with $N_{\mathbf{k}_s}$ and thus reduces

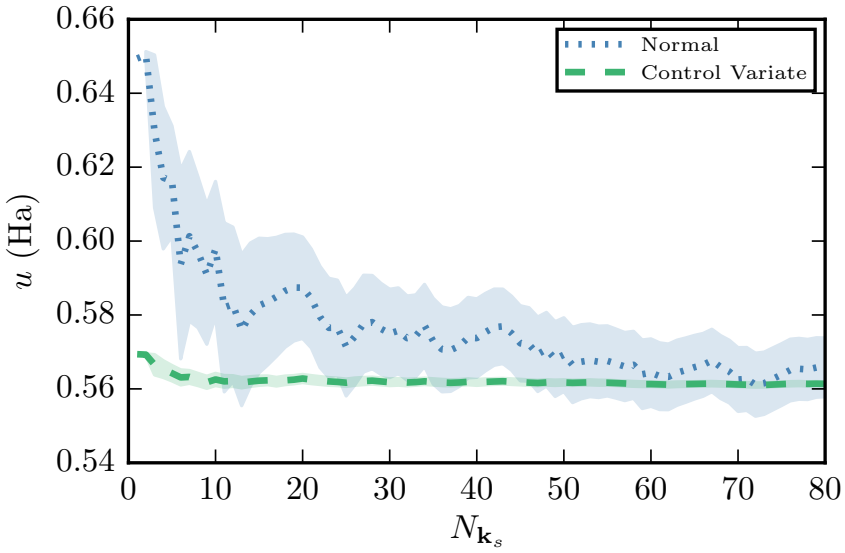

Figure 5.2: Variation of the twist averaged i-DMQMC internal energy per particle of with and without using a control variate. Here we used Eq. (5.10) in calculating the average. The system shown here is an $N = 14$, $\zeta = 0$, $r_s = 1$ UEG with $M = 1045$ and $N_w = 7 \times 10^6$ at $\Theta = 0.0625$. The shaded region around the lines represents the error bar calculated from Eq. (5.5).

the computational overhead of twist averaging, here by about a factor of 30.[3]

5.1.2 Many-particle size effects

The second type of finite size error is fundamentally many-body in nature and does not have an analogue in mean-field methods. Converging these residual size errors, which manifest themselves in the expectation values of both one- and two-body operators, can be slow and usually represents the largest remaining source of systematic error in any QMC simulation of a bulk system. One of the most successful approaches to reducing these sources of error, particularly for the UEG, has been with the development of analytic size corrections usually derived from approximate many-body theories [162, 166, 167].

Consider the exact expression for the finite size correction to the (electron-electron) potential energy calculated in a QMC simulation. This is defined as the difference between the potential energy evaluated in the infinite and the N-particle system, i.e.,

$$\Delta V = \langle \hat{V} \rangle_\infty - \langle \hat{V}_{\mathrm{Ew}} \rangle_N \tag{5.12}$$

$$= V_\infty - V_{\mathrm{QMC}}. \tag{5.13}$$

[3] Another recent development which could further reduce this burden is to use 'special' twist vectors [164]; these use the fact that for certain values of $\tilde{\mathbf{k}}_s$, $t_0^N(\tilde{\mathbf{k}}_s) = t_0^\infty$.

To allow for a clearer analysis of the sources of error in Eq. (5.12), it is helpful to rewrite Eq. (5.12) in terms of the static structure factor. The static structure factor is defined as

$$S(\mathbf{q}) = N^{-1}\langle\hat{\rho}_{-\mathbf{q}}\hat{\rho}_{\mathbf{q}}\rangle, \tag{5.14}$$

where

$$\hat{\rho}_{\mathbf{q}} = \sum_{\mathbf{k}\sigma}\hat{c}^{\dagger}_{\mathbf{k}+\mathbf{q}\sigma}\hat{c}_{\mathbf{k}\sigma}, \tag{5.15}$$

is the density fluctuation operator. Inserting Eq. (5.15) into Eq. (5.14) we can write

$$S(\mathbf{q}) = \sum_{\sigma\sigma'} S_{\sigma\sigma'}(\mathbf{q}), \tag{5.16}$$

where the spin-resolved structure factor can be written as

$$S_{\sigma\sigma'}(\mathbf{q}) = \frac{N_\sigma}{N}\delta_{\sigma\sigma'} + \frac{1}{N}\sum_{\mathbf{kp}}\left\langle\hat{c}^{\dagger}_{\mathbf{k}+\mathbf{q}\sigma}\hat{c}^{\dagger}_{\mathbf{p}-\mathbf{q}\sigma'}\hat{c}_{\mathbf{p}\sigma'}\hat{c}_{\mathbf{k}\sigma}\right\rangle, \tag{5.17}$$

which is quite clearly related to the potential energy. We can now write the size correction as

$$\Delta V = \frac{N}{2(2\pi)^3}\int d\mathbf{q}\; v_{\mathbf{q}}(S_\infty(\mathbf{q})-1) - \left[\frac{N}{2L^3}\sum_{\mathbf{q}\neq\mathbf{0}} v_{\mathbf{q}}(S_N(\mathbf{q})-1) + \frac{1}{2}N\xi\right], \tag{5.18}$$

where $S_\infty(\mathbf{q})$ and $S_N(\mathbf{q})$ are the structure factors for the infinite and finite systems, respectively.

We can identify three contributions to the error in Eq. (5.18):

1. the discretisation errors from the replacement of an integral with a sum;
2. the omission of the $\mathbf{q} = \mathbf{0}$ term in the sum;
3. the fact that $S_N(\mathbf{q}) \neq S_\infty(\mathbf{q})$.

Previous work by Chiesa *et al.* [166] proposed accounting for these errors by first noting that the structure factor converges remarkably fast as N increases, so that the residual errors must come from effects 1 and 2. They further suggested that the dominant remaining contribution comes from point 2. The size correction can then be approximated using the behaviour of the static structure factor in the low $\mathbf{q}$ limit, which in turn is known exactly from the RPA [28]. The generalisation of these ideas for finite temperatures was first attempted by Brown *et al.* [11] who found that the appropriate size correction for finite temperatures is given by

$$\Delta V = \frac{\omega_p}{4}\coth\left(\frac{\beta\omega_p}{2}\right), \tag{5.19}$$

where $\omega_p = \sqrt{3/r_s^3} = \sqrt{4\pi n}$ is the plasma frequency. Subsequently it was found by Schoof [12] that this size correction was very inaccurate for low r_s and also higher temperatures. To address the reasons for this we need to go back to basics and understand the origin of Eq. (5.19).

5.2 Improved Finite Size Corrections for Warm Dense Matter

The static structure factor can be calculated exactly from the fluctuation-dissipation theorem as

$$S(\mathbf{q}) = -\frac{1}{2\pi n}\int_{-\infty}^{\infty} d\omega\ \mathrm{Im}[\chi(\mathbf{q},\omega)]\coth\left(\frac{\beta\omega}{2}\right), \tag{5.20}$$

where $\chi(\mathbf{q},\omega)$ is the density-density response function of the system [28, 168]. To derive the size correction given by Eq. (5.19) we first need to determine the low $\mathbf{q}$ limit of Eq. (5.20) within the RPA, which is exact in this regime. The RPA expression for the susceptibility is [28]

$$\chi^{\mathrm{RPA}}(\mathbf{q},\omega) = \frac{\chi^0(\mathbf{q},\omega)}{1 - v_{\mathbf{q}}\ \chi^0(\mathbf{q},\omega)} \tag{5.21}$$

$$= \frac{\chi^0(\mathbf{q},\omega)}{\varepsilon^{\mathrm{RPA}}(\mathbf{q},\omega)}, \tag{5.22}$$

where

$$\chi^0(\mathbf{q},\omega) = \frac{1}{L^3}\sum_{\mathbf{k}\sigma}\frac{f_{\mathbf{k}\sigma} - f_{\mathbf{k}+\mathbf{q}\sigma}}{\omega + \varepsilon_{\mathbf{k}\sigma} - \varepsilon_{\mathbf{k}+\mathbf{q}\sigma}}, \tag{5.23}$$

is the usual Lindhard function and here ω is a complex number. One can show by analysing the behaviour of Eq. (5.22) [28], that $\mathrm{Im}\left[\chi^{\mathrm{RPA}}(\mathbf{q},\omega+i\eta)\right]$ is dominated by the plasmon-pole contribution below some critical wavevector $q_c(r_s,\Theta)$ determined when

$$\varepsilon^{\mathrm{RPA}}(\mathbf{q},\omega+i\eta) = 0, \tag{5.24}$$

for η an infinitesimally small positive constant. This simplifies matters considerably as we can replace the imaginary part of the RPA dielectric function by a delta function centred on the (q-dependent) plasma frequency $\Omega_p(\mathbf{q})$. To find the plasma frequency we can solve Eq. (5.24) for $\Omega_p(\mathbf{q})$ and then determine the plasmon oscillator strength (the weight of the delta function) as outlined in Ref. [28]. The result is (for $q \leq q_c$)

$$-\frac{1}{\pi}\mathrm{Im}[\chi^{\mathrm{RPA}}(\mathbf{q},\omega+i\eta)] \approx \frac{\Omega_p(\mathbf{q})}{2v_{\mathbf{q}}}\delta(\omega - \Omega_p(\mathbf{q})), \tag{5.25}$$

and we have the property that $\lim_{\mathbf{q}\to\mathbf{0}}\Omega_p(\mathbf{q}) = \omega_p$ [28]. Inserting Eq. (5.25) into Eq. (5.20) we find

$$\lim_{q\to 0} S_{\mathrm{RPA}}(q) \approx \frac{q^2}{2\omega_p}\coth\left(\frac{\beta\omega_p}{2}\right). \tag{5.26}$$

With the low q behaviour of the static structure factor at hand we can now derive the required size correction. We first replace S_∞ and S_N by $S_\infty^{\rm RPA}$ in Eq. (5.18):

$$\Delta V \approx \frac{N}{2(2\pi)^3}\int d\mathbf{q}\; v_{\mathbf{q}}(S_\infty^{\rm RPA}(\mathbf{q})-1)e^{-\delta q^2} - \frac{N}{2L^3}\sum_{\mathbf{q}\neq\mathbf{0}} v_{\mathbf{q}}(S_\infty^{\rm RPA}(\mathbf{q})-1)e^{-\delta q^2}, \quad (5.27)$$

where we can introduce the convergence factors of $e^{-\delta q^2}$ because both the sum and the integral in Eq. (5.18) converge [162]. It is convenient to further write the size corrections as $\Delta V = \Delta V_1 + \Delta V_2$ where

$$\Delta V_1 = \frac{N}{2(2\pi)^3}\int d\mathbf{q}\; v_{\mathbf{q}} S_\infty^{\rm RPA}(\mathbf{q})e^{-\delta q^2} - \frac{N}{2L^3}\sum_{\mathbf{q}\neq\mathbf{0}} v_{\mathbf{q}} S_\infty^{\rm RPA}(\mathbf{q})e^{-\delta q^2} \quad (5.28)$$

and

$$\Delta V_2 = -\frac{N}{2(2\pi)^3}\int d\mathbf{q}\; v_{\mathbf{q}} e^{-\delta q^2} + \frac{N}{2L^3}\sum_{\mathbf{q}\neq\mathbf{0}} v_{\mathbf{q}} e^{-\delta q^2}. \quad (5.29)$$

We will now show that ΔV_1 gives the desired size correction while ΔV_2 is related to the Madelung constant.

To work out the difference between the sum and the integral in Eq. (5.28) we use the Poisson summation formula [166]

$$\sum_{\mathbf{R}_s} f(\mathbf{R}_s) = \frac{1}{L^3}\sum_{\mathbf{k}} \tilde{f}(\mathbf{k}), \quad (5.30)$$

$$= \frac{1}{L^3}\tilde{f}(\mathbf{0}) + \frac{1}{L^3}\sum_{\mathbf{k}\neq\mathbf{0}} \tilde{f}(\mathbf{k}), \quad (5.31)$$

where $f(\mathbf{r})$ is a smooth and rapidly decaying function of $\mathbf{r}$, and

$$\tilde{f}(\mathbf{k}) = \int d\mathbf{r}\; e^{i\mathbf{k}\cdot\mathbf{r}} f(\mathbf{r}), \quad (5.32)$$

is its Fourier transform. Setting $\tilde{f}(\mathbf{q}) = v_{\mathbf{q}} S_\infty^{\rm RPA}(\mathbf{q})e^{-\delta q^2}$ we can write

$$\Delta V_1 = \frac{N}{2} f(\mathbf{R}_s = \mathbf{0}) - \frac{N}{2}\left(-\frac{1}{L^3}\lim_{\mathbf{q}\to\mathbf{0}}\tilde{f}(\mathbf{q}) + \sum_{\mathbf{R}_s} f(\mathbf{R}_s)\right) \quad (5.33)$$

$$= \frac{N}{2}\left(\frac{1}{L^3}\lim_{\mathbf{q}\to\mathbf{0}}\tilde{f}(\mathbf{q}) - \sum_{\mathbf{R}_s\neq\mathbf{0}} f(\mathbf{R}_s)\right) \quad (5.34)$$

where we have used

$$f(\mathbf{R}_s) = \frac{1}{(2\pi)^3}\int d\mathbf{q}\; e^{-i\mathbf{q}\cdot\mathbf{R}_s}\tilde{f}(\mathbf{q}). \quad (5.35)$$

For ΔV_2 we have

$$\Delta V_2 = \frac{N}{2L^3}\sum_{\mathbf{q}\neq\mathbf{0}} v_{\mathbf{q}} e^{-\delta q^2} - \frac{N}{(2\pi)^3}\int d\mathbf{q}\; v_{\mathbf{q}} e^{-\delta q^2} \tag{5.36}$$

$$= \frac{N}{2}\left(\frac{1}{L^3}\sum_{\mathbf{q}\neq\mathbf{0}} v_{\mathbf{q}} e^{-\delta q^2} - \frac{1}{\sqrt{\pi\delta}}\right) \tag{5.37}$$

$$\approx \frac{N}{2}\xi, \tag{5.38}$$

where we have used the definition of the Madelung constant from **??**, set $\kappa = 1/(2\sqrt{\delta})$ and assumed δ is small.

If we now incorporate ΔV_2 into our definition of V_{QMC}, insert Eq. (5.26) into Eq. (5.34), drop the convergence factors and assume that the real space term in Eq. (5.34) is negligible [162], we arrive at

$$\Delta V = \frac{N}{2L^3}\lim_{\mathbf{q}\to\mathbf{0}} v_{\mathbf{q}} S_\infty^{\mathrm{RPA}}(\mathbf{q}) \tag{5.39}$$

$$= \frac{\omega_p}{4}\coth\left(\frac{\beta\omega_p}{2}\right) \tag{5.40}$$

as desired.

In deriving these corrections we have assumed that we are in the low q limit ($q \leq q_c$), limit so that we can replace the QMC structure factor by that found from the RPA. However, this may not be the case in a QMC simulation where the minimum q is fixed by

$$q_{\mathrm{min}} = \frac{2\pi}{L} \tag{5.41}$$

$$= \frac{2\pi}{\left(\frac{4}{3}\pi N\right)^{1/3} r_s}. \tag{5.42}$$

Generally speaking, $q_c(r_s, \Theta)$ decreases with both increasing Θ and decreasing r_s (increasing density) [169] as shown in Fig. 5.3. We see that for typical simulation sizes available to us, the values of q_c are much smaller than q_{min}, suggesting that the response function is not exhausted by the plasmon-pole and, therefore, that the low-odered correction is not appropriate.

