

Accelerating and Converging Stochastic Quantum Chemistry



Verena Andrea Neufeld

Department of Chemistry

University of Cambridge

This dissertation is submitted for the degree of

Doctor of Philosophy

Jesus College

September 2019

Für Oma

To Elias, my family and my friends

Declaration

The work in this thesis was done by the author under the supervision of Dr Alex Thom partially with collaborators as outlined in the relevant chapters. (Parts of) these chapters have already been published or are to be published, listed below. Please see the individual chapters for details on contributions. Please also see the Acknowledgements for some contributions.

- While most of this paper is only referenced/reviewed as it is mostly the work of others, section 3.4.2 in chapter 3 is a modified version of parts of:

J. S. Spencer, V. A. Neufeld, W. A. Vigor, R. S. T. Franklin, and A. J. W. Thom. *Large scale parallelization in stochastic coupled cluster*, J. Chem. Phys. **149**, 204103 (2018)[1].

(Reproduced (although slightly adapted) from [1], with the permission of AIP Publishing. This article may be downloaded for personal use only. Any other use requires prior permission of the authors and AIP Publishing. This article appeared in Ref. [1] and may be found at <https://doi.org/10.1063/1.5047420>.)

- Most of chapter 4 was published in (slightly modified to fit into the thesis and to include new results from the literature):

V. A. Neufeld and A. J. W. Thom. *A study of the dense uniform electron gas with high orders of coupled cluster*, J. Chem. Phys. **147**, 194105 (2017)[2].

(Reproduced (although slightly adapted) from [2], with the permission of AIP Publishing. This article may be downloaded for personal use only. Any other use requires prior

permission of the authors and AIP Publishing. This article appeared in Ref. [2] and may be found at <https://doi.org/10.1063/1.5003794>.)

- Most of chapter 5 was published in (slightly modified to fit into the thesis):

V. A. Neufeld and A. J. W. Thom. *Exciting determinants in Quantum Monte Carlo: Loading the dice with fast, low memory weights*, J. Chem. Theor. Comput. **15**, 1, 127-140 (2019)[3].

(Reproduced in part with permission from [3]. Copyright 2018 American Chemical Society. ACS Articles on Request author-directed link: <http://pubs.acs.org/articlesonrequest/AOR-S8PS7M2bqqB3jTNrrjav> and doi: <https://doi.org/10.1021/acs.jctc.8b00844>)

- Most of chapter 6 is to published in (slightly modified to fit into the thesis):

V. A. Neufeld, and A. J. W. Thom. *Accelerating stochastic quantum chemistry*, submitted[4].

(Reproduced in part with permission from Journal of Chemical Theory and Computation, submitted for publication. Unpublished work copyright 2019 American Chemical Society.)

- Chapter 7 includes a review of a subsection of the following publication (which is included more as a review and not directly as the other papers in the other chapters are):

J. S. Spencer, N. S. Blunt, [7 other authors], V. A. Neufeld, [7 other authors], and A. J. W. Thom. *The HANDE-QMC project: Open-source stochastic quantum chemistry from the ground state up*, J. Chem. Theor. Comput. **15**, 3, 1728-1742 (2019)[5].

(Reproduced in part with permission from [5]. Copyright 2019 American Chemical Society. ACS Articles on Request author-directed link: <http://pubs.acs.org/articlesonrequest/AOR-bhuCYVv5KiPXUyNFTwsd> and doi: <https://doi.org/10.1021/acs.jctc.8b01217>)

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text.

It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text.

It does not exceed the prescribed word limit for the relevant Degree Committee.

Verena Andrea Neufeld

December 2019

Accelerating and Converging Stochastic Quantum Chemistry

Verena Andrea Neufeld

Abstract: In the last decade, stochastic versions of quantum chemistry methods such as coupled cluster Monte Carlo (CCMC) or full configuration interaction quantum Monte Carlo (FCIQMC) have made highly accurate energy calculations possible that are not accessible to the corresponding deterministic methods (full configuration interaction and coupled cluster) at the same accuracy. CCMC and FCIQMC parallelize well and exploit the sparsity in the wavefunction which decreases memory costs and makes calculations in larger systems tractable. With CCMC it is straightforward to set up high order coupled cluster calculations, such as CCSDTQ5, which includes quintuple excitations explicitly. In this thesis, the convergence of the energy accuracy with the coupled cluster levels up to CCSDTQ5 was tested on the uniform electron gas, a model solid system, for various degrees of electron correlation. This gave information on what coupled cluster level is needed to reach sufficient accuracy when modelling a solid system. Before large solid systems can be modelled, the CCMC and FCIQMC algorithms need to be optimised. The efficiency in one of the crucial steps in these algorithms, the *spawn* step, was improved, keeping computational and memory costs as low as possible. Furthermore, the convergence of CCMC and FCIQMC was accelerated by employing a quasi-Newton propagation. Using the model system information of what coupled cluster level is needed and having made great progress towards accelerating these methods, the computation of highly accurate energies in solid or large molecular systems should be more feasible in the future.

Acknowledgements

First and foremost, I'd like to wholeheartedly thank Dr. Alex J. W. Thom for all his great ideas and insights, for supporting me with publishing my work, for creating such a wonderful research atmosphere and for always being there for me when I had questions. I joined his group with little notice and he quickly agreed to take me on which I am very grateful for.

I'd also like to thank the rest of the Thom group and overall Theory group over the years, particularly Ruth Franklin, Charlie Scott, Nick Blunt, Aron Cohen, Sam Niblett, Bang Huynh, Andreea Filip, Pablo López Ríos (see chapters for more details) and Aleks Reinhardt for their help, such as answering questions about research, formalities, code, etc, and for some fruitful discussions. I'd like to thank Andreea Filip, Nick Blunt, Aleks Reinhardt and Alex Thom for thorough proof reading of (parts of) this thesis. Thanks also goes to the rest of the HANDE QMC code team, especially former lead member James Spencer, which has not only provided me with a clean, working code to run Hubbard model and uniform electron gas calculations to start my PhD off with, but also supported me with reviewing code, fixing bugs, etc. It has been a pleasure to learn so many good coding principles in my PhD work and to contribute back to the code. I am very happy we got a paper presenting the code published! I'd like to thank the two summer students, Adam Přáda (summer 2017) and Jiří Etrych (summer 2018), that I co-supervised, for making progress on the coupled cluster for solid systems venture, some results being published in the HANDE QMC paper. I also like to thank the centre for doctoral training (CDT) cohort 1, that I am part of, for fruitful discussions as well, e.g. when trying to automatically determine the starting

point for reblocking. Another thank you also goes to the staff here at the chemistry department, especially the computing staff that keeps our computers and cluster going.

I would like to acknowledge the EPSRC Centre for Doctoral Training in Computational Methods for Materials Science for funding under grant number EP/L015552/1 and the Cambridge Philosophical Society for a studentship. This work used the ARCHER UK National Supercomputing Service (<http://www.archer.ac.uk>) and the UK Research Data Facility (<http://www.archer.ac.uk/documentation/rdf-guide>). Although not directly used for results in this thesis, the Darwin HPC cluster at Cambridge was used during the PhD. I also would like to acknowledge the various travel/conference/workshop funds that have enabled me to present my research to and get to know the community. These funds include various travel grants from Jesus college, Cambridge, grants from the IOP, including the C R Barber Trust, the Peter Salamon Award for Young Scientists at the Telluride Science Research Center, the Cambridge Philosophical Society, the Department of Chemistry at the University of Cambridge, and funding that the workshops (such as CECAM workshops) themselves provided.

I'd also like to thank the open-source community, e.g. for python packages. This thesis was written with the thesis template (<https://github.com/kks32/phd-thesis-template>) mainly by Krishna Kumar.

I thank my friends for being there for me and supporting me. Apart from friends already named above, this especially goes to Alice, Marise, Volker, Yvonne and the undergraduate crew Alex, Claudia, Ella, Emma, Harrison and Tom.

Last but not least, I wholeheartedly thank my parents, sister and of course Elias for supporting me in every way they can and for never leaving my side. This work would have been so much harder without you.

Table of contents

1	Introduction	1
2	Quantum Chemistry	5
2.1	Introduction	5
2.2	Hamiltonian and Born–Oppenheimer Approximation	8
2.3	Hartree–Fock Theory	9
2.3.1	Basis Sets	10
2.3.2	Fock operator and Self-Consistent Field Procedure	10
2.3.3	Interpretation and Beyond Hartree–Fock	11
2.4	Configuration Interaction	12
2.5	Coupled Cluster Theory	14
2.6	Uniform Electron Gas	16
3	Stochastic Quantum Chemistry	19
3.1	Quantum Monte Carlo	20
3.1.1	Monte Carlo sampling	20
3.1.2	Variational Monte Carlo	22
3.1.3	Projector Monte Carlo – Diffusion Monte Carlo	23
3.2	Full Configuration Interaction Quantum Monte Carlo	25
3.2.1	The Basics of FCIQMC	25
3.2.2	Initiator FCIQMC	29
3.2.3	Further Advances in FCIQMC	30

3.3	Coupled Cluster Monte Carlo	31
3.4	High Performance Computing – Parallelization	34
3.4.1	FCIQMC Parallelization	34
3.4.2	CCMC Parallelization	35
3.5	Calculations and Data Analysis	36
3.5.1	Running a Calculation	36
3.5.2	Analysing the Data	41
4	Converging Coupled Cluster Energies in the Uniform Electron Gas	45
4.1	Introduction	46
4.2	Extrapolation to Complete Basis Set Limit	47
4.3	Results	50
4.4	Discussion	53
4.5	Summary and Conclusions	56
5	Accelerating the Importance Sampling in the <i>Spawn</i> Step	57
5.1	Introduction	58
5.2	Excitation Generators	60
5.2.1	Heat Bath Excitation Generators	63
5.2.2	On-the-fly Power–Pitzer Excitation Generators	67
5.2.3	Pre-computed Power–Pitzer Excitation Generator	71
5.3	Results and Discussion	77
5.3.1	Coupled Cluster Monte Carlo	80
5.3.2	Full Configuration Interaction Quantum Monte Carlo	85
5.3.3	Practical Advice	87
5.4	Conclusion	87
6	Accelerating the Convergence with quasi-Newton	89
6.1	Introduction	90
6.2	Quasi-Newton Method	91

6.3	Deterministic Propagation	93
6.4	Stochastic Propagation	95
6.5	Application to the Chromium Dimer	99
6.6	Conclusion and Further Work	103
7	Concluding Remarks and Future Outlook	105
	Bibliography	109

Chapter 1

Introduction

Where accurate and precise electron system energy estimates are needed, quantum chemical computational methods are of vital importance. This thesis aims to get closer to obtaining accurate energies in (realistically sized) solid systems using stochastic versions of the quantum chemistry methods full configuration interaction (FCI) and coupled cluster (CC). A model solid system, the uniform electron gas, was used to investigate what accuracy coupled cluster gives in solid systems depending on the degree of correlation. To tackle larger systems, the stochastic methods were accelerated in their sampling and convergence.

Computational simulations can be viewed as one of the three ways to conduct science, in addition to experiments and non-computational theoretical work[6]. While experimental verification is crucial, simulations have several advantages. For example, it might be cheaper and quicker to run multiple simulations with slightly different parameters or under different conditions in materials design than it would be to perform the relevant experiments. Simulations can also connect theory with experiment for example by calculating the molecular motion with Newton's laws of motion in a molecular dynamics simulation to find structures[7] that can be experimentally verified. Different computational methods are needed for various problems, balancing accuracy with computational and memory costs[8]. At large length scales, continuum calculations solving macroscopic equations, such as stress and strain relationships[9], might be the

best choice. The most accurate *ab initio* simulation methods are from the field of quantum chemistry, giving energies of multi-electron systems. This accuracy comes with the price that they are limited to short length scales. It is one of the aims of quantum chemistry to increase the length scales that can be tackled by such methods before the calculation becomes infeasible due to computational and memory costs, and less accurate methods have to be used. Increasing the range of quantum chemistry methods can be achieved for example by making use of high-performance computing resources, better approximations or stochastic methods.

Applications where accurate and precise energies are required include material and molecular design with structure prediction, see e.g. Refs [10–13]. More accurate methods can also serve as benchmarks for higher level methods, see e.g. Refs [14–16]. Since quantum chemistry approaches are *ab initio* methods, they can also help to gain a greater scientific understanding of electron systems[17].

There are quantum chemistry methods of different levels of accuracy, scaling and ability to make use of high performance computing resources. Density functional theory (DFT)[18, 19] is a very successful method and has tackled systems of over 10^6 atoms[20]. However, DFT is not systematically improvable and its accuracy is very dependent on the choice of exchange-correlation functional. Reliable energies can be obtained with wavefunction approaches starting at Hartree–Fock (HF) theory, see e.g. textbooks [21, 22], and then including correlation as in the Møller-Plesset[23], coupled cluster[24–28] and configuration interaction theories, see e.g. textbooks [21, 29], for example. These post-Hartree–Fock methods are systematically improvable with higher accuracy coming at the cost of worse scaling with system size. Full configuration interaction (FCI) gives the exact solution in the given basis set. Coupled cluster is size extensive and reaches the FCI result in the limit of allowing all excitations. It has been shown that CCSD(T)[30], coupled cluster singles doubles and perturbative triples, can give chemical accuracy in many systems[27, 31]. The equation often solved by these quantum chemistry methods is the Schrödinger equation in the Born–Oppenheimer

approximation. This approach is followed in this thesis. Nuclear vibrations, relativistic effects, etc, are thus ignored and all calculations are done at zero Kelvin.

Conventional wavefunction methods such as FCI are memory intensive and suffer when limited in memory. The wavefunction representation can also be very sparse, i.e. it contains a significant number of small or zero entries. Projector Monte Carlo methods such as diffusion Monte Carlo[32] have been developed, giving exact energies in the fixed-node approximation, where the nodes of the wavefunction are fixed *a priori*, if the nodes of the wavefunction can be predicted *a priori*. Booth et al.[33] introduced FCI quantum Monte Carlo, FCIQMC, a stochastic version of FCI which decreases memory requirements by dealing with the more important states in the wavefunction with a higher probability and therefore making use of the sparsity, while being unbiased in its original form, since all states have a chance of being selected. FCIQMC is also highly parallelisable[34]. With FCIQMC in the initiator approximation[35], the uniform electron gas in a Hilbert space with a size of 10^{108} was tackled[36]. Such a calculation would be impossible with FCI due to the memory cost. Note that alternative stochastic configuration interaction versions exist[37, 38]. After this, a stochastic version of coupled cluster, coupled cluster Monte Carlo (CCMC), has been developed by Thom[39] which straightforwardly enables attempts at calculations with high coupled cluster levels.

(Deterministic) coupled cluster has not only been applied to molecules but also to solids, see e.g. Refs [40–43, 12] with coupled cluster levels up to CCSD(T). Making use of coupled cluster’s size-extensive nature — compared to truncated configuration interaction methods —[27] and the ability to give systematically improvable energies can be invaluable to material design. Progress has been made with the type of orbitals used, see e.g. Refs [40, 44, 45], and techniques to converge faster to the true energy with the number of \mathbf{k} points[46, 12]. The size of periodic solid that can be studied and the level of coupled cluster are still limited. Ref. [47] gives an overview of the coupled cluster in solids work in the literature. In this thesis, steps towards using CCMC for solid systems are made. First, a model system, the uniform electron gas,

was investigated with coupled cluster levels up to CCSDTQ5, which includes quintuple excitations directly, using CCMC, exploiting its easy set-up of various coupled cluster level calculations. For different levels of correlation, i.e. different electron densities, the accuracies of those coupled cluster levels were studied, giving information about the coupled cluster performance in “real” solids. FCIQMC has been applied to solids as well[41, 5], likewise limited in the system size it can tackle. The CCMC and FCIQMC algorithms were further developed here to make solid studies more feasible in the future. These improvements optimised the importance sampling in the spawn step, i.e. the excitation generator, and accelerated the convergence to the true energy with low extra costs.

The thesis is structured as follows: the following two chapters mostly describe existing quantum chemistry and stochastic quantum chemistry methods relevant to this work. After that, the uniform electron study and the two algorithmic improvements follow in three chapters. The last chapter concludes and offers a future outlook.

Chapter 2

Quantum Chemistry

This chapter gives an overview of conventional quantum chemical methods; Hartree–Fock theory and post-Hartree–Fock methods including coupled cluster theory and configuration interaction. More exhaustive and/or detailed introductions are given in the literature[21, 48–50, 27]. At the end, the uniform electron gas is introduced which will be the subject of several investigations in following chapters. The subsequent chapter will then outline the stochastic versions of coupled cluster[39] and configuration interaction[33], which are further developed and applied in this thesis.

2.1 Introduction

The aim of quantum chemistry is to find properties of systems consisting of interacting electrons *a priori*. Here, we restrict ourselves to systems at zero temperature and focus on ground state energies.

See textbooks, such as Refs [21, 49], for more details for this section. In quantum chemistry, a single electron is completely described by a wavefunction χ , specifying the electron’s distribution in the relevant space and its spin state. As a simple example and for simplicity ignoring spin in this case, in one-dimensional real space with position y , the probability of finding the electron between $y = a$ and $y = b$, with $a < b$, is given by $\int_a^b \chi^*(y)\chi(y)dy$, assuming χ is normalised, i.e. $\int_{\text{all space}} \chi^*(y)\chi(y)dy = 1$. The

mean electron position in this example is $\bar{y} = \int_{\text{all space}} \chi^*(y)y\chi(y)dy$, which in Dirac notation[51] can be written as $\bar{y} = \langle \chi | \hat{y} | \chi \rangle$ where \hat{y} is the position operator. χ and the corresponding energy E of the electron can be found by solving the time-independent Schrödinger equation[52]

$$\hat{H} |\chi\rangle = E |\chi\rangle. \quad (2.1)$$

This is an eigenvalue equation where \hat{H} is the Hamiltonian operator giving energy E when applied to χ . Later, \hat{H} will be defined for quantum chemistry problems. In the following studies, \hat{H} does not act on the spin part of the wavefunction. Note that spin does manifest itself in the Pauli exclusion principle[53], preventing two electrons from sharing the same quantum numbers simultaneously. The single-electron wavefunction can therefore be factorised into independent spatial and spin parts as

$$\chi(y, \sigma) = \varphi(y)s(\sigma). \quad (2.2)$$

For the convenience of a simpler notation, in multi-dimensional real space, the wavefunction is a function of \mathbf{x} which now contains position \mathbf{r} and spin σ , i.e. $\chi(\mathbf{x})$, with

$$\mathbf{x} = \{\mathbf{r}, \sigma\}. \quad (2.3)$$

The single electron wavefunctions $\chi(\mathbf{x})$, also called spinorbitals, can be combined to give many-body wavefunctions of N electrons $\Psi(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4, \dots, \mathbf{x}_N)$. The wavefunctions themselves are not “physical” quantities but their magnitudes squared, representing the electron probability distribution, are. The electrons are indistinguishable particles and so

$$|\Psi(\mathbf{x}_1, \dots, \mathbf{x}_i, \dots, \mathbf{x}_j, \dots, \mathbf{x}_N)|^2 = |\Psi(\mathbf{x}_1, \dots, \mathbf{x}_j, \dots, \mathbf{x}_i, \dots, \mathbf{x}_N)|^2 \quad (2.4)$$

where electron coordinates i and j have been swapped. For Ψ this means that the swap can introduce a sign change for Fermions and the multi-electron wavefunction is




Fig. 2.1 Pictorial representation of two Slater determinants with (spatial) orbitals ordered by energy ϵ . On the left hand side: restricted Hartree-Fock determinant D_0 . To the right hand side: doubly excited determinant with respect to D_0 . The blue curved arrows show those “excitations”. Drawn using Inkscape, <https://inkscape.org/> [Accessed: 11.12.2019].

antisymmetric, i.e.

$$\Psi(\mathbf{x}_1, \dots, \mathbf{x}_i, \mathbf{x}_{i+1}, \dots, \mathbf{x}_N) = -\Psi(\mathbf{x}_1, \dots, \mathbf{x}_{i+1}, \mathbf{x}_i, \dots, \mathbf{x}_N). \quad (2.5)$$

Thus, a simple product of different χ to form Ψ is not a Fermionic wavefunction but linear combinations with applicable signs of such products can be. The wavefunction can be expressed as a Slater determinant or combinations thereof. A Slater determinant $D_{\mathbf{i}}$ with N electrons is defined as

$$D_{\mathbf{i}} = \frac{1}{\sqrt{N!}} \begin{vmatrix} \chi_1(\mathbf{x}_1) & \chi_1(\mathbf{x}_2) & \chi_1(\mathbf{x}_3) & \dots & \chi_1(\mathbf{x}_N) \\ \chi_2(\mathbf{x}_1) & \chi_2(\mathbf{x}_2) & \chi_2(\mathbf{x}_3) & \dots & \chi_2(\mathbf{x}_N) \\ \chi_3(\mathbf{x}_1) & \chi_3(\mathbf{x}_2) & \chi_3(\mathbf{x}_3) & \dots & \chi_3(\mathbf{x}_N) \\ \dots & \dots & \dots & \dots & \dots \\ \chi_N(\mathbf{x}_1) & \chi_N(\mathbf{x}_2) & \chi_N(\mathbf{x}_3) & \dots & \chi_N(\mathbf{x}_N) \end{vmatrix}. \quad (2.6)$$

A swap in a column or a row corresponding to a swap in electron coordinate or spinorbitals will lead to a sign change which ensures that Ψ is antisymmetric and $\frac{1}{\sqrt{N!}}$ is a normalisation factor. Following Ref.[54], \mathbf{i} is a list of occupied orbitals in $\{\chi_1, \chi_2, \chi_2, \dots, \chi_N, \dots, \chi_M\}$ where $M > N$ and is therefore a unique index for the

determinant. A pictorial representation of two possible determinants is shown by figure 2.1. It is assumed throughout this thesis that χ are orthonormal. In the next section, the expression of the Hamiltonian will be discussed including a commonly employed approximation, the Born–Oppenheimer approximation[55].

2.2 Hamiltonian and Born–Oppenheimer Approximation

The form of the wavefunction Ψ was covered in the previous section. The Hamiltonian in the many-body Schrödinger equation is given by (see e.g. Ref. [21] who doesn't include $V_{\text{ext.}}$)

$$H = -\frac{1}{2} \sum_i \nabla_i^2 - \frac{1}{2} \sum_I \frac{1}{M_I} \nabla_I^2 + \sum_i \sum_{j>i} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|} - \sum_i \sum_J \frac{Z_J}{|\mathbf{r}_i - \mathbf{R}_J|} + \sum_I \sum_{J>I} \frac{Z_I Z_J}{|\mathbf{R}_I - \mathbf{R}_J|} + V_{\text{ext.}} \quad (2.7)$$

where atomic units have been used, as in the rest of the thesis unless stated otherwise. Capitalized index I corresponds to nuclei at position \mathbf{R}_I with nuclear charge Z_I and mass M_I and index i to electrons at positions \mathbf{r}_i . $V_{\text{ext.}}$ is the external potential that the nuclei and electrons are in.