To demonstrate this point further, we calculate the structure factor in DMQMC using Eqs. (5.16) and (5.17), which we plot in Fig. 5.4 for $N = 14$, 38, 54, 66 and $\zeta = 0$. As expected, we find that the structure factor converges remarkably fast with N. Furthermore, we see that the low q behaviour of $S(q)$ is not well approximated by Eq. (5.26), shown by the dotted line. Size corrections in the form of Eq. (5.19), which used this form of $S(q)$, can therefore lead to corrections which are too large. In Fig. 5.5 we show this directly by plotting the size-corrected potential energy per particle as a

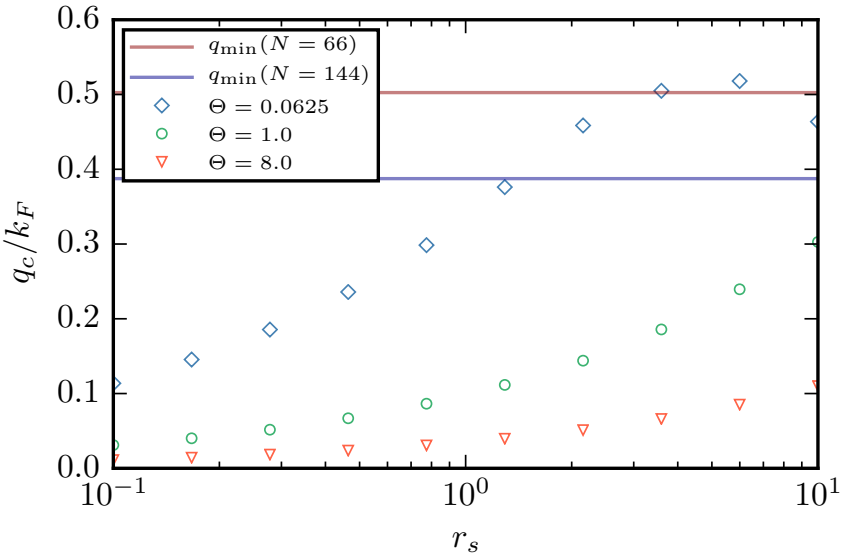

Figure 5.3: Rough determination of $q_c(r_s, \Theta)$ for an unpolarised UEG within the RPA. Here q_c is determined by the maximum q for which $\mathrm{Re}\left[\varepsilon^{\mathrm{RPA}}(q, \omega + i\eta)\right] = 0$ and $\mathrm{Im}\left[\varepsilon^{\mathrm{RPA}}(q, \omega + i\eta)\right] = 0$. Mild numerical issues coupled with the only approximate form of this solution are the cause of the noise at higher r_s. The horizontal lines represent values of $q_{\min} = 2\left(\frac{\pi}{3N}\right)^{1/3} k_F$. The code to evaluate this is available at Ref. [170].

function of N. It is clear that the size corrections to v are very large and that the low-order corrections do not help convergence much. Indeed, the low-order corrections can render the potential energy positive for these values of N and Θ [55]. It should be noted that this problem was understood in Brown's RPIMC work [11, 59], where a classical size correction was applied for the highest temperature and density points. The RPA procedure presented here is more general, as it incorporates both the degenerate and non-degenerate limits.

A better size correction for v can be found directly from the RPA by going beyond the small q expression for the structure factor, Eq. (5.26), and calculating the energy difference in Eq. (5.18) using the full RPA structure factor for both $S_N(\mathbf{q})$ and $S_\infty(\mathbf{q})$. This is again motivated by the fact that S_N converges rapidly with N and that the size correction vanishes in the limit of $N \to \infty$. One can attempt to go further and use local field corrections to the density response function by setting

$$\chi(\mathbf{q}, \omega) = \frac{\chi^0(\mathbf{q}, \omega)}{1 - (1 - G(\mathbf{q}, \omega))v_{\mathbf{q}}\chi^0(\mathbf{q}, \omega)}. \tag{5.43}$$

It turns out that a particularly good choice is to use the STLS scheme [73], which matches more accurately with the QMC structure factor than the full RPA form [55]. In practice, however, either scheme yields satisfactory size corrections [55] and the RPA is easier to implement.

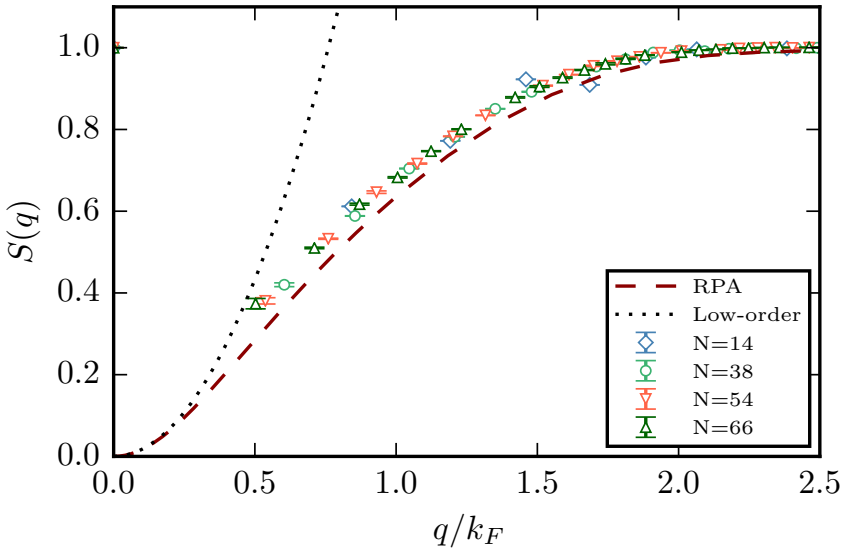

Figure 5.4: Convergence of the static structure factor with particle number calculated using i-DMQMC. The systems shown here is has $r_s = 0.5$, $\Theta = 0.25$ and $\zeta = 0$. Also plotted is the RPA structure factor in the thermodynamic limit using Eq. (5.44) and the low $\mathbf{q}$ form calculated using Eq. (5.26)

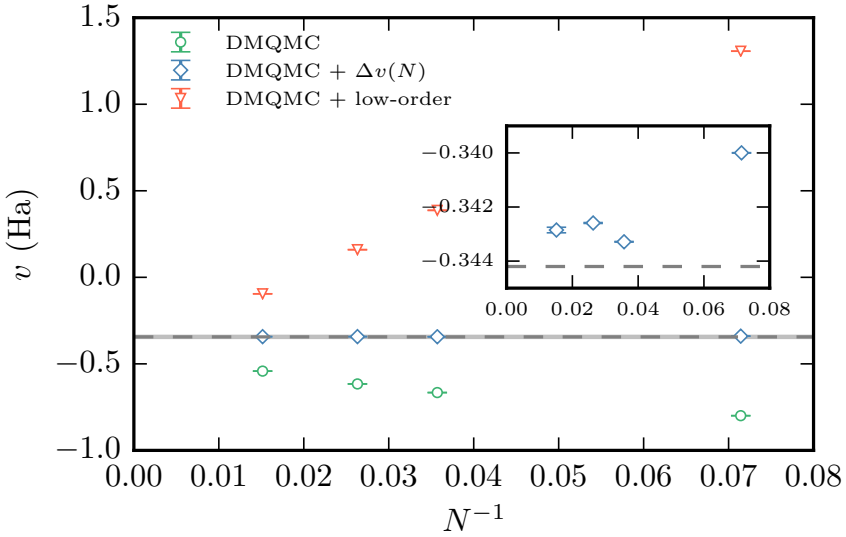

Figure 5.5: The performance of size corrections for the potential energy applied to DMQMC data for $r_s = 0.5, \Theta = 8$. Here $\Delta v(N)$ is the correction found from Eq. (5.18) using the full RPA structure factor, while the low-order correction is that of Eq. (5.19). The horizontal dashed line represents the thermodynamic limit result calculated in Ref. [55] using the same procedure applied to CPIMC data. The shaded region around the dashed line represents a 1% error; the full-RPA corrected DMQMC data fall within this band for as few as 14 electrons. The inset shows a close up of this region showing the residual a $1/N$ error, which generally requires higher electron numbers to resolve fully [55].

In evaluating the structure factor from Eq. (5.20) it is convenient to perform the frequency integral along the imaginary frequency axis [73], so that we can write

$$S(\mathbf{q}) = -\frac{1}{n\beta}\sum_{l\in\mathbb{Z}}\chi(\mathbf{q}, i\omega_l) \tag{5.44}$$

where $i\omega_l = 2i\pi l/\beta$ are the Bosonic Matsubara frequencies (the poles of $\coth\frac{\beta\omega}{2}$). Working with reduced units of $x = q/k_F, y = q/k_F$, χ^0 can be evaluated (in the thermodynamic limit) as [73]

$$\chi^0(x, i\omega_l) = -\frac{3n}{4E_F x}\int_0^\infty dy \frac{y}{e^{y^2/\Theta-\eta_0}+1}\log\left|\frac{(2\pi l\Theta)^2+(x^2+2xy)^2}{(2\pi l\Theta)^2+(x^2-2xy)^2}\right|. \tag{5.45}$$

Eqs. (5.22) and (5.45) allow for a simple evaluation of Eq. (5.44), code for which can be found at Ref. [170]. Convergence of the summation and integrals need to be carefully checked with respect to $l_{\max}$ and the upper limit, $q_{\max}$, of the momentum integral, a task aided by the use of various low-order forms for Eq. (5.45) [73]. We benchmarked our results against those found in Ref. [73] for various thermodynamic quantities. For the size corrections a useful check is to ensure that the full RPA size corrections agree with the low-order form for high r_s and low Θ.

Going back to Fig. 5.5 we demonstrate the improved convergence of the potential energy with N when this new size correction is used. Any remaining finite-size errors are too small to be visible on the scale of the figure. The inset shows the residual bias; to remove this fully one needs to go to larger electron numbers so that a reliable extrapolation can be performed [55].

5.2.1 Size corrections to the free energy

To parametrise a DFT functional one needs data for $f_{\rm xc}$ across the whole (r_s, Θ) plane. There are at least three ways to achieve this:

1. evaluate $f_{\rm xc}$ directly;
2. evaluate $f_{\rm xc}$ via thermodynamic integration of v [51]:

$$f_{\rm xc} = \frac{1}{r_s^2}\int_0^{r_s} dr_s'\; r_s' v(r_s', \Theta); \tag{5.46}$$

3. evaluate $f_{\rm xc}$ via $u_{\rm xc}$ [51]:

$$f_{\rm xc}(r_s, \Theta) - \Theta\left(\frac{\partial f_{\rm xc}(r_s, \Theta)}{\partial\Theta}\right)_{r_s} = u_{\rm xc}(r_s, \Theta). \tag{5.47}$$

Option 3 was deemed the best route for the original $f_{\rm xc}$ fit of Karasiev *et al.* [51] when fitting to the RPIMC data [11]. It has the advantage that the internal energy is simpler

to calculate using most QMC methods but the disadvantage that size corrections for the kinetic energy cannot be simply expressed in terms of a rapidly converging and easily calculable function such as the structure factor[4].

Option 2, on the other hand, benefits from two advantages. The statistical errors in v are smaller compared to u_{xc}; and we now know how to obtain reliable and accurate size corrections for v. Given the limitations of the various existing methods in terms of regions of validity, option 2 has the additional advantage that all methods can evaluate v (to within uncontrolled biases), which is not the case for f_{xc}. The downside is that performing the thermodynamic integration potentially introduces additional biases due to integration errors or by fitting to a functional form of unknown quality. Note that this is also a problem for option 3.

Option 1 is in some sense the ideal and most general route given that the Helmholtz free energy f is the fundamental quantity in canonical statistical mechanics and that all other properties of the system can be obtained from it using thermodynamic relations. Also, the entire phase diagram will not always be required or calculable, so a method capable of providing thermodynamic limit data for f_{xc} is important as more complex systems are studied. We can conclude that for the purpose of parametrising f_{xc} for DFT functionals option 2 is perhaps the best route, but that size corrections to f_{xc} bear some attention.

Size corrections to f_{xc} can in principle be found using thermodynamic integration using the identity [14, 51] (see also Section 3.2.4)

$$F_{\mathrm{xc}} = \int_0^1 d\lambda\ \lambda^{-1}\langle\lambda\hat{V}\rangle_\lambda \tag{5.48}$$

$$= \frac{1}{r_s^2}\int_0^{r_s} dr_s'\ r_s' V(r_s', \Theta), \tag{5.49}$$

so that

$$\Delta F_{\mathrm{xc}} = \frac{1}{r_s^2}\int_0^{r_s} dr_s'\ r_s' \Delta V(r_s', \Theta), \tag{5.50}$$

which has the advantage that ΔV is readily available. A computationally simpler approach is to instead analytically carry out the coupling constant integral which removes the need to perform the integral in Eq. (5.50) approximately. Using Eq. (5.48) and the relationship between V and the structure factor, we can write [14, 171]

$$F_{\mathrm{xc}} = \frac{N}{2(2\pi)^3}\int_0^1 d\lambda\ \lambda^{-1}\int d\mathbf{q}\ \lambda v_{\mathbf{q}}\left[-\frac{1}{n\beta}\sum_l \frac{\chi^0(\mathbf{q}, i\omega)}{1-\lambda v_{\mathbf{q}}\chi^0(\mathbf{q}, i\omega)} - 1\right], \tag{5.51}$$

$$= \frac{N}{2(2\pi)^3}\int d\mathbf{q}\left[\frac{1}{n\beta}\sum_l \log\left(1 - v_{\mathbf{q}}\chi^0(\mathbf{q}, i\omega_l)\right) - v_{\mathbf{q}}\right], \tag{5.52}$$

[4] This was first pointed out by Simon Groth.

which can be further broken down into an exchange component

$$F_{\rm x} = \frac{N}{2(2\pi)^3}\int d\mathbf{q} v_{\mathbf{q}} \left[-\frac{1}{n\beta}\sum_l \chi^0(\mathbf{q}, i\omega_l) - 1 \right] \tag{5.53}$$

$$= \frac{N}{2(2\pi)^3}\int d\mathbf{q} v_{\mathbf{q}} (S_{\rm HF}(\mathbf{q}) - 1), \tag{5.54}$$

where

$$S_{\rm HF}(\mathbf{q}) = 1 - (2-\zeta)\frac{V}{N(2\pi^3)}\int d\mathbf{k} f_{\mathbf{k}+\mathbf{q}} f_{\mathbf{k}}, \tag{5.55}$$

and a correlation component

$$F_{\rm c} = -\frac{V}{2(2\pi)^3\beta}\sum_l \int d\mathbf{q} \left[\log\left(1 - v_{\mathbf{q}}\chi^0(\mathbf{q}, i\omega_l)\right) + v_{\mathbf{q}}\chi^0(\mathbf{q}, i\omega_l) \right]. \tag{5.56}$$

Thus, we can write the size corrections to the exchange-correlation free energy as[5]

$$\Delta F_{\rm xc} = (F_{\rm x}^{\infty} - F_{\rm x}^{N}) + (F_{\rm c}^{\infty} - F_{\rm c}^{N}), \tag{5.57}$$

where we have

$$F_{\rm x}^{N} = \frac{N}{2V}\sum_{\mathbf{q}\neq\mathbf{0}} v_{\mathbf{q}}(S_{\rm HF}(\mathbf{q}) - 1) + \frac{N}{2}\xi, \tag{5.58}$$

and

$$F_{\rm c}^{N} = -\frac{1}{2\beta}\sum_l \sum_{\mathbf{q}\neq\mathbf{0}} \left[\log\left(1 - v_{\mathbf{q}}\chi^0(\mathbf{q}, i\omega_l)\right) + v_{\mathbf{q}}\chi^0(\mathbf{q}, i\omega_l) \right], \tag{5.59}$$

where both $S_{\rm HF}(\mathbf{q})$ and $\chi^0(\mathbf{q}, i\omega_l)$ are again calculated using the infinite system result. The assumptions here are that the structure factor has converged with respect to N so that the only differences between the finite and infinite system size free energies are due to the replacement of the sum with an integral and the omission of the $\mathbf{q} = \mathbf{0}$ term. In Fig. 5.6 we show the convergence of the size correction with the momentum space cutoff $q_{\rm max}$, showing that the corrections converge rapidly and can be made quite accurate.