Solving the Schrödinger equation with this Hamiltonian due to the number of variables involved is highly expensive. The Born–Oppenheimer approximation[55] can reduce the complexity drastically. Since the mass of an electron is less than 0.1% of the mass of a nucleon[56], the electrons and nuclei can be approximately decoupled. From the perspective of the electrons the nuclei are stationary and the nuclei feel the electrons instantly adapting to their movement. Focussing on the electrons, the wavefunction as introduced in the previous section is determined for a particular configuration of nuclei, for example for a particular atomic separation in the case of the chromium dimer in chapter 6. The simplified electronic Hamiltonian considers the nuclei to be part of the background potential and so they only appear in the Coulomb interaction between electrons and nuclei. It is given by (again, see Ref. [21] with a slightly different

$$\begin{aligned}
\langle D | \hat{H} | D \rangle &= \sum_{i \text{ occ. in } D} \langle i | \hat{h} | i \rangle + \frac{1}{2} \sum_{ij \text{ occ. in } D, j \neq i} \langle ij || ij \rangle \\
\langle D | \hat{H} | D_i^a \rangle &= \langle i | \hat{h} | a \rangle + \sum_{j \text{ occ. in } D, j \neq i} \langle ij || aj \rangle \\
\langle D | \hat{H} | D_{ij}^{ab} \rangle &= \langle ij || ab \rangle \\
\text{all other terms} &= 0
\end{aligned}$$

Table 2.1 Slater–Condon rules[57, 58, 21] for Hamiltonian elements $\langle D_{\mathbf{m}} | \hat{H} | D_{\mathbf{n}} \rangle$ made up of orthonormal spinorbitals χ [21], ignoring the constant term in \hat{H} . The labels for the orbitals that differ between $D_{\mathbf{m}}$ and $D_{\mathbf{n}}$ are shown, i.e. D_{ij}^{ab} is the same as determinant D except that spinorbitals i and j , i.e. χ_i and χ_j , have been replaced by orbitals χ_a and χ_b , ordered such that $i < j$ and $a < b$. $\langle ij || ab \rangle = \langle ij | ab \rangle - \langle ij | ba \rangle$, where $\langle ij | ab \rangle = \delta_{\sigma_i \sigma_a} \delta_{\sigma_j \sigma_b} \int \int \frac{\chi_i^*(\mathbf{r}_1) \chi_j^*(\mathbf{r}_2) \chi_a(\mathbf{r}_1) \chi_b(\mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} d\mathbf{r}_1 d\mathbf{r}_2$.

expression)

$$H = \underbrace{V'_{\text{ext.}}}_{\text{const. term, } = V_{\text{nn}} + V_{\text{ext.}}} - \underbrace{\frac{1}{2} \sum_i \nabla_i^2 - \sum_i \sum_J \frac{Z_J}{|\mathbf{r}_i - \mathbf{R}_J|}}_{\text{one-body term, } = \sum_i h(\mathbf{r}_i)} + \underbrace{\sum_i \sum_{j>i} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|}}_{\text{two body term}}, \quad (2.8)$$

where V_{nn} is the nuclear-nuclear repulsion and the two body term will form the Coulomb and exchange integrals.

In subsequent sections, Hamiltonian matrix elements $\langle D_{\mathbf{m}} | \hat{H} | D_{\mathbf{n}} \rangle$ will have to be evaluated. Slater–Condon rules[57, 58, 21] exploit the orthogonality of the spinorbitals χ and so drastically simplify these calculations. Table 2.1 gives the expressions as a function of the number of differing orbitals between $D_{\mathbf{m}}$ and $D_{\mathbf{n}}$.

2.3 Hartree–Fock Theory

In Hartree–Fock theory (see textbooks such as Refs [21, 22, 59] for this section), the many-body wavefunction Ψ is expressed as a single Slater determinant D_0 . The molecular orbitals χ that make up D_0 are optimised variationally to minimise the Hartree–Fock energy $E_{\text{HF}} = \langle D_0 | \hat{H} | D_0 \rangle$, which is an upper bound for the true ground state energy in that basis set. First, choosing the appropriate atomic orbitals to start

with is described below, followed by the self-consistent field algorithm for finding those optimised molecular orbitals χ and additional information.

2.3.1 Basis Sets

Before a Hartree–Fock calculation is run, the basis set for the atomic orbitals and the parameters in the Hamiltonian have to be chosen, e.g. the position and kind of nuclei in a molecular simulation. In a uniform electron gas calculation, discussed later, the basis set often consists of plane waves, see later. In non-periodic, molecular calculations, a simple basis set is for example STO-3G[60] where three Gaussian functions mimic each Slater function representing an atomic orbital. Basis sets can also include higher order orbitals that have polarising capabilities for example. It might also for example be necessary to represent two p orbitals with different basis functions since they differ in that particular problem in a split valence basis set[50], such as in 6-31G[61]. Another split valence set used here is Ahlrich’s SV basis set[62]. Dunning’s[63] series of cc-pVYZ ($Y = \{D, T, Q, \dots\}$) basis set represents all valence functions with two (double - “D”), three (triple - “T”), etc, sets of Gaussian functions. It also includes correlating functions for virtual orbitals for post-Hartree–Fock calculations that give parts of the correlation energy[64]. A molecular orbital χ is then a linear combination of these atomic orbitals φ with the coefficients found using Hartree–Fock theory.

2.3.2 Fock operator and Self-Consistent Field Procedure

By using a Lagrangian minimisation where $E_{\text{HF}} = \langle D_0 | \hat{H} | D_0 \rangle$ is optimised with respect to χ under the constraint that the spinorbitals χ are orthonormal, the following equation with Fock operator \hat{F} can be written down. The Fock operator is a single particle operator where electrons are affected by the other electrons in a mean-field manner. Hartree–Fock is therefore a mean-field theory. Spin is not explicitly shown here as it is

not directly affected. See reference [21] for details. The Fock operator is given by

$$\hat{F}\chi_i(\mathbf{r}_1) = \sum_j \epsilon_{ij}\chi_j(\mathbf{r}_1), \quad (2.9)$$

where ϵ_{ij} are the Fock values, i.e. elements of the Fock matrix, and

$$\hat{F} = \hat{V}'_{\text{ext.}} + \hat{h} + \hat{J} - \hat{K} \quad (2.10)$$

where \hat{h} is the one-body operator introduced previously and Coulomb operator \hat{J} acts on orbital $\chi_i(\mathbf{r}_1)$ as

$$\hat{J}\chi_i(\mathbf{r}_1) = \sum_j \int \frac{\chi_j^*(\mathbf{r}_2)\chi_i(\mathbf{r}_1)\chi_j(\mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} d\mathbf{r}_2 \quad (2.11)$$

and the exchange operator \hat{K} acts as

$$\hat{K}\chi_i(\mathbf{r}_1) = \sum_j \int \frac{\chi_j^*(\mathbf{r}_2)\chi_i(\mathbf{r}_2)\chi_j(\mathbf{r}_1)}{|\mathbf{r}_1 - \mathbf{r}_2|} d\mathbf{r}_2, \quad (2.12)$$

When $i = j$, the effects of \hat{J} and \hat{K} exactly cancel. In a canonical basis, the Fock values are diagonal, i.e. zero unless $i = j$, and equation 2.9 is an eigenvalue equation.

\hat{F} depends on $\{\chi_i\}$ in \hat{J} and \hat{K} . Therefore, equation 2.9 has to be solved self consistently. First, the form of $\{\chi_i\}$ (or the related charge density) is guessed, then the Fock matrix with elements $\langle\chi_j|\hat{F}|\chi_i\rangle = \epsilon_{ij}$ is evaluated. This information can then be used to update the set of spinorbitals $\{\chi_i\}$. This cycle is repeated until convergence is reached.

2.3.3 Interpretation and Beyond Hartree–Fock

When using canonical orbitals or orbitals where occupied and virtual orbitals were localised independently, equation 2.9 implies that $\langle\chi_i|\hat{F}|\chi_a\rangle = 0$ unless $i = a$ which is significant since therefore $0 = \langle\chi_i|\hat{F}|\chi_a\rangle = \langle\mathbf{D}_0|\hat{H}|\mathbf{D}_i^a\rangle$ where \mathbf{D}_i^a differs from \mathbf{D}_0 by a single excitation. This is Brillouin’s theorem whose result also holds for i occupied

and a virtual in the reference where occupied and virtual orbitals have been localised independently, a result interesting for chapter 5 when approximating Hamiltonian matrix elements for sampling weights. Having optimised the molecular orbitals χ , the many-body wavefunction Ψ is expressed as a single Slater determinant. However, E_{HF} is only an upper bound to the true ground state energy E_0 . At the complete basis set limit, the difference is defined as the correlation energy $E_{\text{corr.}}$ as

$$E_0 = E_{\text{HF}} + E_{\text{corr.}} \quad (2.13)$$

Since $E_{\text{HF}} \geq E_0$, $E_{\text{corr.}} \leq 0$. Throughout this thesis, $E_{\text{corr.}}$ is defined as above even in finite basis sets. Post-Hartree–Fock methods determine or approximate $E_{\text{corr.}}$ and therefore are more accurate than Hartree–Fock. The following two sections cover the post-Hartree–Fock methods configuration interaction and coupled cluster where stochastic versions were applied and further developed in this thesis.

2.4 Configuration Interaction

See textbooks, such as Refs [21, 29], for this section. To reach higher levels of accuracy than Hartree–Fock theory, configuration interaction includes multiple Slater determinants in the wavefunction Ψ . Starting with molecular orbitals optimised by Hartree–Fock theory, Ψ is expressed as a linear combination of Slater determinants $D_{\mathbf{i}}$ with coefficients $c_{\mathbf{i}}$ as

$$|\Psi\rangle = \sum_{\mathbf{i}} c_{\mathbf{i}} |D_{\mathbf{i}}\rangle. \quad (2.14)$$

In full configuration interaction, all possible Slater determinants conserving the number of electrons are included. Full configuration interaction gives the exact energies given the basis set. As is shown shortly, this space can be reduced by only including Slater determinants of the same symmetry as the Hartree–Fock determinant. However, even that space is often too large to run calculations on it. Unless the system is very strongly correlated, the determinants closest to the reference in terms of excitations,

i.e. Hartree–Fock, determinants contribute most to the ground state energy. Note that only double excitations directly contribute to the projected energy. Full configuration interaction can therefore be truncated by excitation level. For example, configuration interaction singles and doubles, CISD, only includes the reference determinant and determinants that differ from it by single or double excitations.

Having determined what D_i to include, the corresponding coefficients c_i are found by Lagrangian optimisation where the Lagrangian \mathcal{L} is

$$\mathcal{L} = \langle \Psi | \hat{H} | \Psi \rangle - E(\langle \Psi | \Psi \rangle - N) \quad (2.15)$$

with Lagrangian multiplier E and arbitrary normalisation constant N . Differentiating this with respect to all the c_i^* , equating to zero and rearranging gives

$$\mathbf{H}\mathbf{c} = E\mathbf{S}\mathbf{c}, \quad (2.16)$$

where \mathbf{H} is the Hamiltonian matrix with elements $\langle D_i | \hat{H} | D_j \rangle$, overlap matrix \mathbf{S} has elements $\langle D_i | D_j \rangle$ and is equal to the identity matrix here, \mathbf{c} is the eigenvector here containing the coefficients with eigenvalue E which is the energy of that state.

Provided the Hilbert space is small enough and the matrices can be stored in memory, equation 2.16 can be diagonalised directly or, if only a few energies such as the ground state energy are required, can be solved iteratively for those energies with a Newton propagation for example, see e.g. Ref. [29]. Sampling it stochastically can drastically reduce these memory requirements enabling calculations of larger systems. Shepherd et al.[36] have determined the energy of a uniform electron gas system with a Hilbert space of 10^{108} determinants using (initiator) full configuration interaction Quantum Monte Carlo[33, 35], a Hilbert space size that is impossible to store with a deterministic FCI calculation.

Configuration interaction is not size consistent if truncated. That means that for example the CISD energy of two infinitely separated Li^+ ions is not equal to twice

the CISD energy of one Li^+ ion. Coupled cluster, presented in the following section, overcomes this shortcoming and is size consistent.

2.5 Coupled Cluster Theory

Coupled cluster theory[24–28] also constructs its wavefunction as a sum of the Hartree–Fock determinant and multiple excitations thereof. However, unlike truncated configuration interaction it is size consistent. The wavefunction ansatz is

$$|\Psi\rangle = N \exp(\hat{T}) |\Psi_{\text{ref.}}\rangle, \quad (2.17)$$

with normalisation N , a reference, starting wavefunction $\Psi_{\text{ref.}}$, which in this thesis is set to D_0 , and

$$\hat{T} = \sum_{\mathbf{i}} t_{\mathbf{i}} \hat{a}_{\mathbf{i}} \quad (2.18)$$

where $\hat{a}_{\mathbf{i}}$ with amplitude $t_{\mathbf{i}}$ is an excitation operator that creates determinant $|D_{\mathbf{i}}\rangle$ as $|D_{\mathbf{i}}\rangle = \hat{a}_{\mathbf{i}} |D_0\rangle$. There is a one-to-one mapping between $\hat{a}_{\mathbf{i}}$ and $|D_{\mathbf{i}}\rangle$. As in configuration interaction, the space can be restricted. For example, in CCSDT, coupled cluster singles, doubles and triples, $\{\hat{a}_{\mathbf{i}}\}$ is restricted to only contain $\hat{a}_{\mathbf{i}}$ that create at most triple excitations from the reference determinant.

The equations to solve are the rearranged Schrödinger equation, projected onto a determinant $|D_{\mathbf{i}}\rangle$ which is in the allowed space, i.e. in CCSDT, $|D_{\mathbf{i}}\rangle$ only differs by at most a triple excitation from $|D_0\rangle$. To be explicit, the equations solved for amplitudes $t_{\mathbf{i}}$ are

$$\langle D_{\mathbf{i}} | \hat{H} - E | \Psi \rangle = \langle D_{\mathbf{i}} | (\hat{H} - E) \exp(\hat{T}) | D_0 \rangle = 0. \quad (2.19)$$

In conventional coupled cluster, often the linked version of this equation is solved, i.e.

$$\langle D_{\mathbf{i}} | \exp(-\hat{T}) (\hat{H} - E \delta_{0\mathbf{i}}) \exp(\hat{T}) | D_0 \rangle = 0, \quad (2.20)$$

since it reduces the number of terms until termination when expanded[65]. The energy E can be found by projecting $\langle D_0 |$ onto the action of the Hamiltonian on the coupled cluster wavefunction (see e.g. Ref. [28]), i.e.

$$E = \frac{\langle D_0 | \hat{H} \exp(\hat{T}) | D_0 \rangle}{\langle D_0 | \exp(\hat{T}) | D_0 \rangle}, \quad (2.21)$$

where the denominator is 1 and could therefore be dropped. However, it is left in here for illustration purposes since it is needed when calculating the projected energy in Monte Carlo versions. Following Ref. [28] and again considering the linked expression, vectorising and expanding the left hand side of equation 2.19 for $D_0 \neq D_i$ around excitor amplitudes \mathbf{t} , setting it to zero and ignoring terms higher than linear in $\delta\mathbf{t}$, gives

$$\delta\mathbf{t} = -\frac{d\mathbf{b}(\mathbf{t})}{d\mathbf{t}}^{-1} \mathbf{b}(\mathbf{t}) \quad (2.22)$$

where $\mathbf{b}_i = \langle D_i | \exp(-\hat{T})(\hat{H}) \exp(\hat{T}) | D_0 \rangle$. In a Newton propagation, this update equation is used to converge \mathbf{t} to the ground state amplitudes \mathbf{t}_0 . Since the first derivative is costly to evaluate, it is approximated by sums of Fock values in quasi-Newton propagations[28]. This will be used in chapter 6 where FCIQMC is accelerated in a similar manner. The unlinked expression (equation 2.19) is the equation solved in the original coupled cluster Monte Carlo formalism. A linked coupled cluster Monte Carlo version exists[65] but does not form part of this thesis. Details of the linked expression are therefore not further elaborated on here.

A popular form of coupled cluster theory is CCSD(T)[30], coupled cluster singles, doubles and perturbative triples. It has been shown to be able to give “chemical accuracy”, i.e. 1 kcal/mol[50, 27], for several systems with weak correlation. Note that it is still less accurate than FCI which is exact in the given basis set. However, the computational scaling is significant, $\mathcal{O}(m^{2l+2})$ [27], where l is the highest excitation level allowed in \hat{T} , e.g $l = 2$ for CCSD, and m is the system size here, meaning occupied or virtual orbitals. CCSD(T) has a scaling of $\mathcal{O}(m^7)$ [27].

2.6 Uniform Electron Gas

Finally, a technical introduction is given to the uniform electron gas (UEG)[66–69] which will be the subject of investigations in several subsequent chapters.

The uniform electron gas is a model solid system where the positive charge of the nuclei is smeared out to give a positive uniform background potential. There is no additional external potential $V_{\text{ext.}}$ in the studies here. In the three dimensional (3D) UEG here, the electrons are placed in a cubic box with sides of length L (real space). The UEG studied here is spin non-polarised, leaving the electron number density $\frac{N}{L^3}$ as the only parameter. The density is usually varied via the Wigner–Seitz radius r_s which in 3D is the radius of a sphere containing one electron on average. Density and r_s are related via $r_s = \left(\frac{3L^3}{4\pi N}\right)^{\frac{1}{3}}$. Due to the uniform potential, the solutions to the — non-interacting — Schrödinger equation are plane waves as in the basic quantum mechanical particle-in-a-box problem. The spatial wavefunctions are therefore approximated by

$$\varphi_{\mathbf{k}}(\mathbf{r}) \propto \exp(i\mathbf{k} \cdot \mathbf{r}), \quad (2.23)$$

where \mathbf{k} is the wavevector, representing the momentum, and \mathbf{r} the position in real space. Here, calculations were done in a momentum, \mathbf{k} , space basis. We use second quantisation to write down the expression for the Hamiltonian. A short interlude follows which describes the notation.

In second quantisation[70], a many-body wavefunction in \mathbf{k} space with electrons at \mathbf{k}_1 and \mathbf{k}_2 with spins σ_1 and σ_2 respectively can be written as

$$|\psi(\mathbf{k}_1, \sigma_1, \mathbf{k}_2, \sigma_2)\rangle = \hat{c}_{\mathbf{k}_2\sigma_2}^\dagger \hat{c}_{\mathbf{k}_1\sigma_1}^\dagger |0\rangle, \quad (2.24)$$

where creation operator $\hat{c}_{\mathbf{k}\sigma}^\dagger$ adds an electron with wavevector \mathbf{k} and spin σ . $|0\rangle$ is the empty space of this system. Contrarily, annihilation operator $\hat{c}_{\mathbf{k}\sigma}$ removes an electron with wavevector \mathbf{k} and spin σ . To obey Pauli’s principle,

$$\hat{c}_{\mathbf{k}_i\sigma_i}^\dagger \hat{c}_{\mathbf{k}_i\sigma_i}^\dagger |0\rangle = 0. \quad (2.25)$$

The anti-commutation relation is followed, e.g.

$$\{\hat{c}_{\mathbf{k}_j\sigma_j}, \hat{c}_{\mathbf{k}_i\sigma_i}^\dagger\} = \delta_{\mathbf{k}_i\mathbf{k}_j} \delta_{\sigma_i\sigma_j}. \quad (2.26)$$

Using this notation, the Hamiltonian of the 3D UEG can be written as

$$\hat{H} = \sum_{\mathbf{k}\sigma} \left(\frac{\mathbf{k}^2}{2} + \frac{V_{\text{Mad.}}}{2} \right) \hat{c}_{\mathbf{k}\sigma}^\dagger \hat{c}_{\mathbf{k}\sigma} + \frac{1}{2L^3} \sum_{\mathbf{k}_1\mathbf{k}_2\mathbf{q}} V(\mathbf{q}) \hat{c}_{(\mathbf{k}_2+\mathbf{q})\sigma_2}^\dagger \hat{c}_{(\mathbf{k}_1-\mathbf{q})\sigma_1}^\dagger \hat{c}_{\mathbf{k}_2\sigma_2} \hat{c}_{\mathbf{k}_1\sigma_1} \quad (2.27)$$

where $V(\mathbf{q}) = 0$ if $\mathbf{q} = \mathbf{0}$ and $V(\mathbf{q}) = \frac{4\pi}{q^2}$ otherwise. $V_{\text{Mad.}}$ is the Madelung constant. Note that the Fock values, $\langle \chi_i | \hat{F} | \chi_i \rangle$, which will be used in chapter 6, are (switching back away from second quantised notation)[68]

$$f_{i \text{ in } \mathbf{0}} = \frac{\mathbf{k}_i}{2} + \frac{V_{\text{Mad.}}}{2N} - \sum_{j \text{ in } \mathbf{0}, j \neq i} \langle ij | ji \rangle \quad (2.28)$$

for spinorbitals i in the reference and

$$f_{i \text{ not in } \mathbf{0}} = \frac{\mathbf{k}_i}{2} - \sum_{j \text{ in } \mathbf{0}, j \neq i} \langle ij | ji \rangle \quad (2.29)$$

for spinorbitals not in the reference determinant.

Chapter 3

Stochastic Quantum Chemistry

This chapter introduces quantum Monte Carlo[32, 71], full configuration interaction quantum Monte Carlo (FCIQMC)[33] and coupled cluster Monte Carlo (CCMC)[39].

Most parts of the chapter describe developments of others with the aim to give a background for later chapters. However, I have contributed to parts of the work in section 3.4.2. Figure 3.2 and some similar text to section 3.4.2 has been published in

- J. S. Spencer, V. A. Neufeld, W. A. Vigor, R. S. T. Franklin, and A. J. W. Thom. *Large scale parallelization in stochastic coupled cluster*, J. Chem. Phys. **149**, 204103 (2018)[1].

(Reproduced (although slightly adapted) from [1], with the permission of AIP Publishing. This article may be downloaded for personal use only. Any other use requires prior permission of the authors and AIP Publishing. This article appeared in Ref. [1] and may be found at <https://doi.org/10.1063/1.5047420>.)

The algorithms described in the paper have been implemented and most of the paper has been written by other authors, though I have discussed and contributed to several parts of the paper. I have mainly contributed to section IIIB where the bias in the parallelisation has been investigated. Figure 3 in the paper can be found as figure 3.2 here. I have initiated this section in the paper, collaboratively designed the figure, collected the data in the figure and created the figure.

3.1 Quantum Monte Carlo

This section opens up with a description of Monte Carlo sampling and then shows how it can be used to estimate the energy of electron systems with Variational Monte Carlo (VMC) and Diffusion Monte Carlo (DMC)[32, 71], which are early and highly successful quantum Monte Carlo methods.

3.1.1 Monte Carlo sampling

Monte Carlo sampling, see e.g. Refs [72, 32, 7, 71] for details, can be used to evaluate integrals for example. It usually involves computer simulations with a random number generator. This generator gives out pseudo-random numbers as random as possible with a period as long as possible. In this thesis, a fast Mersenne Twister random number generator[73] was used.

To understand Monte Carlo sampling better, instead of a random number generator, one could imagine a darts game where a square with sides of length 2 dm is inscribed in a square with sides of 4 dm, see figure 3.1a. See websites for examples of Monte Carlo simulations as a darts game¹. The ratio of these areas is

$$\frac{a_{\text{small square}}}{a_{\text{large square}}} = \frac{(2 \text{ dm})^2}{(4 \text{ dm})^2} = \frac{1}{4}. \quad (3.1)$$

This is the probability that, given a dart has hit the large square, it is inside the small square as well, provided darts are thrown uniformly at the square. As the number of darts gets large,

$$\frac{\text{Number of darts in small square}}{\text{Number of darts in large square}} \rightarrow \frac{a_{\text{small square}}}{a_{\text{large square}}} = \frac{1}{4}, \quad (3.2)$$

¹See e.g. *A simple Monte Carlo Method: Compute Pi*: <http://www.mathcs.emory.edu/~cheung/Courses/170/Syllabus/07/compute-pi.html>, *Monte Carlo Without the Math* by Z. Scott: <https://towardsdatascience.com/monte-carlo-without-the-math-90630344ff7b>, *An Overview of Monte Carlo Methods* by C. Pease: <https://towardsdatascience.com/an-overview-of-monte-carlo-methods-675384eb1694> [all accessed: 15.09.2019].

Thus, assuming the normalising area (the area of the large square) is known, the other area (the area of the small square) can be estimated using random, Monte Carlo sampling. Monte Carlo evaluations of integrals scale more favourably with the number of dimensions than other numerical integrations, e.g. quadrature techniques that scale as $\mathcal{O}((\text{number of grid points})^{\text{number of dimensions}})$ [74]. In a highly dimensional system, Monte Carlo thus can be a good choice for integration.