There is an additional complication here as, for a fixed value of N, the rapid convergence of the structure factor depends on the value of r_s, with $S_N(\mathbf{q})$ converging more rapidly to $S_\infty(\mathbf{q})$ for larger values of r_s. This raises the question whether, the assumption allowing us to perform the coupling constant integral with $S_\infty(\mathbf{q})$ is justified[6]. The other slight concern is that we are using expressions for the exchange and correlation energy for the finite system in the grand canonical ensemble, which may lead to further issues (although this has implicitly been the case from the beginning with the potential energy corrections). Nevertheless, since the size corrections vanish in the $N \to \infty$ limit we can test them out. Indeed, as shown in Fig. 5.7, we see that the corrections perform

[5] We have found that this expression converges more rapidly with l and $\mathbf{q}$ than that found using the combined form of Eq. (5.52), although both yield identical results.

[6] This was pointed out to me by Simon Groth.

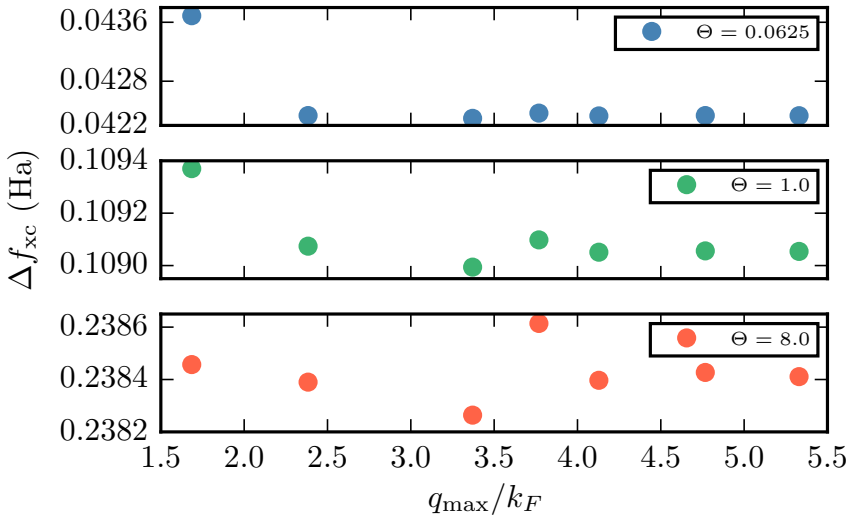

Figure 5.6: Convergence of the size corrections calculated using Eq. (5.57) with q_{max}. The corrections shown here are for a $N = 14$, $\zeta = 0$ and $r_s = 1$ UEG. We see that the corrections converge quite rapidly with q_{max}, particularly at higher temperatures, and also grow in magnitude with increasing temperature. Here we set $l_{\mathrm{max}} = 10000$ for both the sum and the integral.

remarkably well, bringing us to within 1% accuracy of the $N = \infty$ result and seem to reveal a convincing residual N^{-1} behaviour. This suggests that our approach and the functional form used to perform the thermodynamic integration of the size corrected potential energy in Ref. [55] are consistent.

5.2.2 Arbitrary size corrections

The utility of finding corrections for f_{xc} can be appreciated when deriving size corrections for other quantities. We can demonstrate this for the exchange-correlation energy u_{xc} using [51]

$$f_{\mathrm{xc}}(r_s, \Theta) - \Theta \left(\frac{\partial f_{\mathrm{xc}}(r_s, \Theta)}{\partial \Theta} \right)_{r_s} = u_{\mathrm{xc}}(r_s, \Theta), \tag{5.60}$$

where we have used $f_{\mathrm{xc}} = u_{\mathrm{xc}} - T s_{\mathrm{xc}}$, and

$$S = -\left(\frac{\partial F}{\partial T} \right)_{N,V}. \tag{5.61}$$

The size correction Δu_{xc} is then found by replacing f_{xc} by Δf_{xc} in Eq. (5.60) above. The derivative term in Eq. (5.60) can be evaluated using finite differences, a temperature step of $\Delta\Theta = 0.001$ producing sufficiently accurate results.

The performance of these corrections is shown in Fig. 5.8. Again we see the large finite size effects present in the raw QMC data and also the good performance of the analytic corrections. Ref. [55] gives accurate values of f_{xc} in the thermodynamic limit

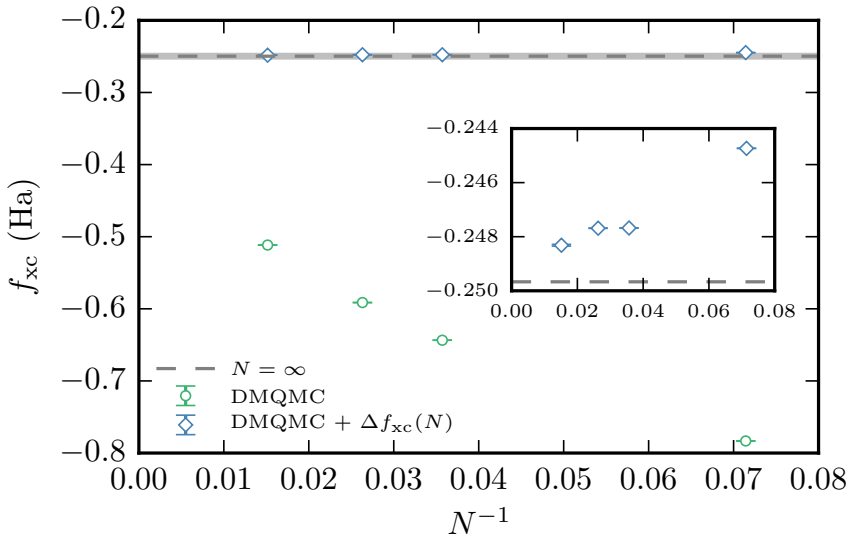

Figure 5.7: Performance of new size corrections outlined in Eq. (5.57) applied to DMQMC data for $r_s = 0.5, \Theta = 8$ as in Fig. 5.5. The horizontal dashed line represents the thermodynamic limit result calculated using a fitted form to the thermodynamic limit v as outlined in Ref. [55]. The shaded region around the dashed line represents a 1% error. The inset shows a close up of this region where we see the residual $1/N$ error, which generally requires higher electron number to resolve fully.

for a finite set of points only. We therefore compare our results to values from the KSDT fit [51] calculated using Eq. (5.60) and the STLS result [73]. The large errors in the KSDT result of about 10% are consistent with similar deviations found in Ref. [55]. In the next section we will investigate the extent of these deviations in more detail.

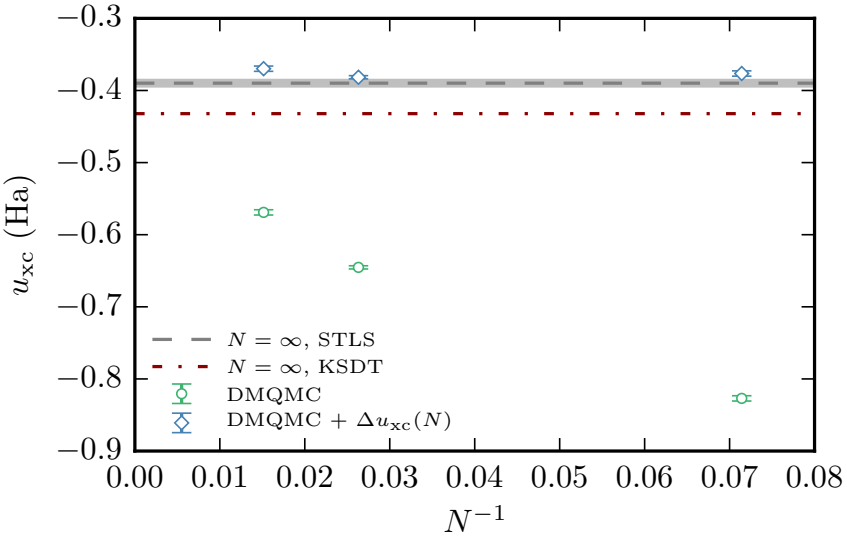

Figure 5.8: Performance of new size corrections outlined in Eq. (5.60) applied to DMQMC data for $r_s = 0.5, \Theta = 8$ and $\zeta = 0$ as in Fig. 5.5. Here, the horizontal dashed lines is the STLS approximation from Ref. [73] and the dot-dashed line is the KSDT result [51]. The shaded region around the dashed line represents a 1% error. No twist averaging was performed.

5.3 Comparison to other methods

The preliminary results from the previous section give us some confidence that we can say something meaningful about the thermodynamic limit of the warm dense electron gas. We mainly focus on relatively high densities and temperatures so that we do not have to use the initiator approximation. We ran DMQMC simulations for $N = 14$, 38 and 66 electrons, all of which are magic numbers, for $\zeta = 0$, $\Theta \geq 1$ and $r_s = 0.5$ and 1. We used the basis set extrapolation techniques discussed previously and also used importance sampling [133]. The error bars in the figures account for any remaining systematic basis set errors. For $\Theta \leq 1$ we also ran i-DMQMC simulations for $N = 14$ electrons in a relatively small basis set of $M = 389$ plane waves using $N_w \approx 10^7$ walkers and applied the basis set corrections for $\Theta \geq 0.25$. The residual basis set errors are estimated to be a few milli-Hartrees per particle and should be insignificant on the scales of the figures and compared to the size corrections. We also performed twist averaging for u_{xc} using the control variate method outlined in Section 5.1.1 for $\Theta \leq 0.125$. Here we used $N_w = 7 \times 10^6$ walkers, $M = 1045$ plane waves and averaged over roughly 100 twist vectors.

A comparison of these results for the potential energy, exchange-correlation free and internal energies for $r_s = 0.5, 1$ are shown in Figs. 5.9 to 5.11. Note that here the

exchange contribution to the internal energy is [14, 169, 172][7]

$$u_{\mathrm{x}} = \frac{3}{2}\mu_{\mathrm{x}} - f_{\mathrm{x}}, \tag{5.62}$$

where

$$\mu_{\mathrm{x}} = -\left(\frac{\partial\Omega_{\mathrm{x}}}{\partial\mu}\right)_{\mu=\mu_0} \Big/ \left(\frac{\partial^2\Omega_0}{\partial\mu^2}\right)_{\mu=\mu_0}, \tag{5.63}$$

$$= -\frac{1}{\sqrt{2}\pi}\beta^{-1/2} I_{-1/2}(\eta_0), \tag{5.64}$$

where $I_{-1/2}(\eta_0)$ is defined as in Eq. (5.8), and that the exchange contribution to the free energy is most conveniently calculated as [173]

$$F_{\mathrm{x}} = -(2-\zeta)\frac{V}{4\pi^3\beta^2}\int_{-\infty}^{\eta_0} d\xi\; I^2_{-1/2}(\xi). \tag{5.65}$$

The first thing to notice is the overall excellent agreement between the DMQMC data and the STLS scheme above $\Theta = 1$. We see from the clustering of points for different particle numbers that the size corrected data, while still containing some residual bias, are very nearly sufficiently converged. In most cases (on this scale) we cannot distinguish between our results and the extrapolated CPIMC results from Ref. [55], which typically used much larger particle numbers and extrapolated the residual particle number bias. We also see the importance of including correlation effects beyond exchange in the description of warm dense matter even in this high density regime.

In agreement with Ref. [55], we find that the KSDT fit for the unpolarised gas can exhibit rather large deviations (up to 10%) from our data at higher temperatures; this can be attributed to the inaccurate RPIMC results and size corrections in the high density regime. Towards lower temperatures the KSDT fit seems to perform better, which can be understood in part by a fortuitous cancellation of the fixed node errors and errors in the finite size corrections and, it must be said, in the use of a good functional form for f_{xc} which also incorporates accurate $T = 0$ data.

For $\Theta \leq 1$, where the initiator approximation was used, we see that, as expected, the i-DMQMC results for v and f_{xc} are significantly biased relative to the CPIMC results. However, as we have seen previously, initiator errors for the internal energy converge more rapidly and thus our results for u_{xc} should be accurate. We see that that the STLS approximation substantially overestimates the depth of the minimum in the warm dense regime. In the $T \to 0$ limit the STLS and KSDT fits agree well with the

[7] There is some subtlety here. The exchange contribution arises from the inversion procedure when passing from the grand canonical to canonical ensemble. By inversion we mean that we first write $\mu = \mu_0 + \mu_{\mathrm{x}} + \mu_{\mathrm{c}}$ to the same order in the interaction as Ω. We then expand Ω and its derivatives about μ_0 and match coefficients to find μ_{x} [14, 25, 27]. One finds that, within the RPA [14], we can write $F_{\mathrm{xc}} \approx \Omega_{\mathrm{xc}}(\mu_0)$, however other quantities (such as the internal energy) may contain contributions from the inversion. A useful formula for the evaluation of derivatives of Ω is $\frac{\partial}{\partial\mu} I_\nu(\eta) = \beta\nu I_{\nu-1}(\eta)$.

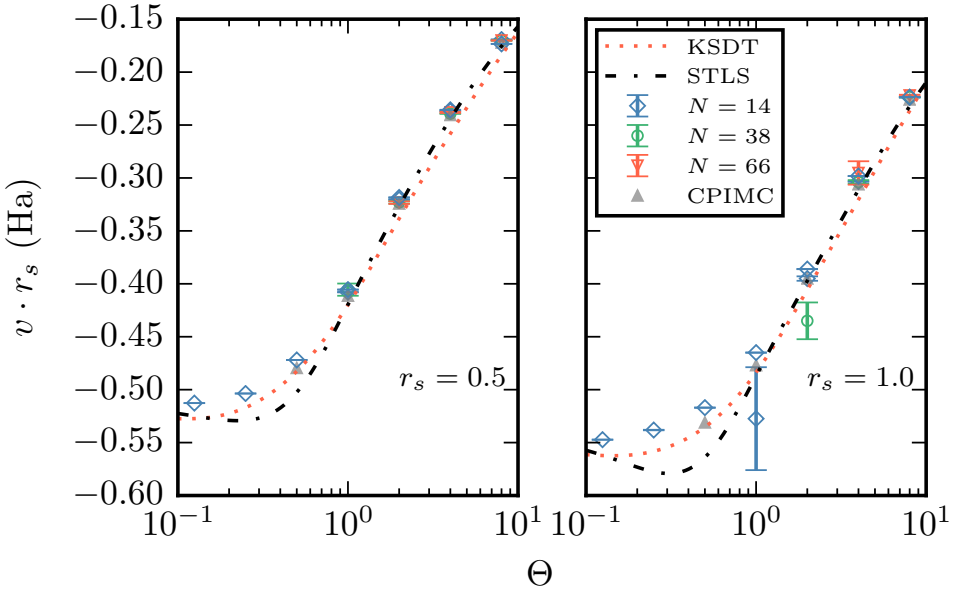

Figure 5.9: Size corrected DMQMC results for v for the $r_s = 0.5$ and $r_s = 1.0$, $\zeta = 0$ UEG. Also plotted are the CPIMC and STLS results. At $\Theta = 1$ we have data with and without the initiator approximation, where the point without the initiator approximation has the larger error bar.

i-DMQMC results, particularly when twist averaging is performed. This is consistent with the excellent agreement between the STLS scheme and $T = 0$ QMC results [28]. That being said, larger system sizes have to be addressed before anything definitive is said here and this is left for future investigation. We also see the importance of having an accurate fit to form of the free energy in the warm dense regime, with derivatives of the free energy being sensitive to the precise functional form.

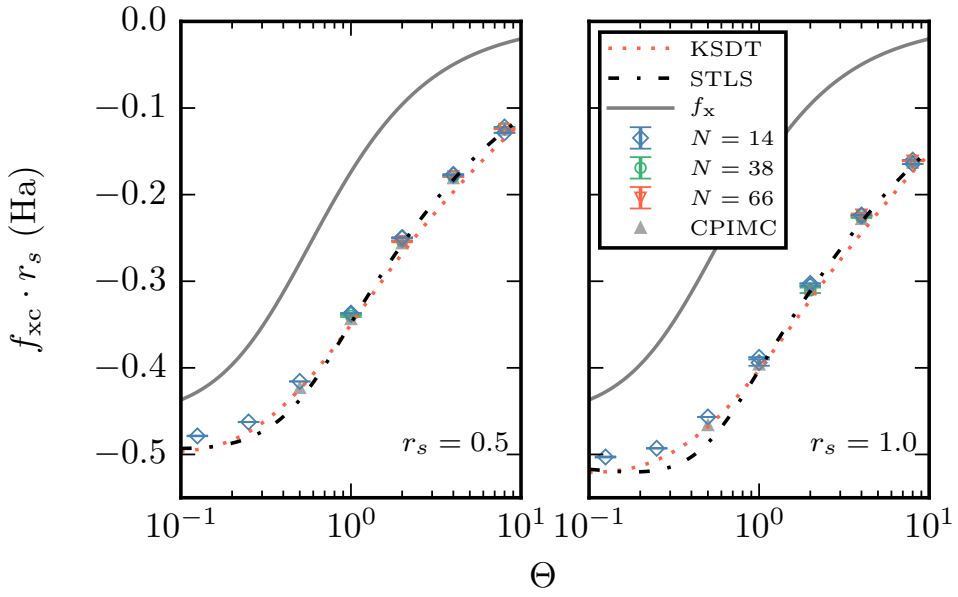

Figure 5.10: Size corrected DMQMC results for f_{xc} for the $r_s = 0.5$ and $r_s = 1.0$, $\zeta = 0$ UEG. Also plotted are exchange only, CPIMC, STLS and KSDT results. At $\Theta = 1$ we have data with and without the initiator approximation, where the point without the initiator approximation has the larger errorbar.

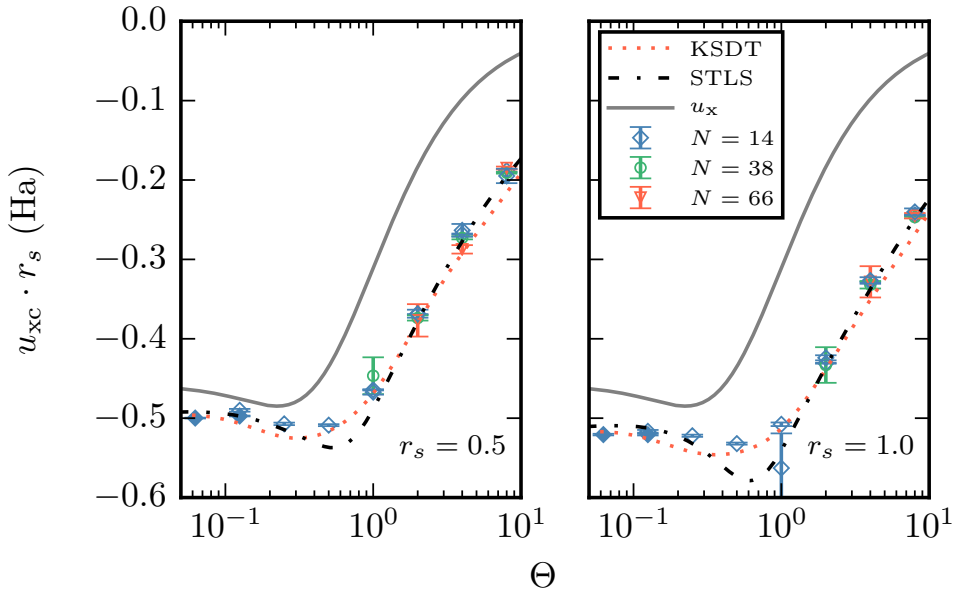

Figure 5.11: Size corrected DMQMC results for u_{xc} for the $r_s = 0.5$ and $r_s = 1.0$, $\zeta = 0$ UEG. Also plotted are exchange only (u_{x}), KSDT and STLS results. At $\Theta = 1$ we have data with and without the initiator approximation, where the point without the initiator approximation has the larger errorbar. The solid markers are twist averaged results for $\Theta \leq 0.125$; we estimate (by comparison to higher temperatures and larger particle numbers) that any residual systematic errors should be accounted for by the size of the markers.

5.4 Summary and Conclusions

In this chapter we investigated the use of analytic finite-size corrections for QMC simulations. We saw that, overall, the residual size errors were the most significant source of remaining error. Existing corrections, which work well at $T = 0$, were found to be not well suited to the high density and high temperature regime. By analysing the origin of these size corrections, we identified the source of this error and suggested an improved scheme that goes beyond leading order. We also suggested a route to calculating accurate size corrections for the exchange-correlation free and internal energies. Using these ideas, we believe size corrections for nearly any thermodynamic estimator can be found using various Maxwell relations. By performing extensive DMQMC calculations for different system sizes, we found that these analytic corrections significantly accelerate the convergence of QMC data with respect to system size in the warm dense regime and should be used in any future investigations. We found good agreement between the size corrected QMC data and the STLS scheme and also the KSDT fit at low temperatures. The deviations seen at higher temperatures are larger but unlikely to be physically important.

Conclusions

One beginning and one ending for a book was a thing I did not agree with.
– unnamed student [174]

In this thesis we have investigated the application of the DMQMC method to the warm dense uniform electron gas. We first showed the necessity in going beyond the original DMQMC algorithm when simulating interacting electrons in large plane wave basis sets. In developing the interaction picture variant, we significantly extended the scope of the domain of applicability of the DMQMC algorithm. This allowed us to evaluate the Helmholtz free energy, a capability that sets it apart from most other QMC methods. The finite basis set corrections we developed also proved crucial and could easily be adapted for other thermal electronic structure theory methods in the future.

We then adapted the initiator approximation to DMQMC to allow us to treat larger system sizes and lower temperatures. We used *i*-DMQMC to resolve the significant disagreement present between CPIMC and RPIMC results for the exchange correlation energy of the UEG and provided new data in the intermediate density regime where none existed before. These results confirmed that the restricted path approximation can introduce significant uncontrolled errors in the high density and low temperature regimes, a drawback that had not previously been appreciated or verified. These data for the polarised UEG at lower temperatures will provide useful input for the reparametrisation of the exchange correlation free energy of the UEG, with the spin polarised results playing a central role in understanding the spin interpolation the thermal LDA [51], and about which there is still some debate [175]. Although *i*-DMQMC has its limitations, it is presently complementary to existing path-integral methods (see Fig. 5.12) and is best suited to lower temperatures and higher densities.

Finally, we addressed the issue of extrapolating with electron number in the warm dense regime. To address the observed poor performance of existing analytic corrections, we showed (in collaboration with Tobias Dornheim, Tim Schoof and Simon Groth) the necessity to go beyond the existing low-order corrections in the high density and high temperature regimes. These relatively simple corrections we derived were shown to vastly improve the convergence of thermodynamic estimators, significantly reducing the computational burden of QMC methods. We further outlined how finite-size corrections to almost any other observable quantity could be derived using similar principles. These

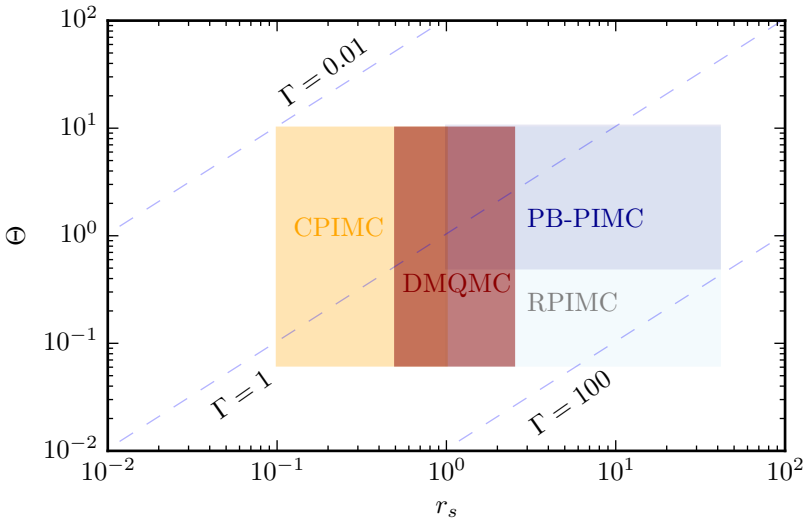

Figure 5.12: Regions of the (r_s, Θ) plane where different QMC methods are most efficient. i-DMQMC can often bridge the gap between low and high densities, particularly at lower temperatures.

results could be important when looking at more complicated systems and also at dense systems at $T = 0$ [176].

Of course this thesis is far from complete and many loose ends remain. Perhaps the most critical issue are the large biases introduced by the initiator approximation when estimating observables that do not commute with the Hamiltonian. Another concern is the numerical instability of the free energy estimator at low temperatures (and its error analysis). Beyond these technical issues, a more user-friendly simulation protocol would be helpful. Too many complicated convergence checks, basis-set corrections, and finite-size corrections are presently required to obtain accurate results.

Nevertheless, and somewhat unexpectedly, we have seen that we can sample the full thermal density matrix in Hilbert spaces containing more elements than there are atoms in the universe. The density matrix contains a wealth of information about a given system and ways of exploiting this should be investigated in the future[8]. As to the problem of warm dense matter, we are now in a position to investigate the performance of thermal density functional theory and its sensitivity to the choice of exchange-correlation free energy functional. Here, as in the ground state, QMC methods will be crucial. As we have seen, thermal density functionals and RPIMC suffer from biases in certain regimes. The next step should be to adapt DMQMC to simulate the full many-electron Hamiltonian of a real solid to address the extent to which these biases carry over and how much they matter.

[8] A novel example of this is given in Appendix B, where we attempt to evaluate spectral properties for one dimensional spin chains.

Appendices

Appendix A

Parallel Improvements to FCIQMC-like Algorithms

A.1 Introduction

An important feature of the FCIQMC algorithm not mentioned in the main text is its parallel efficiency. The spawning and death steps describe walkers evolving independently of each other with only annihilation requiring any communication. As such, walkers, and hence the computational load, can be efficiently distributed across many processors in contrast to conventional diagonalisation techniques. These are important qualities in an algorithm as future computational resources will likely be in the form of very large parallel computers, so it is vital that any new algorithms can take advantage of these machines in a cost and energy efficient manner.

The ability of a code to scale to a given number of processors can be categorised as either strong and weak scaling. Strong scaling involves scaling the code to more processors while keeping the problem sized fixed, while weak scaling involves increasing the problem size in proportion to the number of processors. For example, in FCIQMC if we double the number of cores and kept the number of walkers fixed we would like to see the running time halved, which would be a demonstration of strong scaling. Typically weak scaling is easier to achieve than strong scaling, which is largely down to both giving processors enough to do and also hiding any granularity which emerges when a job is refined as the processor count increases. Naturally there is a limit to the speed up that one can achieve for either scaling type, known as Amdahl's law [177], which states that the maximum parallel speed is given by

$$S(N_p) = \frac{1}{(1-P) + P/N_p}. \tag{A.1}$$

Here $S(N_p)$ is the measured speed up of a code run on N_p processors relative to the serial version and P is the percentage of the code which is parallelisable. Amdahl's law ignores a number of other factors which hamper scaling, principal of which is the overhead associated with communicating information between processors.

QMC methods are generally highly parallelisable with the P fraction being very close to one, so one of the biggest hindrances to achieving improved parallel scaling is this communication overhead. Recently it has been shown that a DMC code CASINO [145] can overcome some of these issues and be scaled to hundreds of thousands of cores by using a combination of load balancing of DMC walkers as well as non-blocking MPI communications [178]. This appendix is concerned with outlining how similar ideas can be applied to FCIQMC-like algorithms[1].

A.1.1 Parallel implementation of FCIQMC

To begin consider the standard parallel implementation of the FCIQMC algorithm, a more complete description of which can be found in [126]. Each processor stores a sorted main list of instantaneously occupied determinants containing the determinant's bit string representation, walker population as well as any simulation dependent flags. For each iteration every walker is given the chance to spawn to another connected determinant, with newly spawned walkers being added to a second spawned walker array. Given that a walker can spawn to any other connected determinant, which may in general reside on any processor, the spawned walker array is subdivided into N_p sections for ease of communication. After evolution a collective `MPI_AlltoAllv` is set up to communicate the spawned walker array to the appropriate processors. The annihilation step is then carried out by merging the subsequently sorted spawned walker array with the main list.

At every point in a simulation one needs to be able to determine which processor a determinant currently resides on as well as where newly occupied determinants should reside. Simply dividing the range of integer values a determinant can take among processors is not likely to ensure an even distribution of determinants given that highly excited determinants would tend to have larger integer representations. To counteract this one can use a hash function which maps the determinants bit string representation to another number within a given range. Determinants are then assigned to processors as

$$p \rightarrow \text{hash}(|D_{\mathbf{i}}\rangle) \bmod N_p, \tag{A.2}$$

where N_p is the number of processors and hash is a hash function. The above procedure, when implemented efficiently, results in an extremely fast highly parallel algorithm.

A.2 Load Balancing

The workload of the algorithm is primarily determined by the number of walkers on a given processor, but the above hashing procedure distributes work to processors on

[1] These ideas were suggested by Dr. James Spencer who also guided me through their implementation in HANDE.

a determinant basis. For hashing to be successful we would need roughly the same number of walkers and determinants on every processor. When scaling a problem of a fixed size to more processors, i.e. strong scaling, one observes that the distribution loses some of its uniformity with certain processors becoming significantly under and over populated which negatively affects the parallelism [126]. This is to be expected as in the limit $N_p \to N_{\text{Dets}}$ there would be quite a pronounced load imbalance unless each determinant's coefficient was of a similar magnitude (which can often be the case for strongly correlated systems). Naturally this limit is never reached, but the observed imbalance is largely a consequence of this increased refinement.

Ideally newly spawned walkers would be distributed to processors so as to achieve a uniform distribution automatically. This would be a computationally expensive task both in deciding where a psip should go and storing the resulting mapping which would have to be distributed to every other processor. A scheme is therefore desired which retains the fast lookup capabilities of hashing while achieving a more even distribution of walkers for a given number of processors. It was shown in [126] that isolating the Hartree-Fock determinant on a dedicated processor can significantly improve scaling by reducing the population imbalance across processors. Here we seek an improvement on this by allowing determinants to dynamically change processor in a simple fashion.