If these areas were weighted, one could imagine a scenario as in figure 3.1b, where sub-square regions of 1 dm^2 are allocated weights. In most cases they are given a unit weight except for one sub-square in the small square which carries 5 units weights. This means that a dart which lands in that heavier weighted sub-square will get five times the weight of a dart in other sub-square. The exact weighted area ratio is now

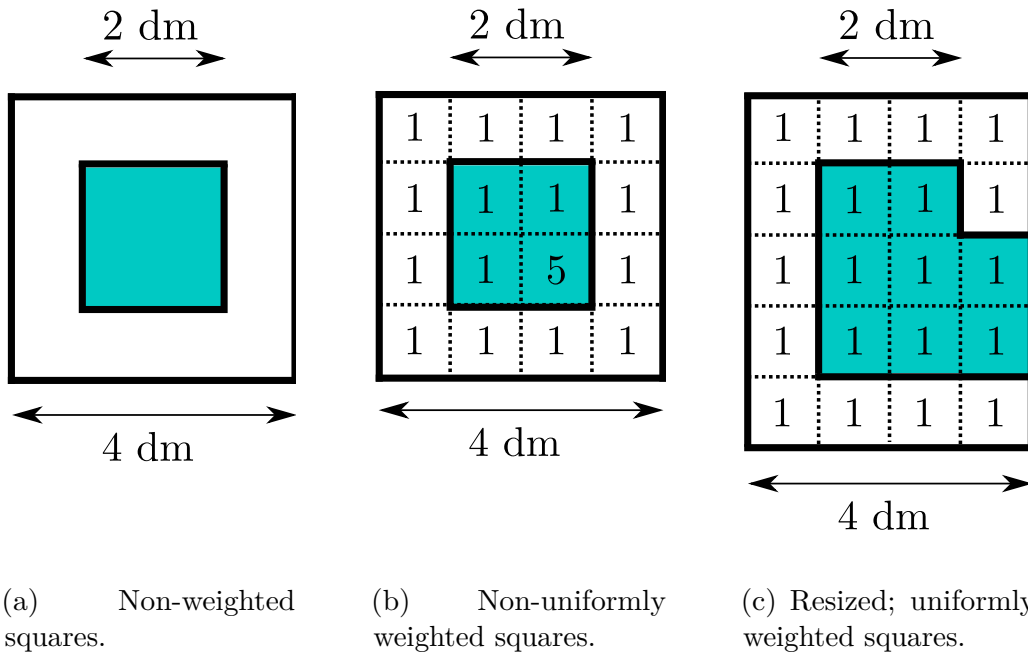


Fig. 3.1 Monte Carlo sampling to calculate the ratio of the shaded blue (weighted) area to total (weighted) area with uniformly thrown darts. (a) shows non-weighted areas, (b) has weights, and (c) demonstrates importance sampling applied to (b) having unified the weights but increased the probability to hit the shaded area accordingly.

$\frac{3 \times 1 + 1 \times 5}{15 \times 1 + 1 \times 5} = \frac{8}{20} = \frac{2}{5}$. When estimated with darts and giving them weights depending on their location, there are fluctuations in that estimate depending on how often the highly weighted sub-square gets hit. If it gets hit more often than expected, the area

estimate will be too high and vice versa. One could imagine an even more extreme case with a weight of 1000 instead of 5 in that sub-square. To decrease these fluctuations, a dart could be allowed to only “hit a part of that sub-square” and receive a fraction of the weight. Concretely, as done in figure 3.1c, the sub-square with weight 5 is spread out into five sub-squares of unit weight. The total area that originated from the old sub-square is now five times as likely to be hit but that effect is balanced by a five times lower weight. All darts carry the same weight again and fluctuations are decreased.

This process is called *importance sampling* (see section 10.4.3 in Ref. [71]) which decreases the variance and increases the efficiency of Monte Carlo simulations. It is the focus of chapter 5, where the importance sampling of CCMC and FCIQMC in the *spawn* step was improved. More formally, instead of sampling the area under function $a(x)$ with uniform probability, i.e.

$$\int a(x)dx, \quad (3.3)$$

the easier function $\frac{a(x)}{b(x)}$ could be sampled with normalised probability $b(x)$, i.e.

$$\int b(x) \frac{a(x)}{b(x)} dx. \quad (3.4)$$

As $b(x) \rightarrow ka(x)$, where k is a constant, the variance of $\frac{a(x)}{b(x)}$ tends to zero. In the example in the previous paragraph, this meant that all darts, even though they were sampled with a non-uniform probability distribution, contributed with one unit weight. Thus, in FCIQMC and CCMC, when distribution $a(x)$ was sampled with probability $b(x)$, $a(x)$ has to be divided by $b(x)$ to converge to the correct answer.

3.1.2 Variational Monte Carlo

One of the first quantum Monte Carlo methods was Variational Monte Carlo (VMC) which uses the variational principle to evaluate an upper bound for the ground state

energy, i.e. it is estimated by [71, 32]

$$\begin{aligned}
E &= \frac{\langle \Psi | \hat{H} | \Psi \rangle}{\langle \Psi | \Psi \rangle} = \frac{\sum_{\mathbf{t}} \langle \Psi | \psi(\mathbf{t}) \rangle \langle \psi(\mathbf{t}) | \hat{H} | \Psi \rangle}{\sum_{\mathbf{t}'} \langle \Psi | \psi(\mathbf{t}') \rangle \langle \psi(\mathbf{t}') | \Psi \rangle} \\
&= \sum_{\mathbf{t}} \frac{\langle \Psi | \psi(\mathbf{t}) \rangle}{\sum_{\mathbf{t}'} \langle \Psi | \psi(\mathbf{t}') \rangle \langle \psi(\mathbf{t}') | \Psi \rangle} \langle \psi(\mathbf{t}) | \hat{H} | \Psi \rangle \frac{\langle \psi(\mathbf{t}) | \Psi \rangle}{\langle \psi(\mathbf{t}) | \Psi \rangle} \\
&= \sum_{\mathbf{t}} \frac{|\langle \Psi | \psi(\mathbf{t}) \rangle|^2}{\sum_{\mathbf{t}'} |\langle \Psi | \psi(\mathbf{t}') \rangle|^2} \frac{\langle \psi(\mathbf{t}) | \hat{H} | \Psi \rangle}{\langle \psi(\mathbf{t}) | \Psi \rangle} = \sum_{\mathbf{t}} p(\mathbf{t}) E_{\text{proj.}}(\mathbf{t}),
\end{aligned} \tag{3.5}$$

summing over points \mathbf{t} occurring with probability $p(\mathbf{t})$ in the relevant space parameterized by \mathbf{t} . $\psi(\mathbf{t})$ is the wavefunction at point \mathbf{t} and Ψ is a combination of $\psi(\mathbf{t})$. These points are sampled with probability distribution $p(\mathbf{t})$ and the projected energy at those points, $E_{\text{proj.}}(\mathbf{t})$, are combined to determine energy E . This sampling of points \mathbf{t} might be done with the Metropolis algorithm [32, 75].

3.1.3 Projector Monte Carlo – Diffusion Monte Carlo

The accuracy of VMC depends on the form of the wavefunction chosen to optimise. Diffusion Monte Carlo (DMC) adds more flexibility to this, even though — in the fixed node approximation — the locations of the nodes of the trial wavefunction fix the nodes of the estimate of the ground state wavefunction, making them inflexible. DMC projects out the ground state wavefunction Ψ_0 , it is a projector Monte Carlo method.

In such quantum projector Monte Carlo methods, the projection is done in imaginary time τ , solving the imaginary time Schrödinger equation, see e.g. Ref. [32]. The resulting energy by the action of the Hamiltonian was shifted by the ground state energy E_0 ,

$$-\frac{d\Psi(\tau)}{d\tau} = (\hat{H} - E_0)\Psi(\tau). \tag{3.6}$$

Since \hat{H} is time-independent, the equation can be solved straightforwardly, with solution

$$\Psi(\tau) = \exp(-\tau(\hat{H} - E_0))\Psi(\tau = 0). \tag{3.7}$$

$\hat{H}\Psi(\tau \rightarrow \infty) = E_0\Psi(\tau \rightarrow \infty)$, i.e. $\Psi(\tau \rightarrow \infty) = \Psi_0$, since the exponent is never negative due to E_0 being the lowest energy. Writing $\Psi(\tau = 0) = \sum_i d_i \Psi_i$ with eigenfunctions of the Hamiltonian Ψ_i and their coefficients d_i , one could imagine considering the evolution of each Ψ_i individually, rewriting equation 3.7 as

$$\Psi(\tau) = \sum_i d_i \exp(-\tau(\hat{H} - E_0))\Psi_i \xrightarrow{\tau \rightarrow \infty} d_0 \Psi_0. \quad (3.8)$$

This shows that the projector $\hat{P} = \exp(-\tau(\hat{H} - E_0))$ projects the ground state wavefunction out of $\Psi(\tau = 0)$. Instead of using an exponential projector, a linear projector can also be used, i.e. $\hat{P} = 1 - \delta\tau(\hat{H} - E_0)$ with time step $\delta\tau = \frac{\tau}{n}$, where n is large, consult e.g. Ref. [76], see full configuration interaction quantum Monte Carlo (FCIQMC)[33] or coupled cluster Monte Carlo (CCMC)[39] section.

In DMC, equation 3.6 is seen as a diffusion equation (remember that the kinetic term in $\hat{H}\Psi(\tau)$ is of the form $-\nabla^2\Psi(\tau)$). A solution to this equation is sampled by delta functions, “walkers”, signed Monte Carlo particles. A projector helps to reach the ground state wavefunction. The magnitude of the solution also decides whether a walker “survives” to the next time step and if yes, whether an additional walker is created. See Ref. [32] for more details. Unfortunately, this algorithm converges to the bosonic ground state solution, having no nodes. One approach to combat this problem and reach the fermionic ground state solution, is the fixed node approximation where the positions of the nodes is fixed *a priori*. Walkers in the pockets between nodes have the same sign and the sign changes across a node. The closer the guessed nodes are to the true nodes in the ground state wavefunction, the better the energy estimate [32]. Note that this is a real space method, whereas the following approaches work in determinant space.

3.2 Full Configuration Interaction Quantum Monte Carlo

Full configuration interaction quantum Monte Carlo (FCIQMC)[33] solves full configuration interaction (FCI) stochastically. Truncated CIQMC, e.g. stochastic CISD, is also straightforwardly possible but not presented here.

3.2.1 The Basics of FCIQMC

FCIQMC, as FCI, expresses wavefunctions in Slater determinant space. While DMC in the fixed node approximation is limited by the knowledge of the node locations, FCIQMC does not need such *a priori* knowledge of the wavefunction. However, it is restricted by the basis set used and it is assumed that the symmetry of the initial wavefunction is compatible with the ground state wavefunction. The evolution to the estimate for the ground state wavefunction is also done in imaginary time. The signed “walkers”, or particles as they will be called in this thesis, then reside on those determinants. The initial wavefunction guess is often a few particles on the Hartree–Fock determinant, the “reference”. This corresponds to an initial wavefunction $\Psi(\tau = 0) \propto D_0$ with magnitude corresponding to the sum of the particle amplitudes. The wavefunction is then propagated by the linear projector \hat{P} . Applying \hat{P} n times,

$$\Psi(\tau = n\delta\tau) = \hat{P}^n \Psi(\tau = 0) = (1 - \delta\tau(\hat{H} - E_0))^n \Psi(\tau = 0) \propto \Psi_0 \quad (3.9)$$

as $n \rightarrow \infty$. As long as $|1 - \delta\tau(E_{\max.} - E_0)| < 1$, the projection to Ψ_0 is successful. $E_{\max.}$ is the maximum eigenvalue of \hat{H} . This ensures that every eigenvector except for Ψ_0 gets reduced by applying \hat{P} , while Ψ_0 is unaffected. $\delta\tau$ therefore has to be set such that $\delta\tau < \frac{2}{E_{\max.} - E_0}$ [76, 33]. As described in the previous chapter, the FCI wavefunction is

$$|\Psi(\tau)\rangle = \sum_{\mathbf{i}} c_{\mathbf{i}}(\tau) |D_{\mathbf{i}}\rangle, \quad (3.10)$$

where the dependence on the evolution with τ has been made explicit. Pre-multiplying equation 3.9 by $\langle D_i |$ and only showing the effect of one time step, the evolution of the coefficients c_i can be written as

$$c_i(\tau + \delta\tau) = c_i(\tau) - \delta\tau \sum_j \langle D_i | \hat{H} - E_0 | D_j \rangle c_j(\tau). \quad (3.11)$$

This is the equation sampled by FCIQMC to evolve its particle populations and to converge to an estimate of the ground state wavefunction. E_0 is (initially) unknown and needs to be estimated. The following paragraph describes the sampling details of basic non-initiator FCIQMC.

In the implementation used here, as in Booth et al.[33], FCIQMC does not sample which particles to propagate, all particles are evolved at every time step. The Hartree–Fock energy E_{HF} is first subtracted of the diagonal elements of $\langle D_i | \hat{H} | D_j \rangle$, making all diagonal elements non-negative[33]. The resulting FCIQMC energy estimate is therefore the correlation energy of E_0 . After a particle on some coefficient c_k has been selected, the *spawn* and *death* steps follow, see below. Notice that this implies that only occupied determinants, i.e. with non-zero c_i , are “active” and can lead to the creation of new particles. FCIQMC can therefore drastically reduce memory requirements as the other determinants do not have to be stored. Note that particles on the same determinant evolve independently in the *spawn* and *death* steps.

- *Spawn* Step (off-diagonal): This step deals with the contributions from the off-diagonal element,

$\langle D_i | \hat{H} - E_0 | D_j \rangle c_j = \langle D_i | \hat{H} | D_j \rangle c_j$. The selected particle on c_k takes the role of a part of c_j here. Starting from $|D_j\rangle$, another determinant $|D_i\rangle$ with a non-zero connection to $|D_j\rangle$, $\langle D_i | \hat{H} | D_j \rangle$, is chosen with probability $p_{\text{gen.}}$. Then, a new particle is created (“spawned”) on $|D_i\rangle$, i.e. added to c_i after this Monte Carlo cycle, with probability $p_{\text{spawn}} = \delta\tau \frac{|\langle D_i | \hat{H} | D_j \rangle|}{p_{\text{gen.}}}$, where the importance sampling requires the division by $p_{\text{gen.}}$ to unbiased the result. If $p_{\text{spawn}} > 1$, $\lfloor p_{\text{spawn}} \rfloor$ particles are spawned and another with probability $p_{\text{spawn}} - \lfloor p_{\text{spawn}} \rfloor$. The sign of the

particle(s) spawned is the opposite sign of $\langle D_i | \hat{H} | D_j \rangle c_j$. The reason why the selected particle on c_k took the role of a part of c_j here and not of c_i , is to reduce spawns with zero spawn probability (which — considering equation 3.11 — would be the case if the connected determinant had no particles on it, whereas now an empty determinant can be made “active” by spawning a particle onto it). Initially, $p_{\text{gen.}}$ was a combination of uniform probabilities in the selection of orbitals that differ between the two determinants. Symmetry was considered as well to reduce selections with $\langle D_i | \hat{H} | D_j \rangle = 0$ [33, 34], or $\frac{|\langle D_i | \hat{H} | D_j \rangle|}{p_{\text{gen.}}}$ was brought as closely to a constant as possible [77, 78, 3] to increase efficiencies. Ref. [3] forms the content of chapter 5.

- *Death Step* (diagonal): This step deals with the diagonal element, $\langle D_i | \hat{H} - E_0 | D_i \rangle c_i$. E_0 is approximated by shift S described later. The probability of removing or cloning the selected particle is $p_{\text{death}} = \delta\tau |\langle D_i | \hat{H} - S | D_i \rangle|$. If $\langle D_i | \hat{H} - S | D_i \rangle > 0$, the particle is removed, otherwise it is cloned on the same determinant, after this Monte Carlo cycle.

After all particles have undergone the *spawn* and the *death* step, the Monte Carlo cycle ends with the clean-up step, the *annihilation* step, which has been shown to help overcome the sign problem and let the wavefunction converge to its final estimate [79];

- *Annihilation Step*: Here, the newly created particles are combined with the existing particles and the particles to die are removed. If there are two particles of opposite sign on the same determinant, they annihilate each other and so both get deleted.

The initial shift S is often set to $S = 0$. After a certain stable target particle population N has been reached (discussed later), S is varied at the end of every B Monte Carlo cycles. Eventually, S will converge to E_0 , the correlation energy of the ground state energy. Remember that E_{HF} was subtracted off the diagonal before the calculation

started. S is varied according to

$$S(\tau) = S(\tau - B\delta\tau) - \frac{\gamma}{B\delta\tau} \ln \frac{N(\tau)}{N(\tau - B\delta\tau)}. \quad (3.12)$$

The closer S is to E_0 , the fewer particles get cloned or removed during the *death* step and so the S stabilises and controls the particle population. At equilibrium, when data for analysis is taken, the population stays approximately constant, only varying due to stochastic noise. In addition to S , E_0 can also be estimated by the projected energy $E_{\text{proj.}}$,

$$E_{\text{proj.}} = \frac{\langle \mathbf{D}_0 | \hat{H} | \Psi \rangle_{\text{average}}}{\langle \mathbf{D}_0 | \Psi \rangle_{\text{average}}}, \quad (3.13)$$

where the Hartree–Fock energy had been subtracted of the diagonal of the Hamiltonian so that $E_{\text{proj.}}$ is also an estimate for the correlation energy. The averages are taken over various time steps in equilibrium using reblocking analysis[80] to account for autocorrelation, see section 3.5.2. In chapter 6, the notion of an instantaneous projected energy is also used, which is a quantity $\frac{\langle \mathbf{D}_0 | \hat{H} | \Psi \rangle}{\langle \mathbf{D}_0 | \Psi \rangle}$ at a particular time step. Taking its average as the energy estimate ignores the covariance between numerator and denominator and is therefore to be treated as an approximation. However, Blunt et al.[81], for example, commented that usually there is no evidence to suggest that the instantaneous projected energy differs significantly from the projected energy. Ichibha et al.[82] also introduced analysis methods that take the average of the instantaneous projected energy. The default in this thesis is not to use the instantaneous projected energy.

There has been significant further developments since the original paper[33]. The next subsection introduces the initiator approximation[35] after which further improvements to FCIQMC for the ground state energy are listed.

3.2.2 Initiator FCIQMC

Shortly after the first FCIQMC paper[33], Cleland et al.[35] developed the initiator approximation that is now extensively used in FCIQMC calculations as it enables tackling calculations in (even) larger systems.

The initiator approximation restricts the spawning from “less important” determinants and so decreases the space occupied by particles, focussing on the more significant determinants. Remember that the more states, i.e. determinants, are occupied, the higher the memory requirements are and the more particles there are in the system, the higher the computational cost is. Concretely, in HANDE[5] the *spawn* step is modified such that a spawn onto an unoccupied determinant is only successful if done by a particle that occupies an initiator determinant or — under certain conditions — if multiple spawn attempts are made to the same determinant (see HANDE code²). A determinant is an initiator determinant if it is occupied by more than $n_{\text{thres.}}$ particles at that time step. The user can set $n_{\text{thres.}}$, but the default used here was 3, see documentation³.

Due to this restriction, a bias in the energy is introduced. This bias can be reduced by increasing the number of particles or by decreasing $n_{\text{thres.}}$. When running an initiator calculation, a particular target population is chosen (no “shoulder” plot as discussed later is used) and the energy determined. This is repeated for several target populations, ideally until the energy has converged with respect to population. There are a few ways to then estimate the converged energy, they all require a graph of particle population against energy estimate. One technique, which is not used here, is to assume that as long as the energy estimates at the highest populations agree within error bars (ideally on a log scale to have sufficient distance), the energy estimate at the highest population can be chosen as the final estimate, see Ref. [83] that also considered N_0 . Another technique, used by Booth et al.[84, 54] and in chapter 6, is to fit a curve, such as an exponential or a polynomial, to all the data points and the infinite limit is the energy

²See <http://www.hande.org.uk/> and <https://github.com/hande-qmc/hande>.

³See <https://hande.readthedocs.io/en/latest/> for details.

estimate. Alternatively, another technique, used in chapter 4[2], is to fit a constant, horizontal line through the last few points that agree with each other.

3.2.3 Further Advances in FCIQMC

Another development for ground state FCIQMC applied here is the use of floating point weights on particles[85, 86]. Instead of holding an integer weight of 1 each, the particles can be assigned a floating point weight. When spawning, a new particle no longer created with probability p_{spawn} with weight ± 1 but instead created with probability 1 with weight p_{spawn} . However, to avoid creating particles on less significant determinants with very small weights, if p_{spawn} is below a threshold — the “spawn cutoff” — an original spawn attempt is made, i.e. a particle of weight ± 1 is created with a probability p_{spawn} . In HANDE[5], in the *annihilation* step, particles with absolute weight less than 1 are rounded down to 0 or up to 1 stochastically with a probability of its current weight. For a particle with weight less than 1, the decision whether it participates in the *spawn* step is done stochastically as well.

A further improvement to the algorithm is semi-stochastic FCIQMC[85, 87] which is usually activated after equilibrium has been reached to reduce sampling noise. The evolution of particles on the more populated determinants is done deterministically, while the remaining particles evolve stochastically as is the default. The size or type of space of deterministic determinants is chosen by the user. This feature has not been applied in any investigations described in this thesis.

Recently, a transcorrelated approach to FCIQMC which modifies the Hamiltonian to decrease problems due to electron cusps has been developed[88]. Other advances include a projector Monte Carlo technique which uses optimisation techniques employed by the machine learning community, such as adding “momentum” to improve convergence[89], or applying a perturbative correction to improve initiator convergence to the true energy[90]. They do not form part of this work.

Finally, as mentioned previously, work has been done in improving the *spawn* step[77, 78, 3] and the propagator[81, 4]. I have contributed to this development and so chapters 5 and 6 are dedicated to it.

3.3 Coupled Cluster Monte Carlo

Another method, related to FCIQMC, is coupled cluster Monte Carlo (CCMC)[39, 54], which uses the same propagator as FCIQMC and similar Monte Carlo steps to stochastically evaluate the coupled cluster ground state energies. Any coupled cluster level can be easily attempted by the user, e.g. CCSD or CCSDTQ5. However, perturbative levels, e.g. CCSD(T), are not implemented yet. Some developments made for FCIQMC also apply to CCMC and vice versa, such as floating point weights[85, 86], increased the efficiency of the *spawn* step[77, 78, 3] (see chapter 5) or improved the propagator[90, 4] (see chapter 6).

As outlined in the previous chapter, the estimate of the ground state Ψ_0 is created from the initial wavefunction as

$$|\Psi_0\rangle = t_0 \exp \left(\sum_{\mathbf{i} \neq \mathbf{0}} \frac{t_{\mathbf{i}}}{t_0} \hat{a}_{\mathbf{i}} \right) |\Psi_{\text{init.}}\rangle, \quad (3.14)$$

with normalisation t_0 , the amplitude on the reference determinant. For simplicity, this normalisation is ignored in the following algorithm outline. The fact that null excitors are not included in \hat{T} is made explicit. Even though a multireference approach exists[91], in this thesis, $|\Psi_{\text{init.}}\rangle = |\mathbf{D}_0\rangle$, i.e. Monte Carlo particles start on the reference excitor \hat{a}_0 and their distribution then spreads. Monte Carlo particles are located on the excitors $\hat{a}_{\mathbf{i}}$, not on determinants as with FCIQMC, with their sum on $\hat{a}_{\mathbf{i}}$ representing the corresponding amplitude $t_{\mathbf{i}}$. The unlinked coupled cluster equations, see chapter 2,

$$\langle \mathbf{D}_{\mathbf{i}} | \hat{H} - E | \Psi_0 \rangle = \langle \mathbf{D}_{\mathbf{i}} | (\hat{H} - E) \exp(\hat{T}) | \mathbf{D}_0 \rangle = 0, \quad (3.15)$$

are then solved iteratively for amplitudes $t_{\mathbf{i}}$, similarly to FCIQMC, ignoring quadratic or higher terms in $t_{\mathbf{i}}$ [54] and having assumed that the Hartree–Fock energy has already been subtracted off the Hamiltonian matrix diagonal,

$$t_{\mathbf{i}}(\tau + \delta\tau) = t_{\mathbf{i}}(\tau) - \delta\tau \langle \mathbf{D}_{\mathbf{i}} | \hat{H} - E | \Psi(\tau) \rangle, \quad (3.16)$$

Again, E is estimated by shift S . Franklin et al.[65] have modified this update equation to include the instantaneous projected energy $E_{\text{proj.,inst.}}$, which accelerates convergence,

$$t_{\mathbf{i}}(\tau + \delta\tau) = t_{\mathbf{i}}(\tau) - \delta\tau (\langle \mathbf{D}_{\mathbf{i}} | \hat{H} - E_{\text{proj.,inst.}} | \Psi(\tau) \rangle - (E_{\text{proj.,inst.}} - S)t_{\mathbf{i}}(\tau)), \quad (3.17)$$

where the last term scales all amplitudes equally and therefore should not bias the wavefunction estimate. This is discussed again in chapter 6 for the quasi-Newton acceleration.

Given that the selection step involves selecting a term in Ψ that is non-linear, it is more involved than the selection step in FCIQMC, where simply every particle was selected once in each Monte Carlo cycle. In CCMC, excitors and clusters, i.e. combinations of them, are selected, not particles. In the *spawn* or *death* steps, the probability to act is multiplied by the product of the amplitudes in the cluster. Clusters of excitors with a size up to the coupled cluster level plus two, i.e. a size of four when doing CCSD, can be selected. That is because the Hamiltonian matrix element $\langle \mathbf{D}_{\mathbf{i}} | \hat{H} | \mathbf{D}_{\mathbf{j}} \rangle = 0$ if $|\mathbf{D}_{\mathbf{i}}\rangle$ and $|\mathbf{D}_{\mathbf{j}}\rangle$ differ by more than two orbitals (see Slater–Condon rules in the previous chapter). A simple selection for CCSD for example used in the original CCMC algorithm[39] is to do as many selections as the sum of the amplitudes of the excitors, $N_{\text{tot.}}$. Cluster sizes up to a size below the maximum are chosen with probability $p_{\text{size}} = \frac{1}{2^{1+\text{size}}}$. Given that size 4 is the maximum size of a cluster, it is also chosen with chance $\frac{1}{16}$ — as size 3 is — so that p_{size} is normalised[54]. Using this decay in p_{size} , smaller cluster sizes are favoured. Note that excitors are unique in the excitation they create and that anti-commutation rules have to be followed when collapsing the action of a cluster on the Hartree–Fock determinant[1]. Within

the same cluster size, excitors are chosen with probability p_{exc} proportional to their amplitude[54]. A more advanced selection scheme is the *full non-composite* selection[1]. Clusters of size one are selected deterministically, while composite clusters, i.e. clusters of size greater than 1, are sampled stochastically. This implies that when stochastically selecting clusters, only composite clusters are selected. This reduces the minimum number of particles required for a stable calculation[1]. A further improvement on top of that is the *even selection* sampling[92]. As in the *full non-composite* algorithm, single excitors are selected deterministically. Selecting the size of the cluster (size greater than 1) and the kind of clusters is then done in a more involved weighted manner to evenly balance the selection over all possible clusters.

After the cluster, e.g. $\hat{a}_i\hat{a}_j$, has been selected, it is collapsed to a single excitor \hat{a}_m , such that $\hat{a}_i\hat{a}_j|D_0\rangle = \pm\hat{a}_m|D_0\rangle = \pm|D_m\rangle$. The sign has to be carefully determined using anti-commutation rules when combining the excitors. If $|D_m\rangle$ is further than coupled cluster level + 2 (e.g. 4 in the case of CCSD) excitations away from $|D_0\rangle$, the excitor is disregarded and a new cluster has to be chosen. This is prevented in *even selection* by choosing allowed clusters only[92]. The product of the particle amplitudes on \hat{a}_i and \hat{a}_j is A . The next steps are, similarly to FCIQMC,

- *Spawn Step*: From determinant $|D_m\rangle$, another determinant $|D_n\rangle$ is generated with probability $p_{\text{gen.}}$. If $|D_n\rangle$ is outside of the coupled cluster level space, e.g. further than a double excitation away from $|D_0\rangle$ in the case of CCSD, the spawn will not happen. Otherwise, with a probability $p_{\text{spawn}} = \delta\tau \frac{|A\langle D_n|\hat{H}|D_m\rangle|}{N_{\text{tot.}}p_{\text{size}}p_{\text{exc}}p_{\text{gen.}}}$, a particle is spawned onto excitor \hat{a}_m . If greater than 1 or if using floating point weights on particles, this is modified similarly to FCIQMC. The sign of the particle created is the opposite of the sign of $A\langle D_n|\hat{H}|D_m\rangle$. When not using *even selection*, the *multispawn* feature[1] reduces “blooms”, when multiple particles are spawned simultaneously, which make the energy estimate harder to analyse. When A is above a threshold multiple spawn attempts are made, which decreases p_{spawn} by increasing $p_{\text{gen.}}$. The same “excitation generators” as in FCIQMC can be employed to generate determinant $|D_n\rangle$ from $|D_m\rangle$.

- *Death Step*: In the full non-composites algorithm, for single excitors $p_{\text{death}} = \delta\tau \frac{|t_{\mathbf{m}}(\langle \mathbf{D}_{\mathbf{m}} | \hat{H} - E_{\text{proj.,inst.}} | \mathbf{D}_{\mathbf{m}} \rangle - (S - E_{\text{proj.,inst.}}))|}{N_{\text{tot.}} p_{\text{size}} p_{\text{exc}}} = \delta\tau \frac{|t_{\mathbf{m}} \langle \mathbf{D}_{\mathbf{m}} | \hat{H} - S | \mathbf{D}_{\mathbf{m}} \rangle|}{N_{\text{tot.}} p_{\text{size}} p_{\text{exc}}}$ and for composite clusters $p_{\text{death}} = \delta\tau \frac{|A \langle \mathbf{D}_{\mathbf{m}} | \hat{H} - E_{\text{proj.,inst.}} | \mathbf{D}_{\mathbf{m}} \rangle|}{N_{\text{tot.}} p_{\text{size}} p_{\text{exc}}}$. The particle created with probability p_{death} has the opposite sign of $A \langle \mathbf{D}_{\mathbf{m}} | \hat{H} - E_{\text{proj.,inst.}} | \mathbf{D}_{\mathbf{m}} \rangle$. Note that a particle cannot be killed, i.e. removed, as in FCIQMC but rather a particle of opposite sign gets created. This is because, in a composite cluster, the collapsed excitor $\hat{a}_{\mathbf{m}}$ might not have a particle on it that could be removed.

Again, after a Monte Carlo cycle, annihilation happens for particles of opposite sign on the same excitor. The initiator approximation has also been developed for CCMC[54] but has not been applied in this thesis so is not mentioned further here.

3.4 High Performance Computing – Parallelization

To enable calculations on large systems, it is crucial to be able to utilise high performance computing resources and so the algorithms should scale well with the number of cores. In this section, the FCIQMC parallelization[34] is briefly mentioned, followed by a more extensive discussion on parallel CCMC[1].

3.4.1 FCIQMC Parallelization

Booth et al.[34] developed an FCIQMC MPI parallelization. FCIQMC parallelises well and does not suffer from bias issues when having a high number of MPI processes to system size ratio. Each determinant is assigned to a particular MPI process via a hash function. As in the serial implementation, newly spawned particles are combined with the existing particles in the *annihilation* step. They are sent to the corresponding process using MPI communication. Particles created on previously empty determinants are assigned to an MPI process. To balance the work load on the MPI processes, the number of particles has to be considered[34, 5] and the communication can be optimised to non-blocking communication[5]

3.4.2 CCMC Parallelization

Spencer et al.[1] developed an MPI parallelization for CCMC. Due to the non-linearity of the wavefunction ansatz, the CCMC MPI parallelization does not scale as favourably as the FCIQMC equivalent and does suffer from energy bias issues if not applied sensibly. Thus, shared memory parallelization with OpenMP is used within the same node where possible as this does not create a bias in the energy estimate. In HANDE[5], the loop over the selection of different excitors and their action in the following Monte Carlo steps is OpenMP parallelized[1]. To cross nodes, MPI parallelization is employed[1]. As with determinants in FCIQMC, each excitor is assigned to an MPI process using a hash function. A bias arises since not every cluster of excitors can be created during the selection process, only those containing excitors on the same MPI process, as the other excitors are “not visible” at that point to that process. To limit the bias and to allow every cluster a chance of being formed at some point in the calculation, the excitors can move from one MPI process to another every $2^{\nu_{\text{move}}}$ iterations, for some integer ν_{move} . Again, at the *annihilation* step, spawned particles are moved to MPI processes with the excitors they belong to and combined with the particles on that process. Newly occupied excitors get assigned a process and it is decided whether excitors move process or not then. The MPI parallelisation is shown to scale well. It is important to get the ratio of system size to number of MPI processes right for efficiency and perhaps more importantly for bias purposes, see below[1].

The CCMC MPI parallelization energy bias has been investigated for the effects of various calculation settings[1]. A priori it can be expected that

- the greater the number of MPI processes the larger the bias since the proportion of excitors visible to a particular process decreases,
- the larger ν_{move} , the greater the bias since then the number of iterations taken for a currently impossible cluster to be created increases,
- and the larger the time step $\delta\tau$, the less fine grained the simulation is, allowing fewer clusters to be selected in a unit of imaginary time.

For a small system, the three dimensional spin non-polarised uniform electron gas (UEG) with 66 spinorbitals and $r_s = \{0.5, 5\}a_0$, was investigated for the effect of these three simulation parameters on the energy estimate, see figure 3.2, with CCSDT using full non-composite selection[1]. The Hilbert space size is only 22969, so this system is small. Figure 3.2 shows that, with the given settings of other parameters, the energy bias is visible but containable. As expected, the bias increased with the number of MPI processes, ν_{move} , and $\delta\tau$. Using hybrid parallelisation with OpenMP helps decreasing the bias by decreasing the number of MPI processes, see figure 3.2b. Note that using 240 MPI processes on this small system does not represent a realistic simulation. Ref. [1] investigates more realistic calculations as well. As long as there is a critical number of excitors (at equilibrium) per MPI process, this bias is not significant within error bars.


It was therefore concluded that while there is a bias in the energy, provided that not too many MPI processes are used, it should be within the error bars of the stochastic propagation and therefore not significant. OpenMP should be used as much as possible. For calculations that are not small and need multiple nodes, Ref. [1] recommends to not run on more than 1 MPI process per approximately 10^5 excitors occupied (at equilibrium).

3.5 Calculations and Data Analysis


Having outlined the CCMC and FCIQMC algorithms in the preceding sections, the details on how to conduct a calculation and how to analyse its output data are discussed here.

3.5.1 Running a Calculation


Before a calculation can start, the basis set and system have to be specified. Unless a model system, such as the UEG, is studied (where the integrals that are part of the evaluation of $\langle D_i | \hat{H} | D_j \rangle$ are evaluated on-the-fly), these integrals have to be pre-



(a) Effect of time step $\delta\tau$ on bias. 240 MPI processes and $\nu_{\text{move}} = 5$ were used.



(b) Effect of number of MPI processes on bias. $\delta\tau = 0.001$ and $\nu_{\text{move}} = 5$. Hybrid calculation used 20 MPI processes with 12 OpenMP threads each.



(c) Effect of move frequency ν_{move} on bias. 240 MPI processes and $\delta\tau = 0.001$ were used.

Fig. 3.2 Investigation of the effect of time step, number of MPI processes and move frequency on the CCMC parallelisation bias using the 3D spin non-polarised UEG at $r_s = \{0.5, 5\}a_0$ with 66 spinorbitals at CCSDT level. Published in Ref. [1] (Figure 3).

calculated and are written into a file and later read in by the program performing the CCMC or FCIQMC calculation. In this thesis, this was performed with PySCF[93] or Psi4[94, 95].

When setting up a calculation, the user has to specify various parameters such as

- time step $\delta\tau$, which should be chosen as high as possible before too many blooms occur and the calculation becomes unstable, which is particularly important for initiator calculations, where the calculation may converge to an answer far away from the true energy with too many blooms. Holmes et al.[77] have shown that $\delta\tau$ affects the calculation efficiency. To avoid blooms above a certain threshold, in FCIQMC, $\delta\tau \frac{|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|}{p_{\text{gen.}}}$ should be below a set limit⁴.
- initial population $N_0(\tau = 0)$ on the reference. If set too high, the required number of total particles, the “shoulder height” in a non-initiator calculation, increases[54]. If too low, the calculation might not be stable enough to reach equilibrium with the chosen time step.
- target population N_{target} where the shift variation starts. In a non-initiator calculation, it should be above the “shoulder height” population, which is discussed later in this section. Setting it too high is wasteful but an order of magnitude or two above the “shoulder height” might be sensible if possible[96]. In an initiator calculation, the calculation is run for multiple target populations and the value at infinite population is estimated, extrapolating out the initiator error.
- shift damping γ . Ref. [97] has shown that $\frac{\gamma}{B\delta\tau}$ can affect the population control bias. Charlie Scott has automated this for HANDE[5] where γ is set such that the variation in the shift follows the variation in the instantaneous projected energy[5].
- the number of Monte Carlo iterations B before the shift is updated again.

⁴See NECI code, https://github.com/ghb24/NECI_STABLE/. Discussions with Pablo López Ríos are gratefully acknowledged.

- the excitation generator to use. See chapter 5 for details and discussion.

It is also recommended that the user chooses to use floating point particle weights[85, 86].

In a non-initiator calculation, N_{target} has to be chosen such that the “shoulder height”, or the “plateau”, is below it. The shoulder plot shows when convergence to the ground state wavefunction has been achieved. An example plot is shown in figure 3.3. A “shoulder” plot is a plot of the total particle population $N_{\text{tot.}}$ against the ratio of the




Fig. 3.3 “Shoulder” plot of a CCSD CCMC calculation on the 3D spin non-polarised UEG at $r_s = 2a_0$ with 1030 spinorbitals (515 plane waves). The data in the figure is from Ref. [2] which forms part of chapter 4. The “shoulder height” is shown with a vertical line, with width twice its error (which is small) and also indicated by an arrow. They were determined as described in Ref. [54], where the mean of the ten points with largest $\frac{N_{\text{tot.}}}{N_0}$ and its standard error are taken to be the “shoulder height” and its error.

total population to the population on the reference N_0 . The explicit dependence on τ is not shown. The dynamics are thoroughly explained by Ref.[54]. In the beginning,

all particles are on the reference. Their distribution then spreads and both the total and the reference particle number increases. The total number of particles grows at a bigger rate until the “shoulder height” where the reference population number growth rates starts exceeding that of the total particle number. At this point, approximate convergence has been reached. At some point, after the shift variation has started, the population oscillates around a constant value due to the population control. On the “shoulder” plot, this is a “blob” where the curve seems to end.

In an initiator calculation, several calculations are run with different N_{target} to extrapolate out the initiator bias. Figure 3.4 shows a sample initiator curve. As




Fig. 3.4 Initiator convergence plot of a *i*FCIQMC calculation on the 3D spin non-polarised UEG at $r_s = 1a_0$ with 1030 spinorbitals (515 plane waves). The data in the figure is from the study in Ref. [2] which forms part of chapter 4. The estimated initiator energy is shown with a horizontal line, with width twice its error. It was estimated by fitting horizontal lines and choosing the best one as described in Ref. [2].

discussed previously, several techniques to determine the energy at infinite population exist; choosing the energy of the calculation with the highest population as long as it

agrees with (an)other calculation(s) with a significantly smaller number of particles, fitting constant horizontal lines and selecting the optimal one as done in the next chapter, or, as in Refs [84, 54] and in chapter 6, fitting a polynomial or exponential function.

A single calculation is run for long enough to be able to estimate the energy and its error. This analysis is discussed in the next section.

3.5.2 Analysing the Data

This section describes how to estimate the mean projected energy and shift and their standard errors. This is complicated by the fact that the data is autocorrelated. Every B Monte Carlo iterations, the instantaneous shift and instantaneous numerator and denominator of the projected energy are printed out. Figure 3.5 shows the behaviour of the shift and the instantaneous projected energy as the calculation progresses. The shift is set to zero first and varied once a target population has been reached, whereas the instantaneous projected energy starts converging immediately.

First, the approximate point of equilibrium is determined. This can be done simply by eye. However, to avoid any bias of the analyser and to speed up analysis, this has been automated. I have implemented a method in collaboration with other HANDE developers to automate the equilibrium finding⁵. It involves taking the iteration at which the error of the standard error estimate divided by the square root of the number of data points left is smallest and then removing another approximate autocorrelation length to ensure equilibrium has been reached. This is done for various printed quantities and the most conservative estimate is then chosen. For figure 3.5, this was between iterations 30000 and 40000 which might seem slightly conservative but the data seems certainly converged there.

The equilibrated data can then be analysed to determine the energy mean and its standard error. A common method to achieve this is the reblocking analysis[80]. It first treats all data points as independent data blocks and calculates the quantity

⁵Fruitful discussions with CDT cohort 1 also acknowledged.




Fig. 3.5 Energy estimators shift, S , and inst. projected energy, $\text{Inst. } E_{\text{proj.}}$, convergence with number of iterations for a CCSD CCMC calculation on the 3D spin non-polarised UEG at $r_s = 2a_0$ with 1030 spinorbitals (515 plane waves). The data in the figure is from Ref. [2] which forms part of chapter 4. The calculation has been run for more than the number of iterations shown. The iteration where the reblocking analysis[80] started was 33580.

mean and its error. Then, neighbouring points are grouped in pairs, combined to one data point and this is repeated. Then neighbouring combined pairs are grouped, and so on. The standard error first increases and then plateaus during this procedure. One of the standard errors at the plateau is then chosen as the true standard error. The mean does not significantly change during this process except for tiny variations due to having to exclude data points to match up the pairings. Other possible analysis methods are described by Ref. [82].

As discussed by Umrigar et al.[98] for DMC and considered for FCIQMC by Vigor et al.[97], population control due to the shift introduces a bias. They introduced a reweighting, post-analysis method that indicates whether a bias is present. The reweight plot allows assessment of whether the projected energy changes within error

bar for different levels of reweighting. If yes, a bias is present. More details are found in Ref. [97]. Figures showing data were drawn with Matplotlib[127] in this thesis.

Chapter 4

Converging Coupled Cluster Energies in the Uniform Electron Gas

The first investigation was how coupled cluster performs in a model solid system, the uniform electron gas, to gain information about what level of coupled cluster might be necessary when modelling more realistic solid systems. This chapter (slightly modified to fit into the format of this thesis and with updates since publication) has been published in

- V. A. Neufeld and A. J. W. Thom. *A study of the dense uniform electron gas with high orders of coupled cluster*, J. Chem. Phys. **147**, 194105 (2017)[2].

(Reproduced (although slightly adapted) from [2], with the permission of AIP Publishing. This article may be downloaded for personal use only. Any other use requires prior permission of the authors and AIP Publishing. This article appeared in Ref. [2] and may be found at <https://doi.org/10.1063/1.5003794>.)

Alex Thom came up with the very initial idea for this study. I conducted this investigation which we both designed in various discussions. I have written the paper/this chapter with edits from Alex Thom. The uniform electron gas capabilities

had already been added to the HANDE code[99, 5] by others. The reviewers also contributed with helpful comments.

4.1 Introduction

Various wavefunction methods have been applied to the three dimensional uniform electron gas (3D UEG)[66, 67, 69], e.g. the random phase approximation (RPA)[100–102] which yields a ground state energy that is equal to the energy output of a version of ring diagram coupled cluster doubles (CCD)[103, 104]. MP2 has been shown to diverge in the thermodynamic limit in the uniform electron gas whereas it is unclear whether coupled cluster singles and doubles (CCSD) does[105, 106]. There exists accurate ground state energy data for the high density regime based on the finite UEG with the FCIQMC[107, 36, 108] and DMC[109–117] methods. Versions of coupled cluster have been applied to the UEG in the thermodynamic limit, see e.g. [103, 118, 119]. CCSD and CCSDT have been applied to the finite 3D UEG[107, 106, 120, 54, 121, 122]. Shepherd[122] has extrapolated finite CCSD/CCD results in the 3D UEG to the thermodynamic limit and has compared them to Ceperley and Alder’s DMC energies[109] (see figure 2c in Ref. [122]). Using these DMC energies as a reference, the extrapolated CCSD correlation energy has an error of under 10% at $r_s = 1.0$ a_0 which increases to about 20% at $r_s = 5.0$ a_0 . Another recent study[54] has performed initiator and non-initiator stochastic coupled cluster in the CCSD and CCSDT levels on the dense 14 electron 3D UEG. The difference between CCSD and CCSDT was found to be significant even in the low correlation regime at $r_s < 1.0$ a_0 . r_s is the radius of a sphere that on average contains one electron. Here, we apply coupled cluster up to the CCSDTQ5 level which included quintuple excitations directly to the 14 electron non-spin-polarized UEG in the range $r_s = 0.5$ to 5.0 a_0 which is representative of some common simple solids (e.g. see Ref. [66]). We compare with (initiator) FCIQMC[108] and MP2[107] results. Using coupled cluster levels from CCSD to CCSDTQ5, we aim to answer the question what coupled cluster level is

needed to accurately model simple finite solids with certain densities, represented by the r_s parameter, with coupled cluster.

For our stochastic calculations, we have made use of development versions of the HANDE code¹. We have used the cluster multispawn feature[1] and the full non-composite cluster selection described in Ref. [1] using one MPI process divided up into OpenMP threads when running CCMC to avoid any risk of distributed memory parallel bias[1]. We have also run some FCIQMC calculations to compare our CCMC results to and we used the conventional and initiator versions for FCIQMC[33, 35] while only using non-initiator CCMC. The error bars of the data presented here were estimated by reblocking analysis[80] using pyblock² and the correlation energies are obtained from the projected energy. Errors were combined in quadrature. We found no significant population control bias using a reweighting scheme used in DMC[98] and adapted to FCIQMC[97].

4.2 Extrapolation to Complete Basis Set Limit

Coupled cluster singles and doubles (CCSD) is the least expensive level of coupled cluster. Owing to momentum and spin conservation, CCSD is equivalent to CCD in the UEG. At first, we extrapolated CCSD calculations to the complete basis set (CBS) limit for the 14 electron UEG. We then estimated the CBS limit of the other truncation levels studied by extrapolating energy differences between truncation levels and adding this to the CBS CCSD result. This is similar to the idea of focal point analysis as described in e.g. Ref. [123].

Shepherd et al.[107] have shown that for MP2, the correlation energy for a finite basis set with M spinorbitals goes as $1/M$ in the leading order for large M . They and other studies[36, 108, 124, 125, 122, 54] have used this trend and shown that it also holds reasonably well for CCSD and FCI(QMC). These studies have usually excluded points with larger $1/M$ that were no longer in the region in which $1/M$ is a good fit.

¹See Ref. [99, 5] and <http://www.hande.org.uk/> for information and code.

²See <https://github.com/jsspencer/pyblock> for information and code.