To start define an array p_{map} as

$$p_{\text{map}}(i) = i \bmod N_p, \tag{A.3}$$

so that its entries cyclically contain the processor IDs, $0, \ldots, N_p - 1$. Determinants are then initially mapped to processors as

$$p \to p_{\text{map}}(\text{hash}(|D_{\mathbf{i}}\rangle) \bmod N_p \times M). \tag{A.4}$$

Equation (A.4) reduces to Equation (A.2) when $M = 1$.

The walker population in each of these M bins on each processor can be determined and communicated to all other processors. In this way, every processor knows the total distribution of walkers across all processors. In redistributing the $N_p \times M$ bins we adopt a simple heuristic approach by only selecting bins belonging to processors whose populations are either above or below a certain user defined threshold. By redistributing bins in order of increasing population we can, in principle, isolate highly populated determinants while also allowing for a finer distribution. This procedure translates to a simple modification of p_{map} so that its entries now contain the processor IDs which give the determined optimal distribution of bin as demonstrated in Fig. A.1.

Finally, the walkers which reside in the chosen bins have to be moved to their new processor, which can simply be achieved using a communication procedure similar to that used for the annihilation stage. Some care needs to be taken that all determinants

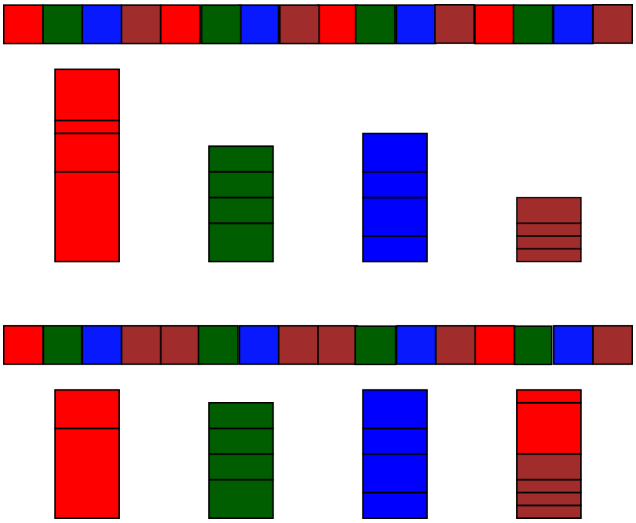

Figure A.1: Demonstration of load balancing procedure. Slots of walkers are distributed to achieve as even as load balance as possible.

are on their correct processors at a given iteration so that annihilation takes place correctly.

Once the population of walkers has stabilised the distribution across processors should be roughly constant, although small fluctuations will persist. With this in mind redistribution should only occur after this stabilisation has occurred and also should not need to occur too frequently. This ensures that the computational cost associated with performing load balancing is fairly minor in a large calculation. Additionally as M is increased the optimal distribution of walkers should be approached, although with an increase in computational effort.

A.3 Non-Blocking Communication

While load balancing should help in reducing the workload imbalance between processors and thus reduce latencies when performing communication, there is additional overhead in physically communication information between nodes across a network. Asynchronous or non-blocking communication can help alleviate some of these effects by overlapping communication with computation. The use of non-blocking communications was shown to greatly improve scaling in [178] and should potentially have a significant impact here. In contrast to DMC walkers, walkers are not completely independent as they must undergo annihilation at the same point in simulation time, so some modification of the algorithm is required.

To achieve this the following procedure was devised: Consider the evolution of walkers from τ to $\tau + \Delta\tau$, then for each processor the following steps are carried out:

1. Initialise the non-blocking receive of walkers spawned onto the current processor

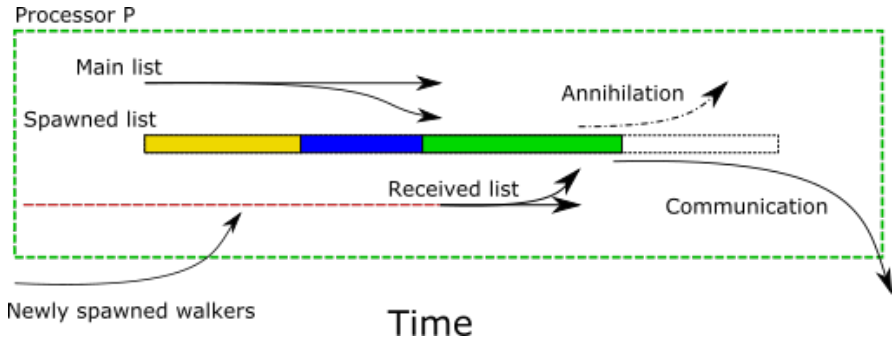

Figure A.2: Demonstration of modified evolution and annihilation procedure for non-blocking communications. The main list can be evolved while waiting for all new walkers to be spawned onto the current processor P, illustrated with a dashed red line. After which the received list is evolved and annihilation takes place whereby the received list and the spawned list is meged with the main list. The spawned list is divided into sections accorinding to which processor a psip resides on, this is indicated here by the different colour blocks present in the spawned list. Walkers which don't reside on the current processor, i.e. the yellow and blue blocks, are then communicated to the appropriate processors.

from time τ.

2. Evolve the main list to time $\tau + \Delta\tau$.
3. Complete the receive of walkers.
4. Evolve the received walkers to $\tau + \Delta\tau$.
5. Annihilate walkers spawned from the evolution of the two lists as well as the evolved received list with the main list on this processor.
6. Send remaining spawned walkers to their new processors.

Fig. A.2 illustrates these steps.

While this requires more work per iteration, it should result in improved efficiency if the time take to complete this work is less than the latency time. This also ensures faster processors can continue doing work, i.e. evolving the main list, while waiting for other processors to finish evolving their main lists. For communications to be truly overlapping the slowest processor would need to complete the steps above before the fastest processor reaches step (3), otherwise there will be latency as the received list cannot be evolved before all walkers spawned onto a given processor are received.

It should be pointed out that walkers spawned onto a processor at time τ are only annihilated with the main list after evolution to $\tau + \Delta\tau$, which differs from the normal algorithm. While annihilation is vital to attaining converged results [104, 114] the times at which it takes place is somewhat arbitrary, once walkers are annihilated at the same point in simulation time [179]. Communication between processors is also required when collecting statistics, however the usual collectives required for this can

simply be replaced by the corresponding non-blocking procedures. This does require that information is printed out in a staggered fashion but this is of minor concern.

Appendix B

Dynamics from DMQMC

Some of the most difficult to calculate properties of correlated electron systems are their excited states [180]. Excited states are of great interest due to their relationship to spectroscopic experiments [64] but traditional techniques struggle with the increased complexity of the problem.

Band structures calculated from DFT often differ quite markedly from experiments: the classic example being the overestimation of band gaps in semi-conductors (see, for example, [72])[1]. QMC methods mostly work in imaginary time making extracting dynamics difficult. Some success has been enjoyed by using the maximum entropy method [181] to aid in analytically continuing data onto the real axis however the presence of statistical errors somewhat limit this to sign-problem free models and small but finite temperatures. DMC and AFQMC can be applied to excited states in both solid state and molecular systems although the fixed node and fixed phase approximations results in larger statistical errors than corresponding ground state calculations and are generally restricted to the first few lowest lying states [82, 83, 182]. Recently dynamical mean field theory (DMFT) has enjoyed much success in treating strongly correlated electron systems in the thermodynamic limit [183]. DMFT is thought to take local dynamical correlations into account accurately although the theory is only strictly exact in the limit of infinite dimensions [183]. Cluster extensions [184] as well as attempts to use DMFT to improve upon DFT [185] are active aread of research. Other techniques such as Time-dependent DFT, GW and Bethe-Salpether theory enjoy success with treating excitations in weakly to moderately correlated systems [64].

Perhaps the most robust way to calculate excited states is by exact diagonalisation [180]. While this is even more restricted by matrix size than a corresponding ground state calculation, often quite a number of indicative features of the many-body spectrum can be deduced by studying small clusters [186]. Typically one uses a Lanczos type method to calculate the first few lowest lying excited states and uses some physical insight to deduce the general features of the corresponding spectral function. A technique similar to Lanczos, which has been widely used for calculating the local density of states

[1] Interpreting the Kohn-Sham eigenvalues as being related to the quasi-particle excitations is also not necessarily straightforward [64].

(LDOS) has been the recursion method [187]. Here we demonstrate how and analogue of the recursion method for operators can be used in the context of DMQMC.

B.1 The Recursion Method

The time evolution of an operator $\hat{A}$ in the Heisenberg picture is given as

$$\hat{A}(t) = e^{i\hat{H}t}\hat{A}(0)e^{-i\hat{H}t} \tag{B.1}$$

$$= e^{i\hat{\hat{L}}t}\hat{A}(0), \tag{B.2}$$

where $\hat{\hat{L}}$ is the Liouvillian super-operator, $\hat{\hat{L}}\hat{A} = [\hat{H}, \hat{A}]$, and $\hbar$ has been set to 1. Then at zero temperature a time dependent correlation function can be written as

$$C(t) = -i\theta(t)\langle\psi_0|e^{i\hat{\hat{L}}t}\hat{A}(0)\hat{A}^\dagger(0)|\psi_0\rangle, \tag{B.3}$$

where $|\psi_0\rangle$ is the ground state eigenfunction of $\hat{H}$ and $\theta(t)$ is the Heaviside step function. Inserting complete sets of eigenstates into Eq. (B.3) one gets

$$C(t) = -i\theta(t)\sum_j |\langle\psi_0|A|\psi_j\rangle|^2 e^{i(E_0-E_j)t}, \tag{B.4}$$

where the fact that $\hat{\hat{L}}|\psi_i\rangle\langle\psi_j| = (E_i - E_j)|\psi_i\rangle\langle\psi_j|$ has been used. Upon Fourier transform of Eq. (B.4) one finds that

$$\tilde{C}(\omega) = \sum_j \frac{\langle\psi_0|\hat{A}|\psi_j\rangle\langle\psi_j|\hat{A}^\dagger|\psi_0\rangle}{\omega + E_0 - E_j + i\eta}, \tag{B.5}$$

where η is a small positive constant introduced to ensure convergence of the integral. One can then define the spectral density of $\tilde{C}$ as

$$\mathcal{A}(\omega) = -\frac{1}{\pi}\lim_{\eta\to 0^+}\mathrm{Im}\{\tilde{C}(\omega)\} \tag{B.6}$$

$$= \sum_j |\langle\psi_0|A|\psi_j\rangle|^2\delta(\omega + E_0 - E_j). \tag{B.7}$$

Thus, the spectral density gives a weighted distribution of the excitations possible in a many body system when acted upon with a perturbation $\hat{A}$. For instance the single particle Green's function can be calculated by replacing $\hat{A}$ with $\hat{c}_i$, from which the local density of states could be calculated (once time ordering and anti-commutators are correctly restored). Naturally, calculating $\mathcal{A}(\omega)$ requires a knowledge of the eigenstates of $\hat{H}$ which is an impossible task so approximations are necessary.

One way to find an approximation to $\mathcal{A}(\omega)$ is via the recursion method [187]. In the past this was an effective technique for calculating the local density of states by trans-

forming the Hamiltonian into a tridiagonal matrix which could be easily diagonalised. An extension of this to tridiagonalise the Liouvillian in a basis of operators was subsequently developed by Lee [188]. A good review of both Liouvillian and Hamiltonian recursion as well as applications to many-body systems can be found in [189].

As in the usual recursion method for the LDOS, the tridiagonal form of the Liovillian superoperator can be expressed in a basis of orthogonal operators, $\{\hat{u}_0, \hat{u}_1, \ldots, \hat{u}_n\}$, in terms of a three term recurrence

$$b_{n+1}\hat{u}_{n+1} = \hat{\hat{L}}\hat{u}_n - a_n\hat{u}_n - b_n\hat{u}_{n-1}. \tag{B.8}$$

The above formula is iterated by specifying a $\hat{u}_0$ ($\hat{u}_{-1} = \hat{0}$), and calculating

$$a_n = (\hat{u}_n, \hat{\hat{L}}\hat{u}_n) \tag{B.9}$$

$$\hat{v}_{n+1} = \hat{\hat{L}}\hat{u}_n - a_n\hat{u}_n - b_n\hat{u}_{n-1} \tag{B.10}$$

$$b_{n+1}^2 = (\hat{v}_{n+1}, \hat{v}_{n+1}) \tag{B.11}$$

$$\hat{u}_{n+1} = \hat{v}_{n+1}/b_{n+1}. \tag{B.12}$$

The inner product , $(\cdot, \cdot)$, chosen above can take many forms depending on the problem at hand, once the Liouvillian remains Hermitian with respect to it. For instance, the Trace-norm inner product

$$(A, B) = \frac{\mathrm{Tr}(A^\dagger B)}{\mathrm{Tr}(\mathbf{I})}, \tag{B.13}$$

can be used for infinite temperature quantities or the more usual ground state inner product

$$(A, B) = \frac{\langle\psi_0|A^\dagger B|\psi_0\rangle}{\langle\psi_0|\psi_0\rangle}, \tag{B.14}$$

for ground state properties. The inner products emerging during the recursion procedure are just (very complicated) expectation values which DMQMC can, in principle, calculate. Applying DMQMC to directly calculate the recursion coefficients above is likely not to succeed as expectation values are only correct on average while Eq. (B.8) requires the evaluation of inner products at every step.

As pointed out in [190] it is possible to work around this by transforming inner product. To see this consider the resolvent matrix element we wish to calculate

$$G_{00}(\omega) = \left(\hat{u}_0, \frac{1}{\omega\hat{\hat{I}} - \hat{\hat{L}}}\hat{u}_0\right), \tag{B.15}$$

where $\mathbf{T}$ is the tridiagonal matrix of $\hat{\hat{L}}$ and ω is a complex number. If the recursion procedure, Eq. (B.8), is carried out to all levels, then we can write

$$\hat{\hat{L}}\hat{u}_n = \sum_m \mathbf{T}_{mn}\hat{u}_m, \tag{B.16}$$

which is true regardless of the inner product used to find the orthogonal operators and the tridiagonal matrix $\mathbf{T}$. It follows then that

$$\frac{1}{\omega\hat{\hat{I}} - \hat{\hat{L}}}\hat{u}_n = \sum_m (\omega\mathbf{I} - \mathbf{T})^{-1}_{mn}\hat{u}_m, \tag{B.17}$$

and inserting Eq. (B.17) into Eq. (B.15), where now the inner product is the one determined by the physics, one finds that

$$G_{00}(\omega) = \sum_n (\omega\mathbf{I} - \mathbf{T})^{-1}_{n0}(\hat{u}_0, \hat{u}_n). \tag{B.18}$$

This allows us to calculate the spectral density at finite or zero temperatures by

1. Finding $\mathbf{T}$ and $\hat{u}_i$ using Liovillian recursion with an inner product such as Eq. (B.13) which is possible to do exactly.
2. Calculating the overlap elements $S_{0n} = (\hat{u}_0, \hat{u}_n)$ using DMQMC.