In this study, we have decided to modify this approach to allow higher orders of $1/M$ to be considered as well. This accounts for the fact that $1/M$ is merely a leading order term and by adding higher orders we allow for correction terms to account for the part of the energy not accounted for by $1/M$. There are two aspects that need to be considered when choosing the best fit curve: What polynomial are we fitting, i.e. what is the highest order of $1/M$ to include, and how many points with high $1/M$ should be excluded from the fit?

Starting with the lowest order polynomial to fit ($1/M$ when fitting CCSD and a constant when fitting coupled cluster differences), we first fit all the data points and then start excluding points with lowest M . For each fit, we calculate χ^2 over number of degrees of freedom #d.o.f.. $\chi^2 = \sum_i \left(\frac{f(x_i) - y_i}{\sigma_i} \right)^2$ where y_i is a data value, $f(x_i)$ is its fitted value and σ_i is the standard deviation of y_i [126]³. As soon as we reach a local minimum in the $\chi^2/\text{\#d.o.f.}$ value, we stop removing points and note down the value at $1/M = 0$ given by the fit at the local minimum. If no local minimum can be found before there are as few data points left as the number of fitting parameters, then the search for a best fit for the first polynomial was unsuccessful. We then repeat this procedure of consecutively removing data points with the next order polynomials, initially starting with a full set of data points again. We fit linear, quadratic and cubic polynomials and a constant as well if we are fitting to differences. Finally, we compare the results of the fits at local minima in the number of points at $1/M = 0$. If the lowest order fit result agrees with the higher order ones within 2σ , we accept it as the CBS result. If it does not agree with all the higher ones, we compare the second lowest order fit result to its higher order fit results, etc. This process can continue up comparing the CBS results from the highest two polynomials. If there is still no CBS result at the end, then the extrapolation was not successful and a CBS value has to be estimated (see results section for individual cases).

As an example, figure 4.1 shows the best fits with the lowest $\chi^2/\text{\#d.o.f.}$ for $r_s = 0.5$ a_0 CCSD and 14 electrons. The linear and the quadratic fit intercepts do not agree

³Discussions with Pablo López Ríos regarding this section and his idea of using $\chi^2/\text{\#d.o.f.}$ gratefully acknowledged.




Fig. 4.1 Extrapolating correlation energy against $1/M$ for $r_s = 0.5$ a_0 CCSD and 14 electrons with the best fit linear line (yellow, dashed with points, excluding three data points) giving $b_0 = -0.58866(5) E_h$, best quadratic fit (green, solid line, excluding two data points) with $b_0 = -0.58850(6) E_h$ and best cubic fit (red, long dashes, excluding one data point), giving $b_0 = -0.58848(7) E_h$. The CBS limit is then taken to be $-0.58850(6) E_h$ from the quadratic fit, as the linear fit and the quadratic fit do not agree within 2σ whereas the quadratic and cubic fits agree within 2σ . The CBS result is shown with a light blue horizontal line that has a thickness of twice its error.

within 2σ . The quadratic and cubic fits agree which meant that we took the quadratic fit intercept as the CBS result. We have used the `curve_fit` function in the SciPy⁴ optimize module for curve fitting and Matplotlib[127] for plotting. The standard errors of the correlation energy were treated as absolute and not relative weights.


⁴E. Jones, T. Oliphant, P. Peterson et al., *SciPy: Open source scientific tools for Python*. See <https://www.scipy.org/> for more information.

4.3 Results


Figure 4.2 shows how the differences in correlation energy between consecutive coupled cluster levels vary with r_s for different numbers of spinorbitals M . As a reference, an accuracy of 0.01 eV/electron = 0.00037 E_h /electron is shown with dashed horizontal lines. This is of a similar order of magnitude as chemical accuracy (ca. 0.04 eV/molecule[32]). To distinguish solid phases from each other, enthalpy differences of about 0.1 eV/atom often need to be resolved and at room temperature an accuracy of 0.01 eV in the energy is desired (see Ref. [128] for details). We have therefore chosen 0.01 eV/electron as a guide for energies to be of sufficient accuracy.

The CCSD to CCSDTQ5 CBS values are summarized in table 4.1. Note that while figure 4.2 quotes energies in energies per electron, table 4.1 shows energies for 14 electrons. First, the CCSD CBS value was found and then the CBS limit of differences between consecutive coupled cluster levels were added on to find the CBS limit of the other truncation levels. For r_s up to 2.0 a_0 , earlier CCSD and CCSDT results[54] are shown as well. MP2 results[107] and FCIQMC are given for comparison. For $r_s = 0.5$, 1.0, 2.0 and 5.0 a_0 FCIQMC values from Shepherd et al.[108] are given and additionally for $r_s = 0.5$ and 1.0 a_0 , new FCIQMC CBS results are presented for comparison. When using the initiator approximation[35], the FCIQMC correlation energies values for a certain number of spinorbitals M were estimated by fitting horizontal lines to energy against number of Monte Carlo particles curves, consecutively removing data points with the least number of particles. The energy at the global minimum in $\chi^2/\#d.o.f.$ when fitting a horizontal line is taken as the energy result. The error in the average number of particles was very small and therefore ignored. For the (*i*)FCIQMC results with $r_s = 0.5$ and 1.0 a_0 , the initiator approximation was used for M greater then 358 and 66 respectively. The initiator method was not used for CCMC calculations in this study. See footnote for more details⁵.


⁵It has been noted that when a calculation was restarted, the projected energy from the last iteration before the restart was not stored correctly (see https://hande.readthedocs.io/en/latest/release_notes/v1.2.html). The effects of this were tested by comparing the mean of the first half of the data of one output file with the mean of



(a) CCSD - CCSDT correlation energy difference



(b) CCSDT - CCSDTQ correlation energy difference



(c) CCSDTQ - CCSDTQ5 correlation energy difference

Fig. 4.2 Coupled cluster energy per electron differences at spinorbits $M = 66, 162, 358, 1030$. The dashed horizontal lines show an accuracy of 0.01 eV/electron.

Table 4.1 Summary of complete basis set extrapolated results for the correlation energy of the 14 electron uniform electron gas in E_h .

	$r_s = 0.5 a_0$	$r_s = 1.0 a_0$	$r_s = 2.0 a_0$	$r_s = 3.0 a_0$	$r_s = 5.0 a_0$
CCSD	$-0.58850(6)/-0.5897(1)^a$	$-0.51450(9)/-0.5155(3)^{1b}$	$-0.4096(10)/-0.4094(1)^1$	$-0.3395(1)$	$-0.2531(3)$
CCSDT	$-0.59457(7)/-0.5965(2)^1$	$-0.5307(2)/-0.5317(3)^1$	$-0.4407(10)/-0.4354(4)^1$	$-0.3780(3)^c$	$-0.2970(4)^d$
CCSDTQ	$-0.59465(8)$	$-0.5311(2)$	$-0.4432(10)^e$	$-0.3833(3)^{35}$	$-0.3015(4)^{4f}$
CCSDTQ5	$-0.5947(2)^5$	$-0.5311(2)^6$	$-0.4434(10)^{56}$	$-0.3837(3)^{356}$	$-0.3025(4)^{46}$
(i) FCIQMC	$-0.59467(9)^g/-0.5969(3)^h$	$-0.5313(2)^i/-0.5325(4)^8$	$-0.4447(4)^8$	$-0.306(1)^8$	
(i) FCIQMC-TC j	$(-0.5948(2))$	$(-0.5309(2))$	$(-0.4440(3))$	$(-0.3078(3))$	
MP2	$-0.575442(1)^k$	$-0.499338(2)^{10}$	$-0.398948(2)^{10}$		$-0.255664(4)^{10}$

^aThis (initiator) CCSD/CCSDT value is from Spencer et al.[54]

^bAlso compare to $-0.5152(5)$ from figure 7 in Shepherd et al.[107] as quoted by Spencer et al.[54]

^cThe CCSDT to CCSD energy difference for $r_s = 3.0 a_0$ was estimated by the mean of a constant, linear, quadratic and cubic fit with lowest $\chi^2/\#d.o.f.$ if multiple fits were available.

^dThe CCSDT to CCSD energy difference for $r_s = 5.0 a_0$ was estimated by the mean of a constant, linear and quadratic fit with lowest $\chi^2/\#d.o.f.$ if multiple fits were available.

^eThe CCSDTQ to CCSDT difference for $r_s = 2.0$ and $3.0 a_0$ and the CCSDTQ5 to CCSDTQ difference for $r_s = 0.5 a_0$ was estimated by the mean of a linear fit and the data point with lowest $1/M$.

^fThe CCSDTQ to CCSDT difference for $r_s = 5.0 a_0$ and the CCSDTQ5 to CCSDTQ difference for $r_s = 1.0, 2.0, 3.0$ and $5.0 a_0$ was estimated by the CCSDTQ to CCSDT difference at 66 spinorbitals.

^g (i) FCIQMC value of $r_s = 0.5 a_0$ was estimated by the CCSDTQ value plus the difference of CCSDTQ to (i) FCIQMC extrapolated value.

^hThis i FCIQMC data is from Shepherd et al.[108]

ⁱ (i) FCIQMC value of $r_s = 1.0 a_0$ was estimated by the CCSDT value plus the difference of CCSDT to (i) FCIQMC extrapolated value. Note that if dropping an i FCIQMC data point that was run at a different time step, this becomes $-0.5312(2)$, which agrees with the quoted result here.

^jTranscorrelated FCIQMC, data from Ref. [129]. Results in parentheses since data was published after this work.

^kThe MP2 data is from Shepherd et al. [107]

4.4 Discussion

Figure 4.2a shows that CCSD gives an accuracy worse than 0.01 eV/electron for r_s greater than $0.5 a_0$ as the difference between CCSD and CCSDT is greater than 0.01 eV/electron. Considering figure 4.2b, CCSDT seems to be sufficient up to $r_s = 2.0 a_0$. As the differences in correlation energy increase in magnitude with M and the $M = 162$ energy for $r_s = 3.0 a_0$ is close to 0.01 eV/electron, one should be cautious about using CCSDT for $r_s = 3.0 a_0$. Figure 4.2c shows that the difference between CCSDTQ and CCSDTQ5 is not negligible for r_s greater than $2.0 a_0$.

Of course, this analysis implicitly assumes that the energy is monotonically decreasing with coupled cluster level. If the difference to the next excitation level is bigger than 0.01 eV/electron, we expect the difference to the true energy also to be greater than 0.01 eV/electron. However, we found that in our case, the energy was monotonically decreasing and the CCSDTQ5 result agrees very well with FCIQMC, see table 4.1. This supports our approach of comparing the energy difference to the next excitation level when assessing accuracies.

Figure 4.3 shows the difference in correlation energy found with CCSD, CCSDT and CCSDTQ to the correlation energy found with CCSDTQ5 as a fraction of the CCSDTQ5 correlation energy. Given that the CCSDTQ5 energy shown in table 4.1 is merely a lower bound for the true magnitude of the CCSDTQ5 energy, the errors presented here are also lower bounds. The error in CCSD is at least 16% for $r_s = 5.0 a_0$ and for CCSDT it is still as big as about 2%. The error of CCSDTQ is small but noticable for $r_s = 5.0 a_0$. This means that for a study of a solid with $r_s \approx 4 a_0$ say, e.g. sodium, CCSD may give a correlation energy that is off by over 12% and the error with CCSDT is still over 1%. As the energy differences between coupled cluster levels increase with r_s , properties such as the lattice parameter or the bulk modulus will be underestimated by low orders of coupled cluster.

the second half. It was found that the effect is not significant as for most of the data tested the two means usually agreed with the energy stated (considering standard errors), see README in <https://doi.org/10.17863/CAM.14336> for more details.




Fig. 4.3 Fractional difference of CCSD, CCSDT and CCSDTQ correlation energies to the CCSDTQ5 correlation energy as a function of r_s . Some coupled cluster correlation energies were estimated as described in table 4.1.

As Shepherd et al.[107] already noted, for low r_s , MP2 performs worse than CCSD and vice versa for higher r_s in the regime studied (see table 4.1). MP2 gives a less accurate answer than CCSDT and higher truncation levels for all studied r_s .

We present new extrapolated FCIQMC results for $r_s = 0.5$ and $1.0 a_0$, which are similar to but do not agree with Shepherd et al.'s[108] values. Similarly, our CCSD and CCSDT values for $r_s = 0.5$ and $1.0 a_0$ do not agree within 2σ with Spencer et al.'s[54] values. Our CBS correlation energies are less negative. We can explain these deviations by considering the shape of the extrapolation curves such as figure 4.1. Our CCSD calculations went up to 18342/11150 spinorbitals for $r_s = 0.5/1.0 a_0$ and that was our starting point to extrapolate higher truncations and FCIQMC from. Shepherd et al.[108] and Spencer et al.[54] only considered M up to 4218 at most. If fewer data points with low $1/M$ are present and a linear fit is employed (as Shepherd et al.[108]

and Spencer et al.[54] did), the intercept with the y axis, the CBS energy estimate, will be more negative than in the case where lower $1/M$ are present and higher fits are allowed. Our FCIQMC values quoted in table 4.1 were found by extrapolating the difference between the CCSDTQ/CCSDT and the FCIQMC values for $r_s = 0.5/1.0$ a_0 as CCSDTQ/CCSDT was the highest coupled cluster data set that contained the highest M used in our FCIQMC study for $r_s = 0.5/1.0$ a_0 respectively. Had we instead extrapolated FCIQMC directly, the results would have been $-0.59497(4)$ E_h (instead of $-0.59467(9)$ E_h) with a linear fit for $r_s = 0.5$. For this direct fit we included spinorbitals up to $M = 4218$ and when we extrapolated differences, we used information from the CCSD result with spinorbitals up to 18342. This shows that it is crucial to include large numbers of virtual orbitals to converge to the correct answer. We believe that the disagreement of the CCSD and CCSDT values for $r_s = 0.5$ and 1.0 a_0 with Spencer et al.'s[54] values may also be due to initiator energies that are not converged fully. We have not used the initiator approximation for coupled cluster data here.

After this study was published in 2017, Luo et al.[129] have compared new (*i*)FCIQMC results (using a transcorrelated method) with the results here, shown in parentheses under *i*FCIQMC-TC in table 4.1. They agree at least within 2σ with our results. In collaboration with others we have used the results in this work when benchmarking CCMC's parallelisation behaviour[1] (see chapter 3) or showing that density matrix quantum Monte Carlo[130–132] results tend to results presented here as the temperature tends to zero Kelvin[5]. Also, Lee et al.[133] compared a version of auxiliary-field quantum Monte Carlo[134, 135] results with finite coupled cluster results shown here. Additionally, Blunt[90] has compared finite initiator results augmented by a perturbative correction to the (*i*)FCIQMC results at 358 spinorbitals at $r_s = 0.5$ a_0 (agreement within 1σ) and $r_s = 1.0$ a_0 (agreement within 2σ) presented here.

4.5 Summary and Conclusions

We have shown that CCSD and CCSDT are limited for modelling finite solids that can be described by the 14 electron uniform electron gas with r_s greater than $2.0 a_0$. A comparison with CCSDTQ5 has shown that if an accuracy of 0.01 eV/electron is desired, CCSDT is required beyond $r_s = 0.5 a_0$ and CCSDTQ is worth considering beyond $r_s = 3.0 a_0$. At $r_s = 5.0 a_0$, CCSD only reproduces up to about 84% of the correlation energy and CCSDT up to about 98%.

This study has demonstrated that there can be a need for coupled cluster orders beyond CCSDT when modelling finite correlated solid-state systems. The next two chapters show algorithmic improvements which can make these future studies more feasible.

Chapter 5

Accelerating the Importance Sampling in the *Spawn* Step

This is the first of the two main method development chapters. Here, the *spawn* step was accelerated using an efficient excitation generator. This chapter (slightly modified to fit into the format of this thesis) has been published in

- V. A. Neufeld and A. J. W. Thom. *Exciting determinants in Quantum Monte Carlo: Loading the dice with fast, low memory weights*, J. Chem. Theor. Comput. **15**, 1, 127-140 (2019)[3].

(Reproduced in part with permission from [3]. Copyright 2018 American Chemical Society. ACS Articles on Request author-directed link: <http://pubs.acs.org/articlesonrequest/AOR-S8PS7M2bqqB3jTNrrjav> and doi: <https://doi.org/10.1021/acs.jctc.8b00844>)

Alex Thom implemented the very first version of the presented excitation generator which was heavily modified in this study. I conducted this investigation — including modifying/further code implementations — which we both discussed/reviewed. I have written the paper/this chapter with edits from Alex Thom. The reviewers also contributed with helpful comments.

5.1 Introduction

After studying the uniform electron gas, it became clear that CCSD might not be sufficient to study solid systems unless they are weakly correlated. Periodic solids can get expensive fast as adding a \mathbf{k} point will add the number of orbitals per \mathbf{k} point to the list of orbitals M . Going from a 111 to a 222 \mathbf{k} point grid, means 8 times more orbitals to deal with. It is important to have an efficient algorithm that uses as little memory as possible.

Here, we propose a change to the *spawn* step in the algorithm to use weighted excitations, inspired by the heat bath algorithm proposed by Holmes et al.[77], and Cauchy–Schwarz weights proposed by Smart et al.[78]. The method introduced here has a lower computational scaling in CCMC than the heat bath excitation generators and a significantly lower memory cost. The *spawn* part of the algorithm explores the space of possible determinants/excitors. For a given determinant, it decides how the determinants connected via the hamiltonian become involved in the wavefunction by becoming occupied. As Holmes et al.[77] already noted, it is not efficient to give all connected determinants/excitors an equal probability of being considered as some are more important for the dynamics than others. They have shown that their heat bath weighting when selecting states to spawn to can greatly improve the overall efficiency.

Remember that a Slater determinant $|\mathbf{D}_{\mathbf{m}}\rangle$ is connected to another determinant $|\mathbf{D}_{\mathbf{n}}\rangle$ by their connecting Hamiltonian element $\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle$ as part of the *spawn* step and the algorithms to choose $|\mathbf{D}_{\mathbf{n}}\rangle$ given $|\mathbf{D}_{\mathbf{m}}\rangle$ are called excitation generators. The probability of this generation is denoted $p(\mathbf{n}|\mathbf{m}) = p_{\text{gen}}$ after which a spawn occurs with probability $p_{\text{spawn}} \propto \delta\tau \frac{|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|}{p_{\text{gen}}}$, with time step $\delta\tau$.

For an efficient calculation, p_{spawn} should have a reasonable value. If $p_{\text{spawn}} > 1$, multiple particles are spawned at the same time, known as a “bloom”, which is undesirable as it leads to less controllable population dynamics. If, however, p_{spawn} is small, determinants are selected with no resulting spawn, and the algorithm is inefficient. p_{spawn} therefore ideally has a constant value, which can be altered by the time step $\delta\tau$. Hence, it is desirable to make p_{gen} proportional to $|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|$ rather

than selecting determinants uniformly. Holmes et al.[77] have introduced a heat bath sampling algorithm which weights the probability of choosing $|\mathbf{D}_{\mathbf{n}}\rangle$ with approximately $|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|$, but requires pre-computation of Hamiltonian elements leading to a significant storage cost of $\mathcal{O}(M^4)$ (which is of the same order as stored integrals if the code does not calculate them on-the-fly) and computational cost of $\mathcal{O}(N)$ where M and N are the size of the basis set and number of electrons respectively. Smart et al.[78] have reported the use of the Cauchy–Schwarz-like inequalities to provide upper bounds for $|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|$ with weights calculated on-the-fly. This reduces the storage cost while being linearly scaling in the number of orbitals.

Inspired by these ideas, excitation generators were investigated with weights generated on-the-fly using Cauchy–Schwarz and Power–Pitzer[136] inequalities to approximate $|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|$ whose computational cost scales linearly with the number of spinorbitals in the basis, M . We then present a new excitation generator that uses this Power–Pitzer inequality but is of low computational order, $\mathcal{O}(N_{\text{ex.}})$ in the case of CCMC or $\mathcal{O}(N)$ when using FCIQMC, with memory cost $\mathcal{O}(M^2)$ which is also below the heat bath memory scaling. $N_{\text{ex.}}$ for a determinant or excitor is the number of electrons excited with respect to the reference. For a truncated coupled cluster theory $N_{\text{ex.}}$ does not scale with system size¹. In a single-reference calculation, the reference determinant carries the most weight in the wavefunction and the majority of spawnings occur from determinants within a few electrons of excitation of this. We therefore may pre-compute excitation weightings based on the reference determinant, which shares the majority of electrons with nearby excited determinants, and then map the excitation to apply to any excited determinant, $|\mathbf{D}_{\mathbf{n}}\rangle$. By this method, similar weights to the heat bath algorithm and weights inspired by a Power–Pitzer inequality are employed and the spread in $\frac{|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|}{p_{\text{gen}}}$ is minimised at a reduced computational and memory cost.

Rather than integer-valued, real-valued excip amplitudes[85, 86] have been used and the full non-composite version of the CCMC algorithm[1] with truncated and even

¹We note that in this suboptimal implementation here, this new excitation generator does not scale as $\mathcal{O}(N_{\text{ex.}})$ with CCMC.

selection[92] has been applied. We have varied the shift damping automatically to reduce the variance of the projected energy². We have also used MPI and OpenMP parallelization[1]. The results here were checked for population control bias using a reweighting scheme by Umrigar et al.[98] and Vigor et al.[97]. Data has been reblocked[80] implemented in pyblock³ to estimate error bars. Our CCMC and FCIQMC calculations were done with the HANDE code[99, 5] which is open source⁴. The integral files needed were created with PySCF[93]. When applicable, localization has been applied using a Boys[137] localization function in PySCF[93].

5.2 Excitation Generators

As mentioned above, in the *spawn* step, the excitation generator selects a determinant $|D_n\rangle$ connected to $|D_m\rangle$ with probability p_{gen} . The spawn probability is proportional to $\frac{\delta\tau|\langle D_n|\hat{H}|D_m\rangle|}{p_{\text{gen}}}$. In this paper, we present a method that aims to use an optimal p_{gen} so that more important determinants are selected with a higher probability. An introduction to excitation generators in FCIQMC is given by Booth et al.[33, 34]. The idea of excitation generation and dividing by the generation probability was also discussed in e.g. Refs[77, 138–141] and a transition with uniform selection is also done by the configuration state function projector Monte Carlo method of Ohtsuka et al.[38] or by the Monte Carlo configuration interaction by Greer[37]. Kolodrubetz et al.[140] used a weighted excitation generator that – among other distributions – used the inverse momentum squared as a weight. Booth et al.[34] also considered weighting the excitation generation by Hamiltonian matrix elements by enumerating a subset of excitations with the magnitudes of these Hamiltonian elements. Due to the cost of finding p_{gen} , this idea was not pursued further. A version of the uniform excitation generators described here, is explained in detail in Ref. [34].

²This feature has been implemented by Charles Scott.

³For code, see <https://github.com/jsspencer/pyblock>

⁴See <http://www.hande.org.uk/> and <https://github.com/hande-qmc/hande> for information and code

The spawn probability is only non-zero if $\langle D_{\mathbf{n}} | \hat{H} | D_{\mathbf{m}} \rangle$ is non-zero. The Hamiltonians, \hat{H} , considered here only contain constant, one body, and two body terms. $\langle D_{\mathbf{n}} | \hat{H} | D_{\mathbf{m}} \rangle$ can therefore only be non-zero if $|D_{\mathbf{n}}\rangle$ and $|D_{\mathbf{m}}\rangle$ differ by at most two orbitals. To select a suitable $|D_{\mathbf{n}}\rangle$ for $|D_{\mathbf{m}}\rangle$ to spawn to, we can create a single or a double excitation from $|D_{\mathbf{m}}\rangle$ to generate $|D_{\mathbf{n}}\rangle$ ($\mathbf{n} \neq \mathbf{m}$). Any other excitation would lead to a zero spawn probability. Except for the “original” heat bath excitation generator, all excitation generators discussed here create a single or double excitation from $|D_{\mathbf{m}}\rangle$ to generate $|D_{\mathbf{n}}\rangle$ with probability p_{single} or $1 - p_{\text{single}}$ respectively. As suggested by Holmes et al.[77], we aim to approximately select $p_{\text{spawn,single}}$ and $p_{\text{spawn,double}}$ by setting p_{single} , such that the distribution of excitations is as best as possible. For a single excitation where electron in spinorbital i is excited to spinorbital a ,

$$p_{\text{gen,single}} = p_{\text{single}} p_{\text{method}} p(i) p(a|i). \quad (5.1)$$

p_{method} contains additional factors depending on the selection method of i and a .

In the case of a double excitation, $ij \rightarrow ab$, as i and j ideally come from the same set of orbitals (those occupied in the determinant) and so do a and b (those unoccupied in the determinant), first ij and then ab are selected in all excitation generators discussed here. That means that for example while the selection order between i and j can vary, a will not be selected before either i and j . The possible orders are therefore $ijab$, $ijba$, $jiab$ and $jiba$. While the first selected occupied is called i and the second j , their indistinguishability has to be taken into account when calculating p_{gen} :

$$\begin{aligned} p_{\text{gen,double}} = & \\ (1 - p_{\text{single}}) p_{\text{method}} & (p(i)p(j|i)p(a|i,j)p(b|a,i,j) + \\ & p(i)p(j|i)p(b|i,j)p(a|b,i,j) + \\ & p(j)p(i|j)p(a|j,i)p(b|a,j,i) + \\ & p(j)p(i|j)p(b|j,i)p(a|b,j,i)). \end{aligned} \quad (5.2)$$

In a rather basic implementation, the spinorbitals with electrons to excite i (and j) and the spinorbitals to excite to a (and b) are selected with uniform probabilities. The excitation generator that we call *not renormalised excitation generator* or simply *no. renorm.* here, when doing a double excitation, first selects i and j as a pair with uniform probability from the set of occupied orbitals. In that case,

$$p_{\text{method}}(p(i)p(j|i) + p(j)p(i|j)) = \frac{2}{N(N-1)}, \quad (5.3)$$

where the number of electron is N . If both i and j have the same spin, σ , then a is uniformly chosen from the set of virtual orbitals of that spin, otherwise it can be any virtual orbital. b is then selected uniformly from the set of orbitals (excluding a) with required spin and symmetry. With this selection of b , it is possible that after the selection of i , j , and a , there are no possible selections of b , or that in fact an occupied orbitals has been selected as b . In such cases, it is a forbidden excitation generation. In that case the spawn attempt will be unsuccessful (we set $|\langle \mathbf{D}_{\mathbf{m}} | H | \mathbf{D}_{\mathbf{n}} \rangle| = 0$).

The choice of how to select which electrons to excite and to which spinorbitals they are excited is entirely arbitrary (assuming all valid excitations are possible), as long as the probability with which this selection has been done is known and p_{gen} is then calculated accordingly. As an alternative to the *not renormalised excitation generator* (*no. renorm.*), most forbidden excitations (which lead to unsuccessful spawns) can be avoided by generating a different excitation and renormalising the appropriate probabilities. This is called the *renormalised excitation generator* or in short, *renorm.*. Again, see Booth et al.[33, 34] for an in-depth description of uniform excitation generators.

In the following subsections, we describe the heat bath excitation generators and the *heat bath/uniform Power-Pitzer* excitation generators which follow the ideas of Alavi and others. Finally, the *heat bath Power-Pitzer ref.* excitation generator is presented, which pre-computes some weights based on the reference determinant and therefore has a very low computational cost not scaling with system size ($\mathcal{O}(N_{\text{ex}})$) when

using CCMC or scaling as $\mathcal{O}(N)$ for FCIQMC instead of $\mathcal{O}(M)$). Its memory cost is significantly less than *heat bath* excitation generators, being $\mathcal{O}(M^2)$ instead of $\mathcal{O}(M^4)$.

Further uniform excitation generators consider the relationship of the spin of ij in more detail. In the case of a double excitation, Hamiltonian matrix elements tend to be bigger if i and j do not have parallel spins. This is because following Slater-Condon rules[57, 58], the Hamiltonian matrix element is reduced to a sum of two terms of opposite sign in the case of parallel spins ($\langle ij|ab\rangle - \langle ij|ba\rangle$) and one term if the spins are not parallel ($\langle ij|ab\rangle$). It might therefore be advisable to select anti-parallel spin electrons with a greater probability than parallel ij . Alavi, Booth and others[34]⁵ had the idea of determining whether spins are antiparallel or parallel first when selecting i and j . The *no. renorm. spin* and *renorm. spin* excitation generators are modifications of *no. renorm.* and *renorm.* excitation generators, where instead of finding i and j as a pair from the set of occupied orbitals, it is first decided whether they should have parallel spins or not. With probability p_{parallel} , ij are either selected as a pair from the set of occupied α (probability $\frac{N_\alpha}{N}$) or from the set of occupied β orbitals (probability $1 - \frac{N_\alpha}{N} = \frac{N_\beta}{N}$) where N_α and N_β are the number of α and β electrons respectively. This can lead to forbidden excitations followed by failed spawning attempts if there is only one electron of one type of spin. Here, p_{parallel} is set as the fraction of H_{ijab} where i and j have parallel spins.

Table 5.1 gives an overview over the weighted excitation generators presented here. This table should be understood together with the following descriptions in the next subsections.

5.2.1 Heat Bath Excitation Generators

The heat bath excitation generators aim to get the orbital selection weights as close as possible to the Hamiltonian matrix element $|\langle D_{\mathbf{n}} | \hat{H} | D_{\mathbf{m}} \rangle|$ with the aim of making part of the spawn probability $\frac{|\langle D_{\mathbf{n}} | \hat{H} | D_{\mathbf{m}} \rangle|}{p_{\text{gen}}}$ as close as possible to a constant. In the case

⁵Personal Communication with Ali Alavi and Pablo López Ríos. This is also implemented in NECI https://github.com/ghb24/NECI_STABLE.

Table 5.1 Overview of weighted excitation generators. C.-S. means Cauchy-Schwarz and P.-P. Power-Pitzer. p.c. is pre-calculated and o.t.f. means on-the-fly. Comp./mem. \mathcal{O} is the computational/memory order the excitation generator scales with. As a method of selection, “heat bath” refers to “selecting those like the heat bath excitation generator”. Single excitations or ij in a double excitation that have been selected “uniformly”, have been selected with the uniform *renorm.* excitation generator. N is number of electrons, M the number of spinorbitals and N_{ex} the excitation level possible from the reference at this coupled cluster level.

	single excitations	ij	ab	comp. \mathcal{O}	mem. \mathcal{O}
<i>heat bath</i>	decision after having selected ija	heat bath	heat bath	N	M^4
<i>heat bath uniform singles</i>	uniformly	heat bath	heat bath	N	M^4
<i>heat bath exact singles</i>	exactly, on-the-fly	heat bath	heat bath	N^2M or $NN_{\text{ex}}M$	M^4
<i>uniform C.-S.</i>	uniformly	uniformly	C.-S. o.t.f.	M	M
<i>uniform P.-P.</i>	uniformly	uniformly	P.-P. o.t.f.	M	M
<i>heat bath C.-S.</i>	uniformly	heat bath	C.-S. o.t.f.	M	M^2
<i>heat bath P.-P.</i>	uniformly	heat bath	P.-P. o.t.f.	M	M^2
<i>heat bath P.-P. ref.</i>	p.c.	heat bath p.c.	P.-P. p.c.	N or N_{ex}	M^2

of a double excitation $ij \rightarrow ab$, p_{gen} can be rewritten as

$$p_{\text{gen,double}} = p(i) \times p(j|i) \times p(a|ij) \times p(b|ija) = \frac{\sum_{jab} H_{ijab}}{\sum_{ijab} H_{ijab}} \times \frac{\sum_{ab} H_{ijab}}{\sum_{jab} H_{ijab}} \times \frac{\sum_b H_{ijab}}{\sum_{ab} H_{ijab}} \times \frac{H_{ijab}}{\sum_b H_{ijab}}, \quad (5.4)$$

where $H_{ijab} = |\langle D_{\mathbf{n}} | \hat{H} | D_{\mathbf{m}} \rangle|$ where $|D_{\mathbf{m}}\rangle$ and $|D_{\mathbf{n}}\rangle$ differ by the excitation $ij \rightarrow ab$. In the case of heat bath excitation generators, $\frac{\sum_{jab} H_{ijab}}{\sum_{ijab} H_{ijab}}$ with certain limits in the sums is an approximation for $p(i)$ and so on.

Here, we distinguish between three different heat bath excitation generators described by/based on Holmes et al.[77]. The “original” heat bath excitation generator as introduced and described in detail by Holmes et al.[77] (in short *heat bath*), the heat bath excitation generator that decides first whether a single or a double excitation is performed and which samples singles uniformly which is mentioned by Holmes et al.[77] (*heat bath uniform singles*) and finally, the one that first decides whether to do a single or double excitation and samples singles exactly according to their Hamiltonian matrix element, *heat bath exact singles*⁶. For more information and an in-depth description, see Ref. [77].

In all three heat bath excitation generators, all possible contractions of H_{ijab} appearing in equation 5.4 are pre-computed and stored. More specifically, $H_i = \sum_{jab} H_{ijab}$, $H_{ij} = \sum_{ab} H_{ijab}$, $H_{ija} = \sum_b H_{ijab}$ and H_{ijab} are pre-computed where i, j, a and b can be any spinorbital in the calculation. In all sums $i \neq j \neq a \neq b$. The alias method[142–145, 77] is used and alias tables are pre-calculated for selecting a (given ij) with weights H_{ija} and one for selecting b (given ija) with weights H_{ijab} (which is of $\mathcal{O}(M^4)$). The look-up time with the alias method is of $\mathcal{O}(1)$. The alias tables for selecting i and selecting i given j are computed on-the-fly using pre-computed weights in $\mathcal{O}(N)$ time. The alias table for selecting i then only considers H_i from the set of

⁶Idea by Alavi and co-workers, this was suggested to us as an alternative by Pablo López Ríos (personal communication).

occupied orbitals for i and when selecting j given i , the alias table only considers H_{ij} with occupied j .

When using the *heat bath* excitation generator to find an excitation, first an alias table is created on-the-fly for i as described above and then i is selected. We proceed similarly for j . Using the pre-computed alias table with weights H_{ija} , a is found. If this orbital is occupied, we have a forbidden excitation and the spawn attempt was unsuccessful. Only at this stage it is decided whether to attempt a single or a double excitation. In the algorithm by Holmes et al.[77], a single excitation is attempted with probability $\frac{H_{ia}}{H_{ia}+H_{ija}}$ and a double excitation is attempted with probability $\frac{H_{ija}}{H_{ia}+H_{ija}}$ if $H_{ia} < H_{ija}$ where $H_{ia} = |\langle D_{\mathbf{k}} | \hat{H} | D_{\mathbf{m}} \rangle|$ with $|D_{\mathbf{m}}\rangle$ and $|D_{\mathbf{k}}\rangle$ connected by the excitation $i \rightarrow a$. However, if $H_{ia} > H_{ija}$, both a double and a single excitation are attempted⁷. In our implementation in HANDE[99, 5], that approach was modified to only allow one excitation attempt per excitation generator call. If $H_{ia} \geq H_{ija}$, instead of choosing to attempt a single ($i \rightarrow a$) and a double ($ij \rightarrow ab$) excitation, a single or a double excitation is attempted with probability $\frac{1}{2}$ respectively. The rest follows Holmes et al.[77]. Either a single excitation $i \rightarrow a$ is attempted now or b is selected from pre-computed weights and a double excitation $ij \rightarrow ab$ (provided b is not occupied) is attempted.

The *heat bath* excitation generator relies on single excitations being less significant. It has the major drawback in that it potentially has a bias if there exists no j to be selected after i and before a if $i \rightarrow a$ is valid. This is explained in more detail in Ref. [77]. Our conservative but robust test for bias as implemented in HANDE, counts the number of j for which $\sum_b H_{ijab}$ is non zero for given ia . If this number is greater than the number of virtual orbitals, then there will always be an occupied j to be selected for allowed $i \rightarrow a$ and there is no bias.

⁷It is not clear from Holmes et al.[77] what happens if $H_{ia} = H_{ija}$

5.2.2 On-the-fly Power–Pitzer Excitation Generators

While bringing $\frac{|\langle D_{\mathbf{n}} | \hat{H} | D_{\mathbf{m}} \rangle|}{p_{\text{gen}}}$ closer to a constant as uniform excitation generators[77], *heat bath* excitation generators suffer from a large memory cost ($\mathcal{O}(M^4)$). Alavi and Smart et al.[78] had the idea of calculating approximate weights on-the-fly in $\mathcal{O}(M)$ calculation time which has a lower memory cost. This is for example mentioned by Blunt et al.[146], Holmes et al.[77], Holmes[147] and Schwarz[148]. Alavi and Smart et al.[78] proposed calculating Cauchy–Schwarz-like upper bounds on the two body part of the Hamiltonian on-the-fly when doing a double excitation. Here, we also describe an excitation generator that uses an inequality derived by Power and Pitzer[136] instead. It effectively differs from *Cauchy–Schwarz* excitation generators as described here by the usage of exchange rather than Coulomb integrals. We note that the *Cauchy–Schwarz* excitation generators mentioned here might/do not quite replicate excitation generators of Alavi et al.⁸ which are yet to be fully published and so our description can only be different, see alternative (differing) descriptions by Schwarz[148] or Holmes[147].

Given that i, j, a and b are different, the only non-zero part of the Hamiltonian element $\langle D_{\mathbf{m}} | \hat{H} | D_{\mathbf{n}} \rangle$ in a double excitation are the Coulomb integral $\langle ij|ab \rangle$ and the exchange integral $\langle ij|ba \rangle$ according to Slater-Condon rules[57, 58]. Here, the notation

$$\langle ij|ab \rangle = \int \frac{\phi_i^*(\mathbf{r}_1)\phi_j^*(\mathbf{r}_2)\phi_a(\mathbf{r}_1)\phi_b(\mathbf{r}_2)d\mathbf{r}_1d\mathbf{r}_2}{|\mathbf{r}_1 - \mathbf{r}_2|}, \quad (5.5)$$

is used with one electron orbitals/spinorbitals ϕ that make up Slater determinants $|D_{\mathbf{x}}\rangle$. An example of such a weight used by Alavi and others for $ij \rightarrow ab$ is a Cauchy–Schwarz upper bound on $\langle ij|ab \rangle$ given by

$$\sqrt{|\langle ia|ia \rangle| |\langle jb|jb \rangle|} \geq |\langle ij|ab \rangle|. \quad (5.6)$$

The weights are such that a can be chosen (almost) independently of b and vice versa which makes the algorithm linear scaling in the number of spinorbitals. A

⁸Personal communication with Ali Alavi and Pablo López Ríos.

Power–Pitzer[136] inequality is

$$\sqrt{|\langle ia|ai\rangle| |\langle jb|bj\rangle|} \geq |\langle ij|ab\rangle|. \quad (5.7)$$

Exchange integrals are lower or equal in magnitude than Coulomb integrals (see e.g. Ref. [149]) which means that exchange integrals are the tighter upper bound for $|\langle ij|ab\rangle|$. The two body term in the Hamiltonian is $\langle ij|ab\rangle - \langle ij|ba\rangle$. When a and b have opposite spin, the two body term reduces to $\langle ij|ab\rangle$ and its Power–Pitzer upper bound is used as the weight. If a and b have the same spin, both orderings, ab and ba will generate the same excitation, and this is included in p_{gen} . This section gives a detailed description of the algorithm.

i and j can be selected uniformly or with the *heat bath* weightings producing a family of excitation generators. We denote by *uniform Cauchy–Schwarz* and *uniform Power–Pitzer* excitation generators which select them uniformly, like the *renorm.* excitation generator, and by *heat bath Cauchy–Schwarz* and *heat bath Power–Pitzer* those which select them as the *heat bath* excitation generators do with pre-calculated weights with memory cost of $\mathcal{O}(M^2)$ ⁹. The computational scaling is $\mathcal{O}(M)$ in both cases.

The *Power–Pitzer* and *Cauchy–Schwarz* excitation generators first decide whether to attempt a single or a double excitation according to p_{single} . For single excitations, the *renorm.* excitation generator is employed. When attempting double excitations, i and j are selected either uniformly or with *heat bath* weights out of the occupied orbitals of $|\mathbf{D}_{\mathbf{m}}\rangle$. Then, a is selected out of the set of virtual spinorbitals $a_{\sigma_i, \text{virt.}}$ with the same spin as i . a is selected with the probability of

$$p(a|ij) = p(a|i) = \frac{\sqrt{|\langle ia|ai\rangle|}}{\sum_{a=a_{\sigma_i, \text{virt.}}} \sqrt{|\langle ia|ai\rangle|}} \quad (5.8)$$

⁹The idea of selecting ij like the *heat bath* excitation generator was communicated by Pablo López Ríos (personal communication).

when using *Power–Pitzer* excitation generators or

$$p(a|ij) = p(a|i) = \frac{\sqrt{|\langle ia|ia \rangle|}}{\sum_{a=a_{\sigma_i, \text{virt.}}} \sqrt{|\langle ia|ia \rangle|}} \quad (5.9)$$

when using *Cauchy–Schwarz* excitation generators. b , the second orbital to excite to, is selected out of the set of spinorbitals $b_{\neq a, \sigma_j, \text{sym.}}$ of the same spin as j and the required symmetry to conserve overall symmetry and not equal to a . The weights are given by $\langle jb|jb \rangle$ (*Cauchy–Schwarz*) or $\langle jb|bj \rangle$ (*Power–Pitzer*). If the total weight when finding b is zero (i.e. there are no spinorbitals with the required spin and symmetry or only the spinorbitals found as a has that spin and symmetry) or if the found b is already occupied, the spawn attempt is unsuccessful. Again, orbitals a and b were selected using their weights with the alias method.

The performance of the four excitation generators described in this subsection, *uniform Cauchy–Schwarz*, *heat bath Cauchy–Schwarz*, *uniform Power–Pitzer*, and *heat bath Power–Pitzer*, were then tested, using a chain of three water molecules in the cc-pVDZ basis[63], whose molecular orbitals have been localized. The excitation generators all come with a low memory cost, which is $\mathcal{O}(M)$ temporarily or $\mathcal{O}(M^2)$ and all scale as $\mathcal{O}(M)$ in computational time. The distribution of $\frac{|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|}{p_{\text{gen}}}$, which should ideally be constant, was compared for the four excitation generators. Figure 5.1 shows the histograms (excluding $\frac{|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|}{p_{\text{gen}}} = 0$) with linear and logarithmic frequency scales. The *bottom* graph shows the all excitation generators have similar looking tails to both sides, the *heat bath Power–Pitzer* having the longest tail at big $\frac{|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|}{p_{\text{gen}}}$. However, the number of events in bins above the maximum $\frac{|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|}{p_{\text{gen}}}$ filled bin for the *uniform Power–Pitzer* excitation generator — which has the lowest maximum — is fewer than 100 events which is not significant relatively speaking so if not using initiator approximations there should not be a noticeable effect. The *top* graph demonstrates that the *heat bath Power–Pitzer* gives the sharpest peak and makes $\frac{|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|}{p_{\text{gen}}}$ closest to a constant of the excitation generators. Only non-zero allowed events are shown in figure 5.1. Table 5.2 shows what fraction that is of the total number of events (second

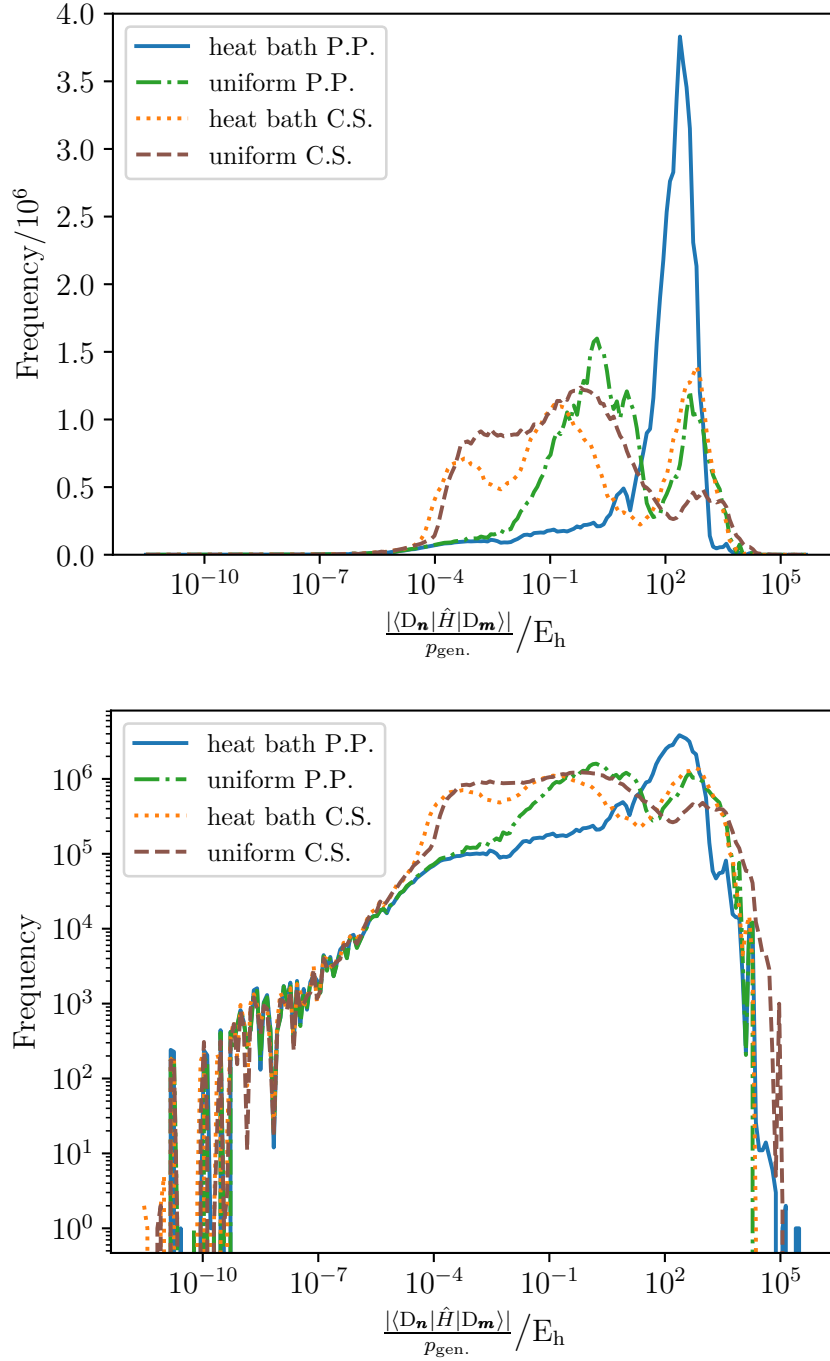


Fig. 5.1 Comparison of the histograms of $\frac{|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|}{p_{\text{gen.}}}$ for the *Cauchy-Schwarz* (C.S.) and *Power-Pitzer* (P.P.) on-the-fly excitation generators. ij are either selected uniformly or using heat bath. The bin middles on the $\frac{|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|}{p_{\text{gen.}}}$ axis are used for the data points. The computational scaling of all excit. gens. here is $\mathcal{O}(M)$. CCSD was performed on three water molecules in the cc-pVDZ basis using localized MOs. The values were logged for one MC iteration. The size of the bins is logarithmic. *Bottom* graph took the log of the frequency whereas the *top* graph did not. They both show the same data. All of them were restarted from the same calculation and then equilibrated before taking data. $\frac{|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|}{p_{\text{gen.}}} = 0$ data is not shown which includes forbidden excitations. p_{single} was set to be the same when running which was corrected in post-processing to make the mean of finite $\frac{|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|}{p_{\text{gen.}}}$ for single and double excitations approximately coincide which did not change p_{single} values by more than about 30%.