The spectral density found by using recursion depends on the number of recursion steps one carries out e.g. three levels of recursion will give a spectral density with three poles. For a many-body system the number of possible transitions, and thus poles, is enormous making a full tridiagonalisation of the Liovillian impossible for all but the smallest systems. This is true for both Liovillian and Hamiltonian recursion; however, what these techniques have in their favour is that for n levels of recursion one gets the first $2n$ moments of the spectral function correct, where the moments are defined as

$$\mu_n = \int_{-\infty}^{\infty} d\omega\omega^n \mathcal{A}(\omega). \tag{B.19}$$

To prove this one can expand Eq. (B.15) as a power series in ω and compare the coefficients appearing in the expansion with the recursion coefficients a_n and b_n. Therefore, the recursion method often gives very accurate results for integrals over the spectral density, which can be used to calculate correlation functions or constrain approximate determinations of a spectral density.

B.2 Implementation

As mentioned above, the general scheme is to first calculate a set of basis operators $\{u_0, \ldots, u_n\}$[2], a_n's and b_n's via the recursion method Eq. (B.8) using a suitable inner product, whereupon we can transform the resolvent element Eq. (B.18) using DMQMC to calculate the overlap elements S_{0n}. The simplest inner product to use is the trace inner product Eq. (B.13) and the easiest interacting Hamiltonian to implement this for

[2] Hats will now be dropped from all (super)operators.

is the Heisenberg Hamiltonian

$$\mathcal{H} = J \sum_{\langle ij \rangle} \mathbf{S}_i \cdot \mathbf{S}_j \tag{B.20}$$

$$= \frac{J}{4} \sum_{\langle ij \rangle} \sigma_i^x \sigma_j^x + \sigma_i^y \sigma_j^y + \sigma_i^z \sigma_j^z, \tag{B.21}$$

where σ_i^α, $\alpha = \{x, y, z\}$, are the Pauli matrices, J is the exchange parameter and the sum runs over all nearest neighbours. The action of L on a given starting vector can be carried out by hand only for a few levels of recursion before becoming intractable so a numerical approach is desired. We follow the implementation of Bohm and Leschke [191] who suggested a compact and fast way in which to carry out the recursion which will briefly be described here.

Consider an operator appearing during recursion. This will be generally be a product of Pauli matrices on the various sites of the lattice. The action of L on this complicated operator can be simplified using the following property of the Liouvillian

$$L(ABCD \dots Z) = L(A)BC \dots Z + \dots + ABC \dots L(Z). \tag{B.22}$$

On each site of a chain there can be one of four operators $\{I, \sigma^x, \sigma^y, \sigma^z\}$, which suggests two bits are required to store this information, explicitly $\{00, 01, 10, 11\}$. In this way any basis element appearing can simply be described using a single integer whose length is twice that of the number of sites. The action of L on a Pauli matrix can be simplified using the commutation relations

$$[\sigma_i^\alpha, \sigma_j^\beta] = i\delta_{ij}\varepsilon_{\alpha\beta\gamma}\sigma_j^\gamma, \tag{B.23}$$

while products of two Pauli matrices on the same site can be reduced to one using

$$\sigma^\alpha \sigma^\beta = \delta^{\alpha\beta} I + 2i\varepsilon_{\alpha\beta\gamma}\sigma^\gamma \tag{B.24}$$

and summation over repeated indices has been assumed. The commutation can be achieved numerically using a `NAND` operation of the on site bits with the mask 11 and working out the resulting permutation directly from the bit's integer representation. Additionally the memory requirements of storing the operators can be significantly reduced by ensuring the lists of operators remain sorted.

The resulting operators can be output to a file and sorted according to the number of σ^x and σ^y operators they contain. This resorting is to take advantage of the fact that for a given density matrix element, ρ_{ij}, only operators which have the same excitation level as ρ_{ij} can give a non-zero contribution to the trace. For example, if $\hat{\rho}_{ij} = |011001\rangle\langle 100110|$ the excitation level is 3 and any operator containing 2 Pauli matrices will automatically be zero in the trace. (Terrible) Code to perform all of this is

available at Ref. [192]. While this results in a large saving for any DMQMC calculation, there are potentially millions of operators contained in a given basis element making large calculations presently intractable.

B.3 Results and Conclusions

As a first example the ground state time dependent correlation function $\langle\psi_0|\sigma_0^z(t)\sigma_0^z(0)|\psi_0\rangle$ was calculated using the recursion method for the one dimensional 8 site anti-ferromagnetic, $J = 1$, chain with periodic boundary conditions.

Fig. B.1 shows the ground-state time-dependent correlation function calculated using the 4 overlaps calculated from DMQMC. Also shown is the 'best possible' result corresponding to eight levels of Liouvillian recursion using the ground state eigenvector found from exact diagonalisation.

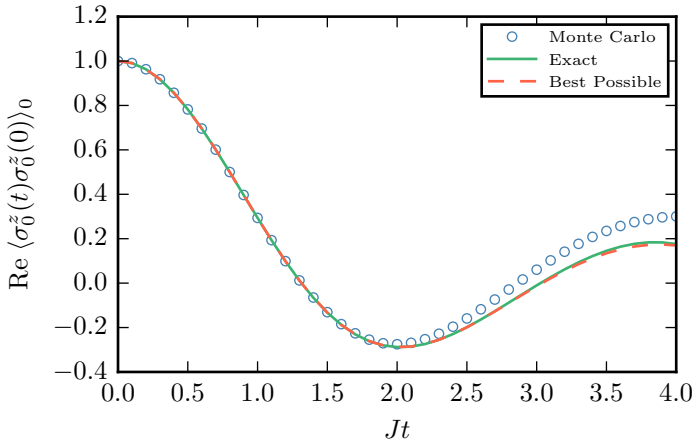

Figure B.1: Re$\{\langle\psi_0|\sigma_0^z(t)\sigma_0^z(0)|\psi_0\rangle\}$ as a function of time. Also shown are the correlation functions calculated from exact diagonalisation as well as the 'best possible' approximation defined in the text.

As it can be seen from Fig. B.1 the Monte Carlo data begins to differ from the exact result at $t \approx 1.8$ while the best possible result fairs substantially better. This is due to the fact that while carrying out recursion directly to $n/2$ levels reproduces the first n moments exactly, it also results in smaller errors for the higher order moments, while the overlap technique only reproduces the first n moments with worse errors beyond this. The overlap method thus results in an 8 term Taylor series approximation to the correlation function which can be seen by noting that

$$C(t) = -i\theta(t)\int_{-\infty}^{\infty} \mathrm{d}\omega \mathcal{A}(\omega) e^{-i\omega t} \tag{B.25}$$

$$= -i\theta(t)\sum_{n=0}^{\infty}\int_{-\infty}^{\infty} \mathrm{d}\omega \mathcal{A}(\omega)\frac{(-i\omega t)^n}{n!} \tag{B.26}$$

$$= -i\theta(t)\sum_{n=0}^{\infty}\frac{(-it)^n}{n!}\mu_n, \tag{B.27}$$

so a finite number of moments reproduces a truncated Taylor series approximation. While it is possible to reproduce the continued fraction coefficients which would have given the best possible result using a technique similar to the recursion method, it turns out this is quite sensitive to noise. For instance, the DMQMC calculation produced overlap elements which had errors of $\mathcal{O}(10^{-3})$ while an error of $\mathcal{O}(10^{-4} - 10^{-6})$ is required to begin to reproduce the best possible result. It is, in principle, possible to improve on this error by running the simulation for longer, however, the current simulation took about eighteen hours running on 36 cores so reducing the error further by brute force is not computationally efficient. A better solution would be to improve how the overlap elements are calculated within DMQMC so as to minimise the amount of bit manipulation required.

Consider a 32 site lattice, i.e. one for which exact diagonalisation begins to become an issue, there are $\mathcal{O}(10^5)$ operators which potentially need to be evaluated for each occupied basis element of which there could be $\mathcal{O}(10^6)$ so it is immediately obvious that this is an expensive task and not possible given the current implementation. For a given basis element most of these overlaps will give a zero contribution, so developing a fast mechanism of deciding which operators don't give a zero contribution is necessary.

Additionally due to these constraints it could be possible that only short time correlation function and only a few (about 5 or 6) excited states could be calculated so applications could be minimal. That being said, other techniques such as the time dependent density matrix renormalisation group rely on short time correlation functions for calculating dynamics [193] while the prospect of calculating essentially exact excited states for small molecules, beyond the reach of exact diagonalisation, could merit further study in this direction.

Bibliography

[1] H. J. Lipkin, "Collective Motion in Finite Many-Particle Systems. Part 1. The Violation of Conservation Laws," Ann. Phys. **9**, 272 (1960).

[2] P. A. M. Dirac, "Quantum Mechanics of Many-Electron Systems," Proc. Royal Soc. A **123**, 714 (1929).

[3] D. Alfè, M. J. Gillan, and G. D. Price, "Composition and Temperature of the Earth's Core Constrained by Combining Ab Initio Calculations and Seismic Data," Earth and Planet. Sci. Lett. **195**, 91 (2002).

[4] A. Benuzzi-Mounaix, S. Mazevet, A. Ravasio, T. Vinci, A. Denoeud, M. Koenig, N. Amadou, E. Brambrink, F. Festa, A. Levy, M. Harmand, S. Brygoo, G. Huser, V. Recoules, J. Bouchet, G. Morard, F. Guyot, T. D. Resseguier, K. Myanishi, N. Ozaki, F. Dorchies, J. Gaudin, P. M. Leguay, O. Peyrusse, O. Henry, D. Raffestin, S. L. Pape, R. Smith, and R. Musella, "Progress in Warm Dense Matter Study with Applications to Planetology," Phys. Scripta **T161**, 014060 (2014).

[5] F. Graziani, M. P. Desjarlais, R. Redmer, and S. B. Trickey, eds., *Frontiers and Challenges in Warm Dense Matter* (Springer, New York, 2014).

[6] J. J. Fortney, S. H. Glenzer, M. Koenig, B. Militzer, D. Saumon, and D. Valencia, "Frontiers of the Physics of Dense Plasmas and Planetary Interiors: Experiments, Theory, and Applications," Phys. Plasmas **16**, 041003 (2009).

[7] S. X. Hu, B. Militzer, V. N. Goncharov, and S. Skupsky, "First-Principles Equation-of-State Table of Deuterium for Inertial Confinement Fusion Applications," Phys. Rev. B **84**, 224109 (2011).

[8] R. Ernstorfer, M. Harb, C. T. Hebeisen, G. Sciaini, T. Dartigalongue, and R. J. D. Miller, "The Formation of Warm Dense Matter: Experimental Evidence for Electronic Bond Hardening in Gold," Science **323**, 1033 (2009).

[9] S. H. Glenzer, O. L. Landen, P. Neumayer, R. W. Lee, K. Widmann, S. W. Pollaine, R. J. Wallace, G. Gregori, A. Höll, T. Bornath, R. Thiele, V. Schwarz, W.-D. Kraeft, and R. Redmer, "Observations of Plasmons in Warm Dense Matter," Phys. Rev. Lett. **98**, 065002 (2007).

[10] S. Mukherjee, F. Libisch, N. Large, O. Neumann, L. V. Brown, J. Cheng, J. B. Lassiter, E. A. Carter, P. Nordlander, and N. J. Halas, "Hot Electrons Do the Impossible: Plasmon-Induced Dissociation of H_2 on Au," Nano Lett. **13**, 240 (2013).

[11] E. W. Brown, B. K. Clark, J. L. DuBois, and D. M. Ceperley, "Path-Integral Monte Carlo Simulation of the Warm Dense Homogeneous Electron Gas," Phys. Rev. Lett. **110**, 146405 (2013).

[12] T. Schoof, S. Groth, J. Vorberger, and M. Bonitz, "*Ab Initio* Thermodynamic Results for the Degenerate Electron Gas at Finite Temperature," Phys. Rev. Lett. **115**, 130402 (2015).

[13] L. D. Landau and E. M. Lifschitz, *Statistical Physics* (Pergamon, Oxford, 1969).

[14] A. L. Fetter and J. D. Walecka, *Quantum Theory of Many-Particle Systems* (McGraw-Hill, New York, 1971).

[15] K. Huang, *Statistical Mechanics*, 2nd ed. (Chichester, New York, 1987).

[16] R. P. Feynman, *Statistical Mechanics* (Westview Press, New York, 1998).

[17] M. Born and R. Oppenheimer, "Zur Quantentheorie der Molekeln," Ann. Phys. **389**, 457 (1927).

[18] R. M. Martin, *Electronic Structure: Basic Theory and Practical Methods* (Cambridge University Press, Cambridge, 2008).

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

[20] A. Szabo and N. S. Ostlund, *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory* (McGraw-Hill, New York, 1989).

[21] S. Raimes, *Many-Electron Theory* (North-Holland Publishing, Amsterdam, 1972).

[22] H. Q. Lin, "Exact Diagonalization of Quantum-Spin Models," Phys. Rev. B **42**, 6561 (1990).

[23] T. Helgaker, P. Jørgensen, and J. Olsen, *Molecular Electronic-Structure Theory* (Wiley, Chichester, 2012).

[24] J. M. Luttinger and J. C. Ward, "Ground-State Energy of a Many-Fermion System. II," Phys. Rev. **118**, 1417 (1960).

[25] W.-D. Kraeft, D. Kremp, W. Ebeling, and G. Röpke, *Quantum Statistics of Charged Particle Systems* (Springer, Berlin, 1986).

[26] J. Vorberger, M. Schlanges, and W.-D. Kraeft, "Equation of State for Weakly Coupled Quantum Plasmas," Phys. Rev. E **69**, 046407 (2004).

[27] D. Kremp, M. Schlanges, and W.-D. Kraeft, *Quantum Statistics of Nonideal Plasmas* (Springer, Berlin, 2005).

[28] G. Giuliani and G. Vignale, *Quantum Theory of the Electron Liquid* (Cambridge University Press, Cambridge, 2005).

[29] N. W. Ashcroft and N. D. Mermin, *Solid State Physics* (Saunders College, Philadelphia, 1977).

[30] E. H. Lieb, "The Stability of Matter," Rev. Mod. Phys. **48**, 553 (1976).

[31] J. M. Luttinger, "Fermi Surface and Some Simple Equilibrium Properties of a System of Interacting Fermions," Phys. Rev. **119**, 1153 (1960).

[32] L. D. Landau, "The Theory of a Fermi Liquid," Soviet Phys. JETP **3**, 920 (1957).

[33] R. Shankar, "Renormalization-Group Approach to Interacting Fermions," Rev. Mod. Phys. **66**, 129 (1994).

[34] V. E. Fortov, I. T. Iakubov, and A. G. Khrapak, *Physics of Strongly Coupled Plasma* (Clarendon Press, Oxford, 2006).

[35] N. D. Mermin, "Stability of the Thermal Hartree-Fock Approximation," Ann. Phys. **21**, 99 (1963).

[36] J. Sokoloff, "Some Consequences of the Thermal Hartree-Fock Approximation at Zero Temperature," Ann. Phys. **190**, 186 (1967).

[37] S. Hong and G. D. Mahan, "Temperature Dependence of the Hartree-Fock Approximation," Phys. Rev. B **50**, 7284 (1994).

[38] T. Sjostrom, F. E. Harris, and S. B. Trickey, "Temperature-Dependent Behavior of Confined Many-Electron Systems in the Hartree-Fock Approximation," Phys. Rev. B **85**, 045125 (2012).

[39] P. Hohenberg and W. Kohn, "Inhomogeneous Electron Gas," Phys. Rev. **136**, B864 (1964).

[40] W. Kohn and L. J. Sham, "Self-Consistent Equations Including Exchange and Correlation Effects," Phys. Rev. **140**, A1133 (1965).

[41] N. D. Mermin, "Thermal Properties of the Inhomogeneous Electron Gas," Phys. Rev. **137**, A1441 (1965).