Table 5.2 Fraction of allowed and fraction of non-zero allowed spawn events, both with respect to total number of spawn events. The latter represents the spawn events depicted in figure 5.1. *heat bath Cauchy-Schwarz* and *uniform Cauchy-Schwarz* and the *Power-Pitzer* excitation generators have been combined to *C.-S.* and *P.-P.* respectively. Individual data points have been rounded to the second decimal place. If a range is given they rounded to either value in the range.

	$\frac{\# \text{allowed}}{\# \text{total}}$ events	$\frac{\# \text{allowed non-zero}}{\# \text{total}}$ events
<i>C.-S.</i>	0.80	0.69–0.72
<i>P.-P.</i>	0.68–0.69	0.68–0.69

column) and what fraction of events are allowed which includes the allowed but zero $\frac{|\langle D_{\mathbf{n}} | \hat{H} | D_{\mathbf{m}} \rangle|}{p_{\text{gen}}}$ events (first column). Both the *Cauchy-Schwarz* and the *Power-Pitzer* excitation generators have a similar fraction of non-zero allowed events. The *Power-Pitzer* excitation generators have more forbidden events but of those that are allowed, more are non-zero. A big source for forbidden events is the selection of b which is selected from the set of orbitals of required spin and symmetry which can be occupied. An event is then forbidden if b selected is occupied. Our implementation could be further improved by excluding occupied orbitals from that selection. In the results section we will let *heat bath Power-Pitzer* represent all these four excitation generators introduced in this subsection.

5.2.3 Pre-computed Power-Pitzer Excitation Generator

Even with their reduced memory requirements, the above excitation generators still add a considerable cost to calculations, and we seek a way to reduce this further. We now introduce an $\mathcal{O}(N)$ *Power-Pitzer* excitation generator, *heat bath Power-Pitzer ref.*, where N is the number of electrons. This can even be modified to be $\mathcal{O}(N_{\text{ex.}})$ where $N_{\text{ex.}}$ is the number of electrons excited with respect to the reference if excitations instead of determinants were stored in our implementation. Within a routine coupled cluster calculation, the maximum $N_{\text{ex.}}$ does not depend on system size. This excitation generator combines advantages of *heat bath Power-Pitzer* where a bias check is not required beforehand (but is with the “original” *heat bath* excitation generator) and

which has a significantly lower memory cost with those of the lower computational scaling of the *heat bath* excitation generators, further improving upon them. We make use of the single-reference nature of coupled cluster where the reference determinant $|D_0\rangle$ is more important than any other determinant by pre-computing some weights based on the reference determinant. Pre-computed weights include *heat bath* and *Power-Pitzer* weights, for selecting the orbitals to excite from and to in a double excitation. Spinorbitals are first found by pretending the reference determinant is the determinant we are exciting from and are then mapped between the current determinant and the reference determinant when it is appropriate. The memory cost is $\mathcal{O}(M^2)$ while the computational cost when spawning is only the mapping of the reference $|D_0\rangle$ to the actual determinant $|D_m\rangle$ which is $\mathcal{O}(N)$. Since weights are based on one determinant, it is not costly to pre-calculate weights for single excitations as well. This is a considerable advantage over the on-the-fly *Power-Pitzer* and *heat bath* excitation generators that either do single excitations uniformly, exactly (which is costly) or partly based on double excitation weights.

In this algorithm, two frames of reference are considered. In the first frame, the reference frame, which is denoted by a prime, excitations are from the reference determinant, i.e. $|D_{m'}\rangle = |D_0\rangle$. In this frame, a double excitation would be $i'j' \rightarrow a'b'$. In the second frame, the simulation frame, the actual frame the calculation is in, excitations are from $|D_m\rangle$ and that excitation is $ij \rightarrow ab$. For selecting some orbitals, the weights of the orbitals in the reference frame are used and its spinorbitals are mapped to the simulation frame to find the actual excitation as explained in the next paragraph.

In HANDE, there is a list of orbitals that are occupied in the reference, usually approximately ordered by one electron energies, and there is an equivalent ordered list with orbitals occupied in the current determinant $|D_m\rangle$. The localized orbitals here were ordered by approximate orbital energies, given by the expectation value of the Fock operator. Every time $|D_m\rangle$ is changed, two new (energy ordered) lists *RD* and *CD* are created, one (*RD*) containing all orbitals that are occupied in the reference




Fig. 5.2 Selecting i and j with *heat bath Power-Pitzer ref.* excitation generator for a double excitation. First i' is selected, occupied in the reference determinant $|D_0\rangle$ and translated to i , occupied in the current determinant, $|D_m\rangle$. i' and i are shown with dashed green circles. In this case, $i' = i$. Then j' is found and translated to j . As j' is not occupied in $|D_m\rangle$, it is mapped to the next orbital of the same spin occupied in $|D_m\rangle$ but not in $|D_0\rangle$. j' and j are shown with solid purple circles. Here $j' \neq j$. Drawn using Inkscape, <https://inkscape.org/> [Accessed: 11.12.2019].

but not in $|D_m\rangle$ and another list (CD) of the same size with all orbitals occupied in $|D_m\rangle$ but not the reference determinant. Orbitals with the same positions in these two lists are made to have the same spin by swapping orbitals in the list CD if necessary. If necessary, orbitals are translated by a one-to-one mapping between these two lists. If i' is not only occupied in the reference but in $|D_m\rangle$ as well, $i' = i$. If not, the position i' has in list RD is translated to the orbital with the same position in list CD . Figure 5.2 shows the translation of i and j in a double excitation in the two frames of reference pictorially. Note that this is the only part of the excitation generator that is not $\mathcal{O}(1)$ but $\mathcal{O}(N)$, mainly arising due to the creation of the two lists. The computational cost is reduced to $\mathcal{O}(1)$ if a determinant is reused. Alternatively, if, as mentioned previously, each excitor is not represented by a determinant but rather the lists RD or CD from the beginning the scaling is reduced to $\mathcal{O}(N_{\text{ex}})$ which is the cost of finding the correct mapping from one list to the other.

The following quantities for single excitations are pre-computed:

$$w_{i',s} = \sum_a \left(\frac{1}{n_{jb}} \sum_{j=j_{\text{occ.ref.}}, b=b_{\text{virt.ref.}}} \langle D_j^b | \hat{H} | D_{i'j}^{ab} \rangle \right), \quad (5.10)$$

where i' is an occupied orbital in the reference and the sum over a is over all orbitals with allowed excitation $i' \rightarrow a$. n_{jb} is $N(M - N)$. $|D_j^b\rangle$ differs from the reference determinant by the single excitation $j \rightarrow b$. We decided to not sum over single excitations from the reference as in the case of self-consistent field reference determinants, Brillouin's theorem would mean that the weights would be (close to) zero. We assume Brillouin's theorem when evaluating the weights. Assuming the system is single referenced, we might assume that a doubly excited determinant might be second most important after the reference determinant. The sum is therefore over all possible double excited determinants trying to connect to a determinant slightly closer to the reference via a single excitation. For selecting a ,

$$w_{a=a_{\sigma,\text{sym.}},i,s} = \frac{1}{n_{jb}} \sum_{j=j_{\text{occ.ref.}}, b=b_{\text{virt.ref.}}} \langle D_j^b | \hat{H} | D_{ij}^{ab} \rangle, \quad (5.11)$$

is pre-computed where i is now an occupied orbital in the current determinant which will have been selected before $w_{a=a_{\sigma,\text{sym.}},i,s}$ is needed. Given that the current determinant is not known at this stage, this is pre-computed for any orbital i . a is then selected from the orbitals of allowed spin and symmetry for which $i \rightarrow a$ is valid. Alias tables are then pre-computed for $w_{i',s}$ and $w_{a=a_{\sigma,\text{sym.}},i,s}$.

When running the excitation generator, it is first decided whether a single or double excitation is attempted with probability p_{single} or $1 - p_{\text{single}}$ respectively. If a single excitation was chosen, i' is first selected in the reference frame from the occupied orbitals in the reference using the alias table constructed with weights $w_{i',s}$. i' is then mapped to the corresponding occupied orbital in the current determinant i in the simulation frame.

Once i is known, a is selected using the pre-computed alias table with $w_{a=a_{\sigma,\text{sym.}},i,s}$. Of course, a could be occupied. If that is the case, the excitation attempt was unsuccessful. Otherwise, $i \rightarrow a$ is found and the generation probability is

$$p_{\text{single}} \times \frac{w_{i',s}}{\sum_{i'=i'_{\text{occ.ref.}}} w_{i',s}} \times \frac{w_{a=a_{\sigma,\text{sym.}},i,s}}{\sum_{a=a_{\sigma,\text{sym.}}} w_{a=a_{\sigma,\text{sym.}},i,s}}. \quad (5.12)$$

For double excitations, four weight tables are pre-computed. For the selection of i and j , heat bath weights are pre-computed, assuming the reference determinant is fully occupied. Two orbitals i' and j' occupied in the reference are found and then mapped to the actual determinant that is occupied. For the virtual orbitals a and b , alias tables based on Power–Pitzer weights are pre-calculated for *all* spinorbitals. Before selecting a , the actual i is known and can be substituted into pre-computed weights $\sqrt{|\langle ia|ai \rangle|}$ to find a . The memory cost is $\mathcal{O}(M^2)$. No mapping is necessary for a and b . Again, if a or b are occupied or b is equal to a or if there is not suitable orbital for b , the spawn attempt was unsuccessful.

Again, orbitals $i'j'$ are part of the reference frame, where the reference determinant is occupied, and ij are the equivalent spinorbitals in the actual frame, where the actual determinant we are exciting from is occupied. $i'j'$ are first found in the reference frame using heat bath weights and then they are mapped to the actual frame. ab are found with Power–Pitzer weights in the actual frame. All weights are pre-computed. This appendix describes the details of generating the double excitation. For i' , the pre-computed weights are

$$w_{i',d} = \sum_{j'=j_{\text{occ.ref.}}, \neq i', a \neq \{i', j'\}, b \neq \{i', j', a\}} H_{i'j'ab} \quad (5.13)$$

i' is selected from the set of occupied orbitals in the reference with a sum over j' , the set of occupied orbitals in the reference other than i' . a and b out of the set of all orbitals (not just virtual) are summed over, provided they don't equal i' , j' or each

other. For j' ,

$$w_{j'i,d} = \sum_{a \neq \{i,j'\}, b \neq \{i,j',a\}} H_{ij'ab} \quad (5.14)$$

is pre-calculated which is of memory scaling order $\mathcal{O}(NM)$. For both $w_{i',d}$ and $w_{j'i,d}$, a minimum weight is set in case the total weight for selecting i' or j' respectively in the reference frame is zero but selecting the equivalent i and j in the simulation frame would be allowed.

To select a and b , Power–Pitzer weights are pre-calculated. For a ,

$$w_{a,i,d} = \sqrt{|\langle ia|ai \rangle|} \quad (5.15)$$

where $w_{a,i,d}$ is zero if $i = a$. ia are from the set of all spinorbitals and a is restricted to the set of the same spin as i . The memory cost is simply $\mathcal{O}(M^2)$. Similarly, for b

$$w_{b,j,\text{sym},d} = \sqrt{|\langle jb|bj \rangle|} \quad (5.16)$$

where $w_{b,j,d} = 0$ if $b = j$ and b is from the set of all spinorbitals with the same spin as j . $w_{b,j,d}$ are arranged in such a way that b 's of the required symmetry later can readily be looked up. Alias tables for all these weights for single and double excitations are pre-computed.

In the case of a double excitation, first i' , an occupied orbital in the reference frame, is selected using $w_{i',d}$. $i' \rightarrow i$ is mapped to an occupied orbital i in the simulation frame if required. Then, j' is found using the pre-computed alias table for $w_{j'i,d}$ and map $j' \rightarrow j$ if needed. i and j are ordered so that j has a higher or equal index in the determinant list as i . Using i and $w_{a,i,d}$, a is found using pre-computed alias tables out of all spinorbitals with the same spin as i . If a is occupied, the spawn attempt was unsuccessful. The symmetry that b should have is then determined and using the pre-calculated alias tables for $w_{b,j,\text{sym},d}$ which give us a b of the correct symmetry (and spin), b is found from the set of all spinorbitals with required spin and symmetry.

Again, if b is occupied or equal to a or if there is not suitable orbital for b , the spawn attempt was unsuccessful.

Overall, this is an excitation generator that is both weighted and can scale as $\mathcal{O}(N_{\text{ex.}})$ in CCMC which does not scale with system size. In FCIQMC the scaling is still low, $\mathcal{O}(N)$. The memory cost is also relatively small, $\mathcal{O}(M^2)$.

5.3 Results and Discussion

To compare the effectiveness of the excitation generators discussed, water chains were then studied in a cc-pVDZ basis set[63] whose MOs have been localized. Figure 5.3 shows a histogram of $\frac{|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|}{p_{\text{gen}}}$ for three waters with the four uniform excitation generators, the *heat bath Power–Pitzer* excitation generator (which had the sharpest peak out of the $\mathcal{O}(M)$ /on-the-fly excitation generators), the *heat bath Power–Pitzer ref.* and the two *heat bath* excitation generators that do not suffer from bias. The “original” *heat bath* excitation generator was rejected by our bias test as it was not clear whether all allowed single excitations can be created. Considering a logarithmic scale in $\frac{|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|}{p_{\text{gen}}}$, the *top* graph in figure 5.3 clearly shows that the uniform excitation generators produce a bigger spread in $\frac{|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|}{p_{\text{gen}}}$ than weighted excitation generators (*Power–Pitzer* or *heat bath*).

The *heat bath* excitation generators produce the sharpest peak. The *heat bath uniform singles* excitation generator, that samples single excitations uniformly, shares the main peak with the *heat bath exact singles* excitation generator, that samples single excitations exactly, but has a larger spread around that peak caused by the uniform sampling of single excitations. The *heat bath exact singles* excitation generator produces two sharp peaks, both containing data from single excitations which were treated exactly here. The reason why this is not one sharp peak is that in an ideal case

$$p_{\text{gen.}} = \left| \frac{\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle}{\sum_{\mathbf{n}} \langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle} \right| \quad (5.17)$$

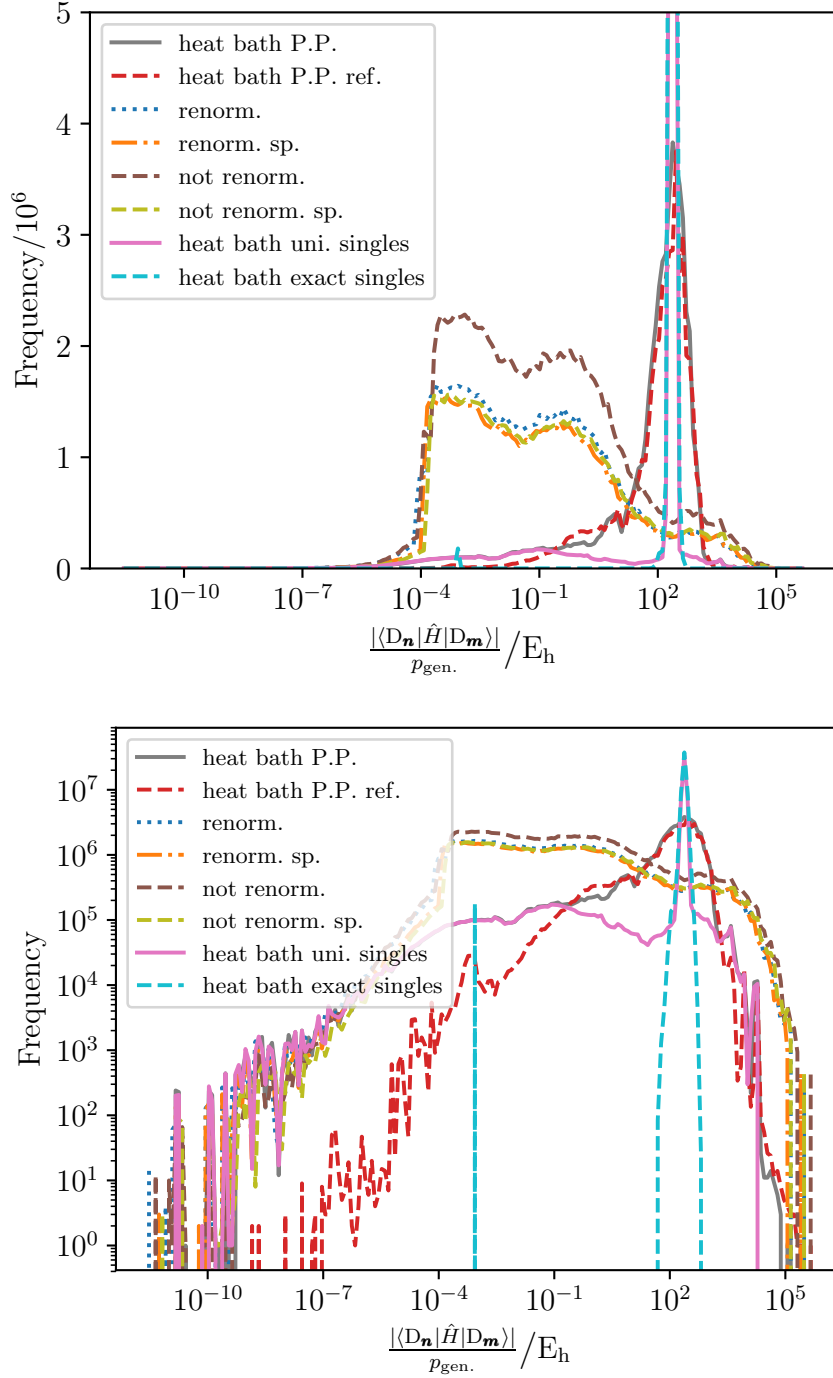


Fig. 5.3 Comparison of the histograms of $\frac{|\langle D_n | \hat{H} | D_m \rangle|}{p_{\text{gen.}}}$ for various excit. gens. The bin middles on the $\frac{|\langle D_n | \hat{H} | D_m \rangle|}{p_{\text{gen.}}}$ axis are used for the data points. CCSD was performed on three water molecules in the cc-pVDZ basis using localized MOs. The values were logged for one Monte Carlo iteration. The size of the bins is logarithmic. *Bottom* graph took the log of the frequency whereas the *top* graph did not. They both show the same data. The frequency axis in the case is truncated in the *top* graph. Most of them were restarted from the same calculation and then equilibrated before taking data. *heat bath exact singles* was restarted from an equilibrated *heat bath uniform singles* but not equilibrated since it is very slow. $\frac{|\langle D_n | \hat{H} | D_m \rangle|}{p_{\text{gen.}}} = 0$ data is not shown which includes forbidden excitations. p_{single} was set to be the same when running which was corrected in post-processing to make the mean of finite $\frac{|\langle D_n | \hat{H} | D_m \rangle|}{p_{\text{gen.}}}$ for single and double excitations approximately coincide which did not change p_{single} values by more than 30%.

Table 5.3 Fraction of non-zero allowed spawn events, both with respect to total number of spawn events. The latter represents the spawn events depicted in figure 5.3. The *renorm.* and *renorm. spin* have been combined to *renorm.* and similarly for *not. renorm.* *P.-P.* means *Power-Pitzer* and *heat b.* is *heat bath*. Individual data points have been rounded to the second decimal place. If a range is given they rounded to either value in the range.

	$\frac{\text{\#allowed non-zero}}{\text{\#total}} \text{ events}$
<i>heat b. P.-P. ref.</i>	0.66–0.67
<i>heat b. P.-P.</i>	0.69
<i>heat b. uniform singles</i>	0.72
<i>heat b. exact singles</i>	0.72
<i>renorm.</i>	0.69–0.72
<i>not. renorm.</i>	0.54–0.57

which means that

$$\frac{|\langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|}{p_{\text{gen}}} \approx \frac{1}{|\sum_{\mathbf{n}} \langle \mathbf{D}_{\mathbf{n}} | \hat{H} | \mathbf{D}_{\mathbf{m}} \rangle|} \quad (5.18)$$

in the case of an ideal excitation generator. This quantity depends on $|\mathbf{D}_{\mathbf{m}}\rangle$ and can therefore not be a constant in general unless the selection step in the CCMC or FCIQMC algorithm is adapted as well. Both *heat bath* excitation generators here have a large memory scaling ($\mathcal{O}(M^4)$) and *heat bath exact singles* which produces the sharpest peak in the histogram has a computational scaling of $\mathcal{O}(MN^2)$ (FCIQMC and in this implementation) or $\mathcal{O}(MNN_{\text{ex}})$ (ideal implementation CCMC) which makes the *heat bath exact singles* excitation generator not practical.

The main peak that the two *Power-Pitzer* excitation generators produce is wider than with the *heat bath* excitation generators but it is significantly more compact than what the uniform excitation generators give. The *heat bath Power-Pitzer ref.* excitation generator has a shorter tail on the low end but a slightly wider tail on the higher end. It has fewer than 250 events in bins with bigger $\frac{|\langle \mathbf{D}_{\mathbf{m}} | \hat{H} | \mathbf{D}_{\mathbf{n}} \rangle|}{p_{\text{gen}}}$ than the highest bin that has an event with the *heat bath uniform singles* excitation generator. The *heat bath Power-Pitzer* excitation generator has fewer than 90 events above the bin with highest $\frac{|\langle \mathbf{D}_{\mathbf{m}} | \hat{H} | \mathbf{D}_{\mathbf{n}} \rangle|}{p_{\text{gen}}}$ in the *heat bath uniform singles* case. The number of finite $\frac{|\langle \mathbf{D}_{\mathbf{m}} | \hat{H} | \mathbf{D}_{\mathbf{n}} \rangle|}{p_{\text{gen}}}$, allowed events are shown in table 5.3. The weighted excitation generators have similar

fractions of allowed non-zero events and the *heat bath Power-Pitzer ref.* excitation generator has the lower computational scaling compared to *heat bath Power-Pitzer* and the *heat bath uniform singles* excitation generator, at least in the case of CCMC. It also does not have the prohibitively large memory scaling of the *heat bath uniform singles* excitation generator.

Next, we move away from abstract performance considerations and consider how the different excitation generators affect the efficiency (as described by Holmes et al.[77]), inefficiency[96], and the position of the shoulder[54] which are all measures of the difficulty of the calculation. The efficiency η is defined as $\eta = 1/(\sigma_E^2 T)$, where σ_E is the statistical error in the energy (here projected energy) and T is the computational time taken to achieve error bar σ_E . In our case here, T was estimated by the CPU time, that sums over OpenMP threads, as determined by the parent MPI process. It is then multiplied by the number of MPI processes. It is therefore to be treated as an approximation. T is the sum of individual times for blocks of iterations and only times of iterations actually used by the reblocking analysis are summed up. We have found T to be highly dependent on implementation so η must be considered carefully. We also consider the (theoretical) algorithmic computational scaling in mind and the inefficiency a as defined by Vigor et al.[96], $a = \sigma_E \sqrt{\delta \tau N_{\text{it.}} \langle N_{\text{tot.}} \rangle}$ where $N_{\text{it.}}$ is the number of iterations considered in the blocking analysis and $\langle N_{\text{tot.}} \rangle$ is the mean number of Monte Carlo particles. When estimating the error in the efficiency and inefficiency, we ignore the correlation in the numerator and denominator of the $E_{\text{proj.}}$.