[42] M. J. Gillan, "Calculation of the Vacancy Formation Energy in Aluminium," J. Phys. Cond. Mat. **1**, 689 (1989).

[43] M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos, "Iterative Minimization Techniques for *ab initio* Total-Energy Calculations: Molecular Dynamics and Conjugate Gradients," Rev. Mod. Phys. **64**, 1045 (1992).

[44] K. Burke, J. C. Smith, P. E. Grabowski, and A. Pribram-Jones, "Exact Conditions on the Temperature Dependence of Density Functionals," Phys. Rev. B **93**, 195132 (2016).

[45] A. Pribram-Jones and K. Burke, "Connection Formulas for Thermal Density Functional Theory," Phys. Rev. B **93**, 205140 (2016).

[46] A. Pribram-Jones, P. E. Grabowski, and K. Burke, "Thermal Density Functional Theory: Time-Dependent Linear Response and Approximate Functionals from the Fluctuation-Dissipation Theorem," Phys. Rev. Lett. **116**, 233001 (2016).

[47] R. G. Parr and W. Yang, *Density-Functional Theory of Atoms and Molecules* (Oxford University Press, New York, 1989).

[48] M. Stoitsov and I. Petkov, "Density Functional Theory at Finite Temperatures," Ann. Phys. **147**, 121 (1988).

[49] S. Pittalis, C. R. Proetto, A. Floris, A. Sanna, C. Bersier, K. Burke, and E. K. U. Gross, "Exact Conditions in Finite-Temperature Density-Functional Theory," Phys. Rev. Lett. **107**, 163001 (2011).

[50] E. W. Brown, J. L. DuBois, M. Holzmann, and D. M. Ceperley, "Exchange-correlation Energy for the Three-Dimensional Homogeneous Electron Gas at Arbitrary Temperature," Phys. Rev. B **88**, 081102 (2013).

[51] V. V. Karasiev, T. Sjostrom, J. Dufty, and S. B. Trickey, "Accurate Homogeneous Electron Gas Exchange-Correlation Free Energy for Local Spin-Density Calculations," Phys. Rev. Lett. **112**, 076403 (2014).

[52] T. Dornheim, S. Groth, T. Schoof, C. Hann, and M. Bonitz, "*Ab Initio* Quantum Monte Carlo Simulations of the Uniform Electron Gas without Fixed Nodes: The Unpolarized Case," Phys. Rev. B **93**, 205134 (2016).

[53] S. Groth, T. Schoof, T. Dornheim, and M. Bonitz, "*Ab initio* Quantum Monte Carlo Simulations of the Uniform Electron Gas without Fixed Nodes," Phys. Rev. B **93**, 085102 (2016).

[54] F. D. Malone, N. S. Blunt, E. W. Brown, D. K. K. Lee, J. S. Spencer, W. M. C. Foulkes, and J. J. Shepherd, "Accurate Exchange-Correlation Energies for the Warm Dense Electron Gas," Phys. Rev. Lett. **117**, 115701 (2016).

[55] T. Dornheim, S. Groth, T. Sjostrom, F. D. Malone, W. M. C. Foulkes, and M. Bonitz, "*Ab Initio* Quantum Monte Carlo Simulation of the Warm Dense Electron Gas in the Thermodynamic Limit," Phys. Rev. Lett. **117**, 156403 (2016).

[56] S. Azadi and W. M. C. Foulkes, "Fate of Density Functional Theory in the Study of High-Pressure Solid Hydrogen," Phys. Rev. B **88**, 014115 (2013).

[57] M. J. Gillan, D. Alfè, and A. Michaelides, "Perspective: How good is DFT for Water?" J. Chem. Phys. **144**, 130901 (2016).

[58] Y. S. Al-Hamdani, M. Ma, D. Alfè, O. A. von Lilienfeld, and A. Michaelides, "Communication: Water on Hexagonal Boron Nitride from Diffusion Monte Carlo," J. Chem. Phys. **142**, 181101 (2015).

[59] E. W. Brown, *Path Integral Monte Carlo and the Electron Gas*, Ph.D. thesis, University of Illinois at Urbana-Champaign (2014).

[60] V. V. Karasiev, L. Calderín, and S. B. Trickey, "Importance of Finite-Temperature Exchange Correlation for Warm Dense Matter Calculations," Phys. Rev. E **93**, 063207 (2016).

[61] S. Zhang, H. Wang, W. Kang, P. Zhang, and X. T. He, "Extended Application of Kohn-Sham First-Principles Molecular Dynamics Method with Plane Wave Approximation at High Energyfrom Cold Materials to Hot Dense Plasmas," Phys. Plasmas **23**, 042707 (2016).

[62] V. V. Karasiev, T. Sjostrom, and S. B. Trickey, "Generalized-Gradient-Approximation Noninteracting Free-Energy Functionals for Orbital-Free Density Functional Calculations," Phys. Rev. B **86**, 115101 (2012).

[63] V. V. Karasiev, D. Chakraborty, O. A. Shukruto, and S. B. Trickey, "Nonempirical Generalized Gradient Approximation Free-Energy Functional for Orbital-Free Simulations," Phys. Rev. B **88**, 161108 (2013).

[64] G. Onida, L. Reining, and A. Rubio, "Electronic Excitations: Density-Functional Versus Many-Body Green's-Function Approaches," Rev. Mod. Phys. **74** (2002).

[65] M. Gell-Mann and K. Brueckner, "Correlation Energy of an Electron Gas at High Density," Phys. Rev. **106**, 364 (1957).

[66] L. Hedin, "New Method for Calculating the One-Particle Green's Function with Application to the Electron-Gas Problem," Phys. Rev. **139**, A796 (1965).

[67] M. S. Hybertsen and S. G. Louie, "Electron Correlation in Semiconductors and Insulators: Band Gaps and Quasiparticle Energies," Phys. Rev. B **34**, 5390 (1986).

[68] J. Harl and G. Kresse, "Accurate Bulk Properties from Approximate Many-Body Techniques," Phys. Rev. Lett. **103**, 056401 (2009).

[69] M. Shishkin, M. Marsman, and G. Kresse, "Accurate Quasiparticle Spectra from Self-Consistent *GW* Calculations with Vertex Corrections," Phys. Rev. Lett. **99**, 246403 (2007).

[70] M. Shishkin and G. Kresse, "Implementation and Performance of the Frequency-Dependent *GW* Method within the PAW Framework," Phys. Rev. B **74**, 035101 (2006).

[71] M. Shishkin and G. Kresse, "Self-Consistent *GW* Calculations for Semiconductors and Insulators," Phys. Rev. B **75**, 235102 (2007).

[72] M. van Schilfgaarde, T. Kotani, and S. Faleev, "Quasiparticle Self-Consistent *GW* Theory," Phys. Rev. Lett. **96**, 226402 (2006).

[73] S. Tanaka and S. Ichimaru, "Thermodynamics and Correlational Properties of Finite-Temperature Electron Liquids in the Singwi-Tosi-Land-Sjölander Approximation," J. Phys. Soc. Jap. **55**, 2278 (1986).

[74] S. Hirata and X. He, "On the Kohn-Luttinger Conundrum." J. Chem. Phys. **138**, 204112 (2013).

[75] X. He, S. Ryu, and S. Hirata, "Finite-Temperature Second-Order Many-Body Perturbation and Hartree-Fock Theories for One-Dimensional Solids: An Application to Peierls and Charge-Density-Wave Transitions in Conjugated Polymers," J. Chem. Phys. **140**, 24702 (2014).

[76] M. R. Hermes and S. Hirata, "Finite-Temperature Coupled-Cluster, Many-Body Perturbation, and Restricted and Unrestricted HartreeFock Study on One-Dimensional Solids: Luttinger Liquids, Peierls transitions, and Spin- and Charge-Density Waves," J. Chem. Phys. **143**, 102818 (2015).

[77] R. Santra and J. Schirmer, "Finite-Temperature Second-Order Many-Body Perturbation Theory Revisited," Chem. Phys. (in press) .

[78] F. Perrot and M. W. C. Dharma-Wardana, "Spin-Polarized Electron Liquid at Arbitrary Temperatures: Exchange-Correlation Energies, Electron-Distribution Functions, and the Static Response Functions," Phys. Rev. B **62**, 16536 (2000).

[79] M. W. C. Dharma-Wardana, "Current Issues in Finite-T Density-Functional Theory and Warm-Correlated Matter," Computation **4**, 16 (2016).

[80] S. G. Brush, "Monte Carlo Study of a One-Component Plasma. I," J. Chem. Phys. **45**, 2102 (1966).

[81] J. P. Hansen, "Statistical Mechanics of Dense Ionized Matter. I. Equilibrium Properties of the Classical One-Component Plasma," Phys. Rev. A **8**, 3096 (1973).

[82] B. L. Hammond, *Monte Carlo Methods in Ab Initio Quantum Chemistry*, edited by W. A. Lester and P. J. Reynolds (World Scientific, Singapore, 1994).

[83] W. M. C. Foulkes, L. Mitas, R. J. Needs, and G. Rajagopal, "Quantum Monte Carlo Simulations of Solids," Rev. Mod. Phys. **73**, 33 (2001).

[84] B. Militzer and E. L. Pollock, "Variational Density Matrix Method for Warm, Condensed Matter: Application to Dense Hydrogen," Phys. Rev. E **61**, 3470 (2000).

[85] D. M. Ceperley, "Path Integrals in the Theory of Condensed Helium," Rev. Mod. Phys. **67**, 279 (1995).

[86] H. F. Trotter, "On the Product of Semi-Groups of Operators," Proc. Am. Math. Soc. **10**, 545 (1959).

[87] M. Suzuki, "Generalized Trotter's Formula and Systematic Approximants of Exponential Operators and Inner Derivations with Applications to Many-Body Problems," Comm. Math. Phys. **51**, 183 (1976).

[88] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, "Equation of State Calculations by Fast Computing Machines," J. Chem. Phys. **21**, 1087 (1953).

[89] W. K. Hastings, "Monte Carlo Sampling Methods Using Markov Chains and Their Applications," Biometrika **57**, 97 (1970).

[90] W. Feller, *An Introduction to Probability Theory and its Applications / Vol.1*, 3rd ed. (Wiley, New York, 1968).

[91] D. M. Ceperley, "Fermion Nodes," J. Stat. Phys. **63**, 1237 (1991).

[92] B. Militzer and K. P. Driver, "Development of Path Integral Monte Carlo Simulations with Localized Nodal Surfaces for Second-Row Elements," Phys. Rev. Lett. **115**, 176403 (2015).

[93] V. S. Filinov, "Analytical Contradictions of the Fixed-Node Density Matrix," High Temp. **52**, 615 (2014).

[94] T. Dornheim, S. Groth, A. Filinov, and M. Bonitz, "Permutation Blocking Path Integral Monte Carlo: A Highly Efficient Approach to the Simulation of Strongly Degenerate Non-Ideal Fermions," New J. of Phys. **17**, 73017 (2015).

[95] T. Schoof, M. Bonitz, A. Filinov, D. Hochstuhl, and J. W. Dufty, "Configuration Path Integral Monte Carlo," Contr. Plasma Phys. **51**, 687 (2011).

[96] T. Schoof, S. Groth, and M. Bonitz, "Introduction to Configuration Path Integral Monte Carlo BT - Complex Plasmas: Scientific Challenges and Technological Opportunities," (Springer International Publishing, Cham, 2014) pp. 153–194.

[97] N. V. Prokof'ev, B. V. Svistunov, and I. S. Tupitsyn, "Exact, Complete, and Universal Continuous-Time Worldline Monte Carlo Approach to the Statistics of Discrete Quantum Systems," J. Exp. Theor. Phys. **87**, 310 (1998).

[98] N. V. Prokof'ev and B. V. Svistunov, "Polaron Problem by Diagrammatic Quantum Monte Carlo," Phys. Rev. Lett. **81**, 2514 (1998).

[99] E. Gull, A. J. Millis, A. I. Lichtenstein, A. N. Rubtsov, M. Troyer, and P. Werner, "Continuous-Time Monte Carlo Methods for Quantum Impurity Models," Rev. Mod. Phys. **83**, 349 (2011).

[100] S. Zhang, J. Carlson, and J. E. Gubernatis, "Constrained Path Quantum Monte Carlo Method for Fermion Ground States," Phys. Rev. Lett. **74**, 3652 (1995).

[101] S. Zhang, J. Carlson, and J. E. Gubernatis, "Constrained Path Monte Carlo Method for Fermion Ground States," Phys. Rev. B **55**, 7464 (1997).

[102] S. Zhang and H. Krakauer, "Quantum Monte Carlo Method Using Phase-Free Random Walks with Slater Determinants," Phys. Rev. Lett. **90**, 136401 (2003).

[103] J. J. Shepherd, G. Booth, A. Grüneis, and A. Alavi, "Full Configuration Interaction Perspective on the Homogeneous Electron Gas," Phys. Rev. B **85**, 081103 (2012).

[104] G. H. Booth, A. J. W. Thom, and A. Alavi, "Fermion Monte Carlo without Fixed Nodes: A Game of Life, Death, and Annihilation in Slater Determinant Space." J. Chem. Phys. **131**, 054106 (2009).

[105] N. Trivedi and D. M. Ceperley, "Ground-state correlations of quantum antiferromagnets: A Green-function Monte Carlo study," Phys. Rev. B **41**, 4552 (1990).

[106] J. B. Anderson, "A Random-Walk Simulation of the Schrödinger equation: H_3^+," J. Chem. Phys. **63**, 1499 (1975).

[107] F. R. Petruzielo, A. A. Holmes, H. J. Changlani, M. P. Nightingale, and C. J. Umrigar, "Semistochastic Projector Monte Carlo Method," Phys. Rev. Lett. **109**, 230201 (2012).

[108] N. S. Blunt, S. D. Smart, J. A. F. Kersten, J. S. Spencer, G. H. Booth, and A. Alavi, "Semi-stochastic Full Configuration Interaction Quantum Monte Carlo: Developments and Applications," J. Chem. Phys. **142**, 184107 (2015).

[109] C. J. Umrigar, M. P. Nightingale, and K. J. Runge, "A Diffusion Monte Carlo Algorithm with Very Small TimeStep Errors," J. Chem. Phys. **99**, 2865 (1993).

[110] H. Flyvbjerg and H. G. Petersen, "Error Estimates on Averages of Correlated Data," J. Chem. Phys. **91**, 461 (1989).

[111] J. Olsen, O. Christiansen, H. Koch, and P. Jørgensen, "Surprising Cases of Divergent Behavior in Möller-Plesset Perturbation Theory," J. Chem. Phys. **105** (1996).

[112] See `http://www.hande.org.uk/` for details on how to acquire the source code.

[113] J. Spencer, N. Blunt, W. Vigor, F. D. Malone, W. Foulkes, J. J. Shepherd, and A. Thom, "Open-Source Development Experiences in Scientific Software: The HANDE Quantum Monte Carlo Project," J. Open Res. Soft. **3** (2015).

[114] J. S. Spencer, N. S. Blunt, and W. M. C. Foulkes, "The Sign Problem and Population Dynamics in the Full Configuration Interaction Quantum Monte Carlo Method," J. Chem. Phys. **136**, 054110 (2012).

[115] M. H. Kolodrubetz, J. S. Spencer, B. K. Clark, and W. M. C. Foulkes, "The Effect of Quantization on the Full Configuration Interaction Quantum Monte Carlo Sign Problem," J. Chem. Phys. **138**, 024110 (2013).

[116] G. H. Booth, D. Cleland, A. J. Thom, and A. Alavi, "Breaking the Carbon Dimer: The Challenges of Multiple Bond Dissociation with Full Configuration Interaction Quantum Monte Carlo Methods," J. Chem. Phys. **135**, 084104 (2011).

[117] D. Cleland, G. H. Booth, C. Overy, and A. Alavi, "Taming the First-Row Diatomics: A Full Configuration Interaction Quantum Monte Carlo Study," J. Chem. Theory Comput. **8**, 4138 (2012).

[118] G. H. Booth, A. Grüneis, G. Kresse, and A. Alavi, "Towards an Exact Description of Electronic Wavefunctions in Real Solids," Nature **493**, 365 (2013).

[119] R. E. Thomas, G. H. Booth, and A. Alavi, "Accurate *Ab Initio* Calculation of Ionization Potentials of the First-Row Transition Metals with the Configuration-Interaction Quantum Monte Carlo Technique," Phys. Rev. Lett. **114**, 033001 (2015).

[120] M. Kolodrubetz and B. K. Clark, "Partial Node Configuration-Interaction Monte Carlo as Applied to the Fermi Polaron," Phys. Rev. B **86**, 075109 (2012).

[121] A. J. W. Thom, "Stochastic Coupled Cluster Theory," Phys. Rev. Lett. **105**, 263004 (2010).

[122] R. S. T. Franklin, J. S. Spencer, A. Zoccante, and A. J. W. Thom, "Linked Coupled Cluster Monte Carlo," J. Chem. Phys. **144**, 44111 (2016).

[123] J. S. Spencer and A. J. W. Thom, "Developments in Stochastic Coupled Cluster Theory: The Initiator Approximation and Application to the Uniform Electron Gas," J. Chem. Phys. **144**, 84108 (2016).

[124] N. M. Tubman, J. Lee, T. Y. Takeshita, M. Head-Gordon, and K. B. Whaley, "A Deterministic Alternative to the Full Configuration Interaction Quantum Monte Carlo Method," J. Chem. Phys. **145**, 44112 (2016).

[125] A. A. Holmes, H. J. Changlani, and C. J. Umrigar, "Efficient Heat-Bath Sampling in Fock Space," J. Chem. Theory Comput. **12**, 1561 (2016).

[126] G. H. Booth, S. D. Smart, and A. Alavi, "Linear-Scaling and Parallelisable Algorithms for Stochastic Quantum Chemistry," Mol. Phy. **112**, 1855 (2014).

[127] W. A. Vigor, J. S. Spencer, M. J. Bearpark, and A. J. W. Thom, "Minimising Biases in Full Configuration Interaction Quantum Monte Carlo," J. Chem. Phys. **142**, 104101 (2015).

[128] W. A. Vigor, J. S. Spencer, M. J. Bearpark, and A. J. W. Thom, "Understanding and Improving the Efficiency of Full Configuration Interaction Quantum Monte Carlo," J. Chem. Phys. **144**, 94110 (2016).

[129] C. Overy, G. H. Booth, N. Blunt, J. J. Shepherd, D. Cleland, and A. Alavi, "Unbiased Reduced Density Matrices and Electronic Properties from Full Configuration Interaction Quantum Monte Carlo," J. Chem. Phys. **141**, 244117 (2014).

[130] R. E. Thomas, D. Opalka, C. Overy, P. J. Knowles, A. Alavi, and G. H. Booth, "Analytic Nuclear Forces and Molecular Properties from Full Configuration Interaction Quantum Monte Carlo," J. Chem. Phys. **143**, 054108 (2015).

[131] N. S. Blunt, A. Alavi, and G. H. Booth, "Krylov-Projected Quantum Monte Carlo Method," Phys. Rev. Lett. **115**, 050603 (2015).

[132] N. S. Blunt, S. D. Smart, G. H. Booth, and A. Alavi, "An Excited-State Approach within Full Configuration Interaction Quantum Monte Carlo," J. Chem. Phys. **143**, 134117 (2015).

[133] N. S. Blunt, T. W. Rogers, J. S. Spencer, and W. M. C. Foulkes, "Density-Matrix Quantum Monte Carlo Method," Phys. Rev. B **89**, 245124 (2014).

[134] D. Pines and D. Bohm, "A Collective Description of Electron Interactions: II. Collective vs Individual Particle Aspects of the Interactions," Phys. Rev. **85**, 338 (1952).

[135] D. Bohm and D. Pines, "A Collective Description of Electron Interactions: III. Coulomb Interactions in a Degenerate Electron Gas," Phys. Rev. **92**, 609 (1953).

[136] J. Bardeen, L. N. Cooper, and J. R. Schrieffer, "Theory of Superconductivity," Phys. Rev. **108**, 1175 (1957).

[137] S. H. Vosko, L. Wilk, and M. Nusair, "Accurate Spin-Dependent Electron Liquid Correlation Energies for Local Spin Density Calculations: A Critical Analysis," Can. J. Phys. **58**, 1200 (1980).

[138] J. P. Perdew and A. Zunger, "Self-Interaction Correction to Density-Functional Approximations for Many-Electron Systems," Phys. Rev. B **23**, 5048 (1981).

[139] J. P. Perdew and Y. Wang, "Accurate and simple Analytic Representation of the Electron-Gas Correlation Energy," Phys. Rev. B **45**, 13244 (1992).

[140] J. P. Perdew, K. Burke, and M. Ernzerhof, "Generalized Gradient Approximation Made Simple," Phys. Rev. Lett. **77**, 3865 (1996).

[141] D. M. Ceperley and B. J. Alder, "Ground State of the Electron Gas by a Stochastic Method," Phys. Rev. Lett. **45**, 566 (1980).

[142] T. Sjostrom and J. Dufty, "Uniform Electron Gas at Finite Temperatures," Phys. Rev. B **88**, 115123 (2013).

[143] P. P. Ewald, "Die Berechnung Optischer und Elektrostatischer Gitterpotentiale," Ann. Phys. **369**, 253 (1921).

[144] L. M. Fraser, W. M. C. Foulkes, G. Rajagopal, R. J. Needs, S. D. Kenny, and A. J. Williamson, "Finite-Size Effects and Coulomb Interactions in Quantum Monte Carlo Calculations for Homogeneous Systems with Periodic Boundary Conditions," Phys. Rev. B **53**, 1814 (1996).

[145] R. Needs, M. Towler, N. Drummond, and P. L. Ríos, "Casino Version 2.1 User Manual," University of Cambridge, Cambridge (2007).

[146] F. D. Malone, N. S. Blunt, J. J. Shepherd, D. K. K. Lee, J. S. Spencer, and W. M. C. Foulkes, "Interaction Picture Density Matrix Quantum Monte Carlo," J. Chem. Phys. **143**, 044116 (2015).

[147] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes : The Art of Scientific Computing* (Cambridge University Press, Cambridge, 1989).

[148] K. F. Riley, M. P. Hobson, and S. J. Bence, *Mathematical Methods for Physics and Engineering*, 3rd ed. (Cambridge University Press, Cambridge, 2006).

[149] D. Chakraborty, J. Dufty, and V. V. Karasiev, "Chapter Two - System-Size Dependence in Grand Canonical and Canonical Ensembles," in *Adv. Quant. Chem.*, Vol. 71 (2015) p. 11.

[150] J. J. Shepherd, G. H. Booth, and A. Alavi, "Investigation of the Full Configuration Interaction Quantum Monte Carlo Method Using Homogeneous Electron Gas Models," J. Chem. Phys. **136**, 244101 (2012).

[151] J. J. Shepherd, A. Grüneis, G. H. Booth, G. Kresse, and A. Alavi, "Convergence of Many-Body Wave-Function Expansions Using a Plane-Wave Basis: From Homogeneous Electron Gas to Solid State Systems," Phys. Rev. B **86**, 035111 (2012).

[152] T. Schoof, S. Groth, and M. Bonitz, "Towards ab Initio Thermodynamics of the Electron Gas at Strong Degeneracy," Contr. Plasma Phys. **55**, 136 (2015).

[153] J. Klimeš, M. Kaltak, and G. Kresse, "Predictive *GW* Calculations Using Plane Waves and Pseudopotentials," Phys. Rev. B **90**, 075125 (2014).

[154] A. Gulans, "Towards Numerically Accurate Many-Body Perturbation Theory: Short-Range Correlation Effects," J. Chem. Phys. **141**, 164127 (2014).

[155] E. Jones, T. Oliphant, P. Peterson, *et al.*, "SciPy: Open source scientific tools for Python," (2001–), [Online; accessed 2017-01-17].

[156] M. Boninsegni, N. Prokof'ev, and B. Svistunov, "Worm Algorithm for Continuous-Space Path Integral Monte Carlo Simulations," Phys. Rev. Lett. **96**, 070601 (2006).

[157] D. Cleland, G. H. Booth, and A. Alavi, "Communications: Survival of the Fittest: Accelerating Convergence in Full Configuration-Interaction Quantum Monte Carlo," J. Chem. Phys. **132**, 041103 (2010).

[158] S. Van Der Walt, S. C. Colbert, and G. Varoquaux, "The NumPy Array: A Structure for Efficient Numerical Computation," Comput. Sci. & Eng. **13**, 22 (2011).

[159] A. Baldereschi, "Mean-Value Point in the Brillouin Zone," Phys. Rev. B **7**, 5212 (1973).

[160] H. J. Monkhorst and J. D. Pack, "Special Points for Brillouin-Zone Integrations," Phys. Rev. B **13**, 5188 (1976).

[161] C. Lin, F. H. Zong, and D. M. Ceperley, "Twist-Averaged Boundary Conditions in Continuum Quantum Monte Carlo Algorithms," Phys. Rev. E **64**, 016702 (2001).

[162] N. Drummond, R. Needs, a. Sorouri, and W. Foulkes, "Finite-Size Errors in Continuum Quantum Monte Carlo Calculations," Phys. Rev. B. **78**, 125106 (2008).

[163] S. Azadi and W. M. C. Foulkes, "Systematic Study of Finite-Size Effects in Quantum Monte Carlo Calculations of Real Metallic Systems," J. Chem. Phys. **143**, 102807 (2015).

[164] M. Dagrada, S. Karakuzu, V. L. Vildosola, M. Casula, and S. Sorella, "Exact Special Twist Method for Quantum Monte Carlo Simulations," Phys. Rev. B **94**, 245108 (2016).

[165] G. G. Spink, R. J. Needs, and N. D. Drummond, "Quantum Monte Carlo study of the Three-Dimensional Spin-Polarized Homogeneous Electron Gas," Phys. Rev. B **88**, 085121 (2013).

[166] S. Chiesa, D. Ceperley, R. Martin, and M. Holzmann, "Finite-Size Error in Many-Body Simulations with Long-Range Interactions," Phys. Rev. Lett. **97**, 076404 (2006).

[167] M. Holzmann, R. C. Clay, M. A. Morales, N. M. Tubman, D. M. Ceperley, and C. Pierleoni, "Theory of Finite Size Effects for Electronic Quantum Monte Carlo Calculations of Liquids and Solids," Phys. Rev. B **94**, 035126 (2016).

[168] D. Pines and P. Nozières, *The Theory of Quantum Liquids: Normal Fermi Liquids*, Vol. 1 (WA Benjamin, 1966).

[169] R. Dandrea, N. Ashcroft, and A. Carlsson, "Electron Liquid at any Degeneracy," Phys. Rev. B **34**, 2097 (1986).

[170] See `http://www.github.com/fdmalone/uegpy/` for details on how to acquire the source code.

[171] F. Perrot and M. W. C. Dharma-Wardana, "Exchange and Correlation Potentials for Electron-Ion Systems at Finite Temperatures," Phys. Rev. A **30**, 2619 (1984).

[172] W. Kohn and J. M. Luttinger, "Ground-State Energy of a Many-Fermion System," Phys. Rev. **118**, 41 (1960).

[173] B. Horovitz and R. Thieberger, "Exchange Integral and Specific Heat of the Electron Gas," Physica **71**, 99 (1974).

[174] F. O'Brien, *At Swim-two-birds* (Dalkey Archive Press, 1939).

[175] S. Tanaka, "Improved Equation of State for Finite-Temperature Spin-Polarized Electron Liquids on the Basis of Singwi-Tosi-Land-Sjolander Approximation," arXiv preprint arXiv:1612.09037 (2016).

[176] J. J. Shepherd, *A Quantum Chemical Perspective on the Homogeneous Electron Gas*, Ph.D. thesis, University of Cambridge (2013).

[177] G. M. Amdahl, "Validity of the Single Processor Approach to Achieving Large-Scale Computing Capabilities," AFIPS Conf. Proc. **483**, 483 (1967).

[178] M. Gillan, M. Towler, and D. Alfé, "Petascale Computing Opens New Vistas for Quantum Monte Carlo," Psik Newsletter **103**, 32 (2011).

[179] J. S. Spencer and W. M. C. Foulkes, "Continuous Time Full Configuration Interaction Quantum Monte Carlo," Unpublished .

[180] P. Fulde, *Electron Correlations in Molecules and Solids* (Springer-Verlag, Berlin, 1993).

[181] J. Gubernatis, M. Jarrell, R. Silver, and D. Sivia, "Quantum Monte Carlo Simulations and Maximum Entropy: Dynamics from Imaginary-Time Data," Phys. Rev. B **44**, 6011 (1991).

[182] W. Purwanto, S. Zhang, and H. Krakauer, "Excited State Calculations Using Phaseless Auxiliary-Field Quantum Monte Carlo: Potential Energy Curves of Low-Lying C_2 Singlet States," J. Chem. Phys. **130**, 094107 (2009).

[183] A. Georges, G. Kotliar, W. Krauth, and M. Rozenberg, "Dynamical Mean-Field Theory of Strongly Correlated Fermion Systems and the Limit of Infinite Dimensions," Rev. Mod. Phys. **68**, 13 (1996).

[184] T. Maier, M. Jarrell, T. Pruschke, and M. H. Hettler, "Quantum Cluster Theories," Rev. Mod. Phys. **77**, 1027 (2005).

[185] G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, and C. A. Marianetti, "Electronic Structure Calculations with Dynamical Mean-Field Theory," Rev. Mod. Phys. **78**, 865 (2006).

[186] E. Dagotto, "Correlated Electrons in High-Temperature Superconductors," Rev. Mod. Phys. **66**, 763 (1994).

[187] R. Haydock, V. Heine, and M. Kelly, "Electronic Structure Based on the Local Atomic Environment for Tight-Binding Bands," J. Phys. C **5**, 2845 (1972).

[188] M. H. Lee, "Orthogonalization Process by Recurrence Relations," Phys. Rev. Lett. **49**, 1072 (1982).

[189] V. S. Viswanath and G. Muller, *The Recursion Method: Application to Many-Body Dynamics* (Springer-Verlag, Berlin, 1994).

[190] J. F. Annett, W. Matthew, C. Foulkes, and R. Haydock, "A Recursive Solution of Heisenberg's Equation and its Interpretation," J. Phys. Cond. Matt. **6**, 6455 (1994).

[191] M. Bohm and H. Leschke, "Dynamic Spin-Pair Correlations in a Heisenberg Chain at Infinite Temperature Based on an Extended Short-Time Expansion," J. Phys. A. **25**, 1043 (1992).

[192] See `http://www.github.com/fdmalone/lrec/` for details on how to acquire the source code.

[193] S. White and I. Affleck, "Spectral Function for the $S = 1$ Heisenberg Antiferromagetic Chain," Phys. Rev. B **77**, 134437 (2008).

Sin é