5.3.1 Coupled Cluster Monte Carlo

All coupled cluster calculations are non-initiator[35, 54]. Figure 5.4 shows the efficiency and inefficiency for chains of two or three waters in the cc-pVDZ basis performing CCSD/CCSDT with localized or canonical molecular orbitals. CCSDT was only run on the weighted excitation generators. When localization has not been applied, i.e. our canonical CCSD run, symmetry has been ignored as it also does not exist in the system with localized orbitals. The systems to study were chosen not to be too large to get

small enough error bars on efficiency and inefficiency. However, the basis set could not be too small since the *heat bath uniform singles* and the *heat bath Power–Pitzer ref.* excitation generators assume that the number of occupied orbitals is small relative to the number of total orbitals, which reflects a realistic calculation. Note that while all of the four types of calculations were run with the same number of MPI processes and OpenMP threads for the different excitation generators, these numbers varied between the types of calculations ¹⁰. The *heat bath exact singles* excitation generator is so slow that it was not possible to take sufficient data with it to produce results.

We now discuss the trends shown in figure 5.4:

- *System size*: The overall trend is that the weighted excitation generators are more efficient and less inefficient than the uniform ones. This becomes more noticeable in the larger system. As expected, modelling three waters is less efficient and more inefficient than two, the difference being more distinct with the uniform excitation generators.
- *Coupled cluster level*: When raising the excitation level to CCSDT, which we did for the weighted excitation generators, the inefficiency increases and the efficiency decreases slightly compared to the CCSD calculation. This is to be expected as the Hilbert space to cover increases. All three weighted excitation generators are affected.
- *Localization*: Using orbitals that have not been localized does not seem to affect the efficiency of the uniform excitation generators and *heat bath uni. singles*. The *heat bath Power–Pitzer (ref.)* excitation generators show a decline in efficiency and increase in inefficiency. In fact, they seem to drop to a similar efficiency level as the uniform excitation generators, *heat bath Power–Pitzer ref.* still being slightly more efficient. The inefficiency means of the *heat bath Power–Pitzer (ref.)* excitation generators are lower than the ones from the uniform

¹⁰To be precise: The dimer calculations were done with 2, 1 and 8 MPI processes for the CCSD localized, CCSD canonical and CCSDT localized calculation respectively. The trimer calculation has been done with 4 MPI processes. The CCSD canonical calculation used 24 OpenMP threads for its MPI process, all the others 12 OpenMP threads per MPI process.

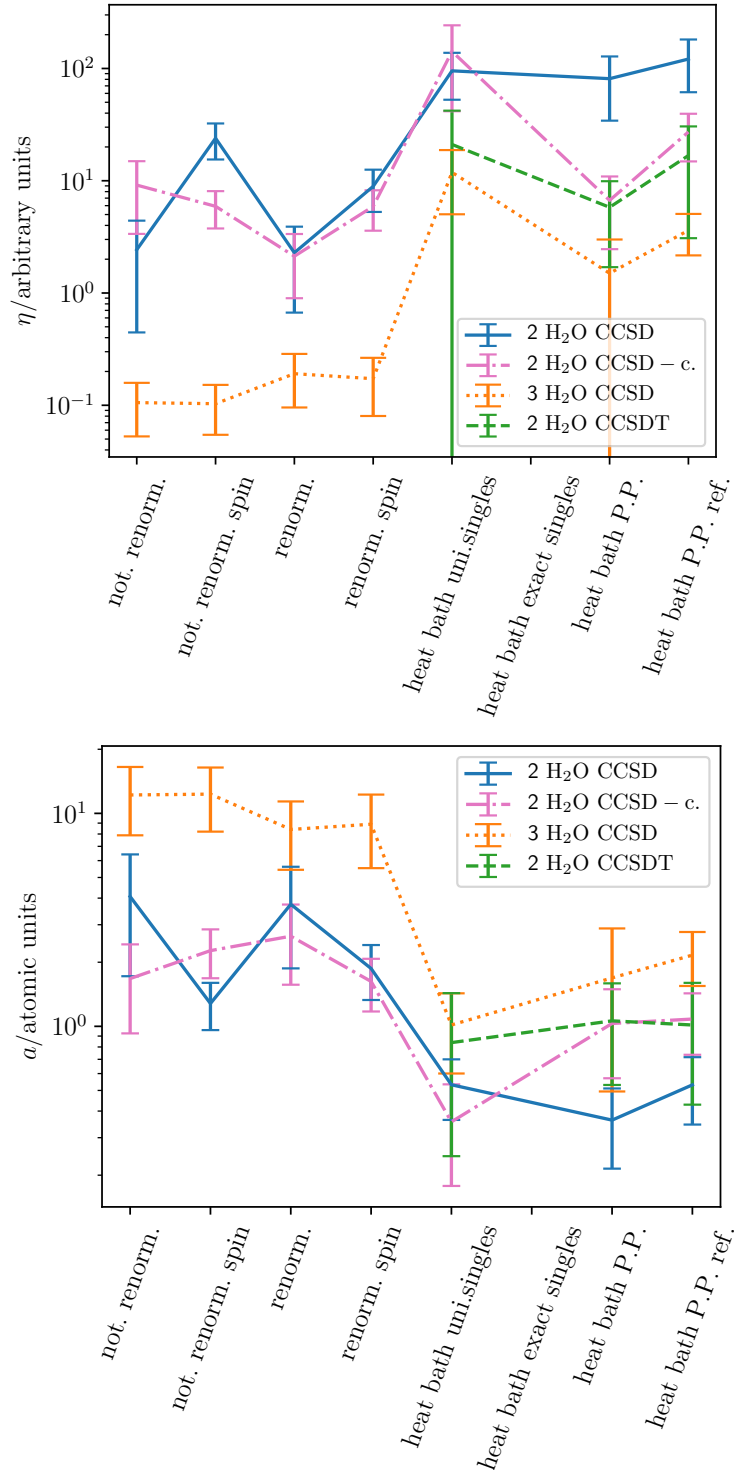


Fig. 5.4 Efficiency η (*top*) and inefficiency a (*bottom*) for chains of two or three water molecules in a cc-pVDZ basis run with CCSD/CCSDT using localized/canonical ('-c.') MOs. Error bars neglect the covariance between numerator and denominator errors in the projected energy. The *heat bath exact singles* excit. gen. was too slow for data to be taken. The different excit. gens. were run under the same conditions with the same time step etc. Only the target population was varied between the calculations. The starting iteration for *heat bath P.P.* was found such that three reblocks could be used in the trimer calculation. The number of reblocks was raised for *heat bath P.P. ref.* for the CCSDT calculation by taking the result of the previous reblocking iteration for the proj. energy. The shift and the proj. energy disagreed by more than 2σ in the canonical CCSD run with *heat bath P.P. ref.* and (*not.*) *renorm. spin* and in the CCSDT run with *heat bath P.P. ref.* excit. gens.

excitation generators, even though the error bars overlap. We expect localization to primarily to affect the weighted excitation generators as, in a double excitation, the weights in the *heat bath uni. singles* are calculated as a sum of Coulomb and — if the spins are parallel — exchange integrals whereas the *heat bath Power–Pitzer (ref.)* excitation generators only use exchange integral weights. Coulomb integrals decay as the inverse of the distance but exchange integrals are more affected by the localization. This explains why the *heat bath Power–Pitzer (ref.)* excitation generator efficiencies are more strongly affected by localization.

The *heat bath uniform singles* excitation generator performs best out of the weighted ones which is expected due to the same low computational scaling as *heat bath Power–Pitzer ref.* which is more favourable than *heat bath Power–Pitzer* while using well approximated weights for double excitations. *heat bath Power–Pitzer ref.* also seems to have higher efficiencies than *heat bath Power–Pitzer*, likely due to the better computational scaling and possibly the more accurate treatment of single excitations, and this might yet be improved by a better *heat bath Power–Pitzer ref.* implementation which scales as $\mathcal{O}(N_{\text{ex}})$ rather than $\mathcal{O}(N)$ computationally.

Next, we consider shoulder heights with CCSDT on two water molecules with localized orbitals. Shoulder heights indicate approximately the minimum number of particles needed in the simulation. Figure 5.5 shows shoulder plots where the difference in shoulder positions between the excitation generators is very clear. The weighted excitation generators again perform best. Their shoulders are significantly lower than those of uniform excitation generators, by a factor of just under 2. Of those studied, the *heat bath Power–Pitzer ref.* excitation generator has the lowest shoulder.

With localized orbitals, the weighted excitation generators all perform better than the uniform ones. The *heat bath Power–Pitzer ref.* excitation generator can scale independently of system size computationally which puts it at a clear advantage over the *heat bath Power–Pitzer* excitation generator. It also has a reduced memory scaling when comparing it to the *heat bath* excitation generators which is significant at bigger systems.




Fig. 5.5 Shoulder plots for two localized waters in a cc-pVDZ basis with CCSDT with various excitation generators. P.P. stands for *Power–Pitzer*, h.b. for *heat bath* and sp. for *spin*. The different excitation generators were run under similar conditions with the same time step etc. The weighted excitations generators started varying the shift after a total population of 20 million whereas the uniform ones did not vary the shift. The vertical lines represent the “shoulder height”, the position of the maximum plus/minus of a standard deviation. To determine the shoulder position, the mean and standard error of the mean of the 10 highest data points were taken[54].

5.3.2 Full Configuration Interaction Quantum Monte Carlo

Next, we turn to FCIQMC. The water chain with two waters in cc-pVDZ basis with localized MOs was considered with initiator FCIQMC. The (in-)efficiencies were determined at one point in the initiator curve (total population against energy). All calculations were started with the same parameters, which included the population at which the shift started varying, and so the eventual equilibrated population of the system is dependent upon the excitation generator. Blooms did happen. For uniform excitation generators it was over 10^7 particles, for the weighted ones 5.6×10^6 . Use of a larger population may lead to a decrease in measured inefficiency[96], so the results from the uniform excitation generators should be regarded as lower bounds for inefficiency. Figure 5.6 shows the efficiency and inefficiency for that system with the particle populations $N_{\text{tot.}}$ explicitly indicated. The weighted excitation generators perform comparably among themselves and all outperform the uniform ones. *heat bath Power-Pitzer ref.* and *heat bath uniform singles* excitation generators both scale linearly in the number of electrons when using FCIQMC. This study has been done on a single point in the initiator curve and we did not investigate whether the behaviour of the initiator curve changed which can affect number of particles needed for convergence. Holmes et al.[77] describe ways to reduce the memory cost by considering spins (we just store zeroes instead of considering the spin when selecting) or by not storing all the weight to select b for example. We have used double precision for the weights. However, even if our implementation is not optimal, it is clear that the *heat bath* excitation generators hit a memory ceiling with big systems significantly earlier than the *heat bath Power-Pitzer ref.* excitation generators. Also, as mentioned earlier, our *heat bath Power-Pitzer* excitation generator implementation could be improved by making sure b is only selected from virtual orbitals. However, even with a more ideal code, the computational scaling of $\mathcal{O}(M)$ remains which becomes prohibitive in large systems.

This suggests that *heat bath Power-Pitzer ref.* is an efficient excitation generator with a low shoulder that can be used in CCMC and FCIQMC as a weighted excitation generator with low computational and memory cost.

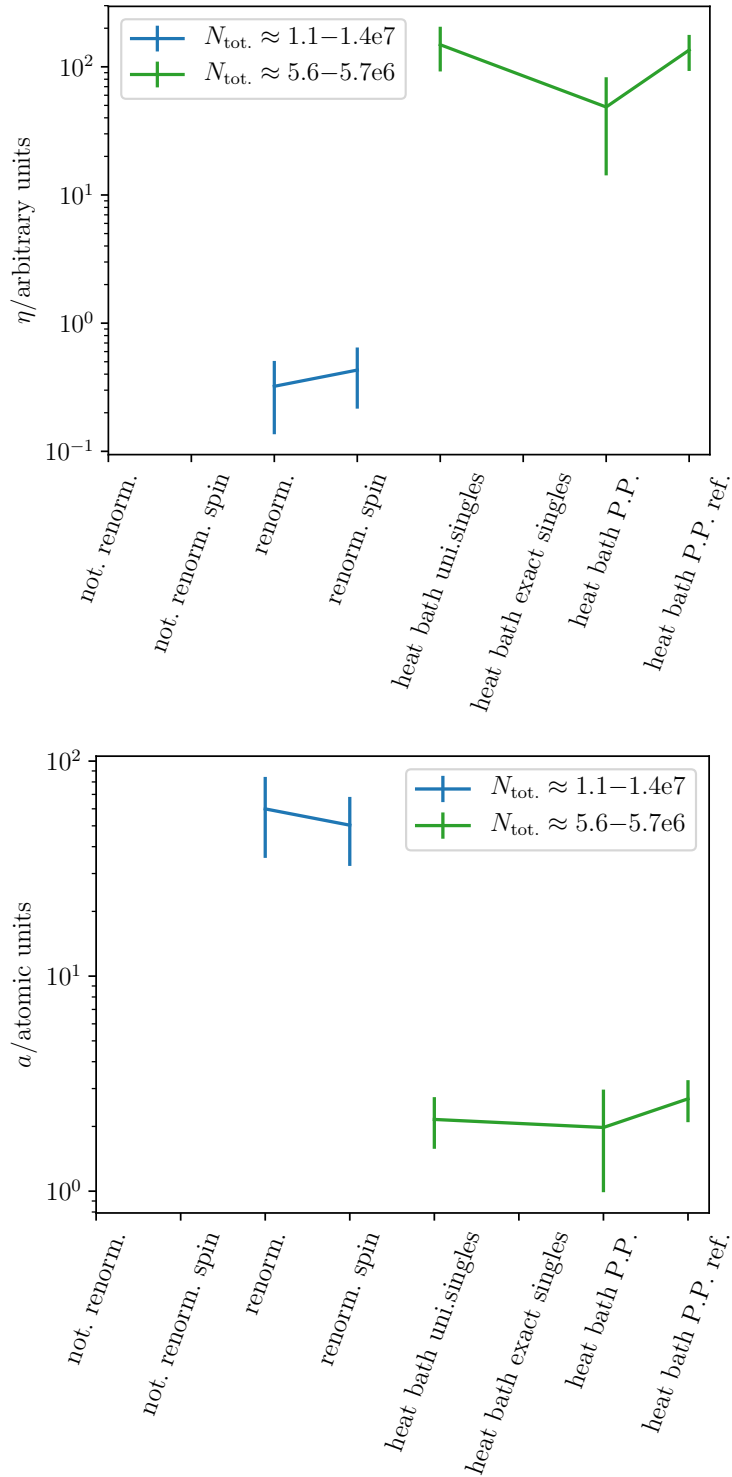


Fig. 5.6 Efficiency η (*top*) and inefficiency a (*bottom*) for a chain of two water molecules in cc-pVDZ basis using localized MOs run with initiator FCIQMC with approximate MC particle populations indicated. Error bars neglect the covariance between numerator and denominator errors in the projected energy and are over-estimates. The *heat bath exact singles* excitation generator was too slow for data to be taken. The different excitation generators were run under the same conditions with the same time step etc. The spawning arrays of the *not. renorm.* excitation generators ran out of memory so the space to store the spawned walkers would need to be increased for those results.

5.3.3 Practical Advice

As long as the memory allows it, it makes sense to make use of the *heat bath* (if no bias is present) or *heat bath uniform singles* excitation generator for FCIQMC as they have the same computational scaling as the *heat bath Power–Pitzer ref.* excitation generator and a better scaling than the *heat bath Power–Pitzer* excitation generator. The results in this paper suggest that they are also at least as efficient as the *heat bath Power–Pitzer ref.* excitation generator. For CCMC with an implementation where the *heat bath Power–Pitzer ref.* excitation generator has a scaling of $\mathcal{O}(N_{\text{ex}})$, the *heat bath Power–Pitzer ref.* excitation generator using localized orbitals is suitable.

As the system size becomes more substantial, the *heat bath* excitation generators will fail due to memory requirements. In that case, the *heat bath Power–Pitzer ref.* excitation generator should be considered, ideally with localized orbitals.

5.4 Conclusion

We have shown that especially when using localized orbitals the *heat bath Power–Pitzer ref.* excitation generator combines the advantages of *heat bath* excitation generators, which are relatively fast and use good weights but struggle with a significant memory cost and a possible bias, and the excitation generators that approximate heat bath weights by inequalities which are calculated on-the-fly reducing the memory scaling but scaling prohibitively computationally in big systems. The *heat bath Power–Pitzer ref.* excitation generator has at worst a low computational order and can be implemented with computational cost independent of system size in coupled cluster with a low memory cost.

This was the first main algorithmic development in this thesis. The next chapter accelerates the convergence. Since periodic solid systems are large, the memory intensive *heat bath* excitation generators are becoming infeasible quickly. The lighter *Power–Pitzer* excitation generators are therefore more realistic approaches for the *spawn* step there.

Chapter 6

Accelerating the Convergence with quasi-Newton

This chapter presents the second algorithmic improvement of the FCIQMC/CCMC algorithm done in my PhD, the acceleration of the convergence to the ground state energy. It is to be published — and a version is on arxiv — (in a slightly different, updated format) as

- V. A. Neufeld, and A. J. W. Thom. *Accelerating stochastic quantum chemistry*, submitted[4].

(Reproduced in part with permission from Journal of Chemical Theory and Computation, submitted for publication. Unpublished work copyright 2019 American Chemical Society. This work has been further updated after submission of the thesis and the most recent version is the one that will be published as a paper.)

Alex Thom designed the original quasi-Newton algorithm which I have expanded. I have implemented the deterministic algorithm in Section 6.3. I have conducted and mainly designed the studies in this paper. Alex Thom helped with useful discussions, providing the integral (intdump) file for the chromium dimer, comments on the manuscript and a very first version of the introduction and method section which I have rewritten with great modifications. Reviewers also contributed with helpful comments.

6.1 Introduction

After gaining knowledge of what coupled cluster level is needed to accurately determine energies in solids and increasing the sampling efficiency with different excitation generators, now the convergence to the ground state energy is optimised using a quasi-Newton scheme commonly used in conventional deterministic coupled cluster[28] which is a computationally inexpensive approximation to the Newton–Raphson method. The hope in the future is to combine faster quasi-Newton convergence of the instantaneous projected energy with a data analysis method such as in Ref. [82] which uses the instantaneous projected energy and less data than reblocking analysis[80] to estimate the energy. This can be of significant help for expensive solids calculations where requiring less data and converging faster can make a great difference in terms of what systems or number of \mathbf{k} points are feasible.

The computational effort for FCIQMC and CCMC to reach equilibration can be very significant and thus it can take a prohibitive amount of time before the energy can be even roughly estimated. The Hessian required for a Newton–Raphson propagation are approximated by using inexpensive Fock expectation value sums by the quasi-Newton method. Since it has been developed and implemented[5] for CCMC and FCIQMC, Blunt et al.[81] have also introduced an alternative Jacobi pre-conditioned propagation[150] to approximate the Hessian. A comparison is made to their method which is computationally more expensive than the approach presented here. Note that other propagator improvements, which are not discussed here, exist, including the use of Chebyshev expansion[151] and techniques used in the machine learning community which have also been applied to Quantum Monte Carlo methods to accelerate convergence[89, 152, 153]. Deustua et al.[154, 155] have used FCIQMC and CCMC to estimate deterministic amplitudes/coefficients and managed to converge to highly accurate energies quickly doing so, see for example the CAD-FCIQMC method[155]. This approach is independent of the convergence acceleration shown here, in fact they can be most likely employed simultaneously to improve convergence.

First, we will describe the quasi-Newton propagation, followed by analysing its convergence behaviour in both *deterministic* and *stochastic* propagations and comparing it to the original, Jacobi, and full Newton propagations. Finally, the quasi-Newton propagation is applied to the chromium dimer in the full Ahlrichs' SV basis[62] demonstrating its capabilities for accurate calculations of large quantum chemical systems.

6.2 Quasi-Newton Method

The quasi-Newton propagation formalism is derived by treating FCIQMC as an optimisation problem. The derivation is similar to a derivation by Davidson[150]. The conclusion also holds for CCMC.

Remember that in FCIQMC, the lowest eigenvalue of the Hamiltonian is found along with an approximation of its eigenvector which is the ground state wavefunction, $\Psi_0 = \sum_{\mathbf{i}} c_{\mathbf{i}} |\mathbf{D}_{\mathbf{i}}\rangle$. The constraint is the normalisation of the wavefunction, $\langle \Psi | \Psi \rangle = N$ for some constant N . As shown in chapter 2 for FCI, a Lagrangian \mathcal{L} with Lagrange multiplier E can therefore be written as

$$\mathcal{L} = \langle \Psi | \hat{H} | \Psi \rangle - E(\langle \Psi | \Psi \rangle - N). \quad (6.1)$$

Differentiating gives the gradient

$$g_{\mathbf{i}} = \frac{\partial \mathcal{L}}{\partial c_{\mathbf{i}}^*} \propto \langle \mathbf{D}_{\mathbf{i}} | \hat{H} - E | \Psi \rangle. \quad (6.2)$$

Setting $\mathbf{g}=\mathbf{0}$ gives the converged (F)CI equations $\langle \mathbf{D}_{\mathbf{i}} | \hat{H} - E | \Psi \rangle = 0$ for all \mathbf{i} . In the original FCIQMC formalism[33], $g_{\mathbf{i}} = \langle \mathbf{D}_{\mathbf{i}} | \hat{H} - E | \Psi \rangle$ is used to propagate from the initial guess to the ground state wavefunction in imaginary time, τ , with an update equation equivalent to steepest descent,

$$\mathbf{c}(\tau + \delta\tau) = \mathbf{c}(\tau) - \delta\tau \mathbf{g}(\tau) \quad (6.3)$$

using time step $\delta\tau$. The optimised wavefunction is Ψ_0 with energy E .

Steepest gradient descent approaches the solution linearly and is therefore inefficient. The quadratically convergent Newton–Raphson method propagates the coefficients towards $\mathbf{g}=\mathbf{0}$ by

$$\mathbf{c}(\tau + \delta\tau) = \mathbf{c}(\tau) - \delta\tau \tilde{\mathbf{H}}^{-1} \mathbf{g}(\tau) \quad (6.4)$$

where the time step $\delta\tau$ was retained for extra flexibility. The elements of the Hessian $\tilde{\mathbf{H}}$ are given by

$$\tilde{H}_{ij} = \frac{\partial g_i}{\partial c_j} \propto \langle D_i | \hat{H} - E | D_j \rangle. \quad (6.5)$$

Since inverting $\tilde{\mathbf{H}}$ is highly expensive, approximations to $\tilde{\mathbf{H}}^{-1}$ are necessary. It may be assumed that the off-diagonal elements in $\tilde{\mathbf{H}}$ are not very significant compared to the diagonal elements and so can be set to zero, leaving an easily invertible diagonal matrix, provided no diagonal elements are zero. Davidson[156] has noted the connection of pre-conditioning to the Newton–Raphson algorithm; while derived differently, this is equivalent to the Jacobi pre-conditioned propagation used by Blunt et al.[81].

Here, the example of coupled cluster theory is followed[28] where Fock expectation values for orbitals i , $\langle i | \hat{F} | i \rangle$, are used in an approximation to the diagonal Hamiltonian elements and off-diagonal elements are ignored. The diagonal elements of $\tilde{\mathbf{H}}$, $\propto \langle D_j | \hat{H} - E_{\text{HF}} - E_{\text{proj.}} | D_j \rangle$, are approximated by the sum of Fock expectation values of occupied orbitals in D_j minus the sum of Fock expectation values of occupied orbitals in the reference,

$$\langle D_j | \hat{H} - E | D_j \rangle \approx \sum_{m \text{ in } j} \langle m | \hat{F} | m \rangle - \sum_{m' \text{ in } \mathbf{0}} \langle m' | \hat{F} | m' \rangle. \quad (6.6)$$

Note that the computational cost of Blunt's Jacobi pre-conditioned propagation[81] is at least $\mathcal{O}(N_{\text{el.}})^1$ whereas the computational cost due to the quasi-Newton propagation is $\mathcal{O}(1)$.

6.3 Deterministic Propagation

To test this approximation, the different propagation techniques were first deterministically tested on a small model system where the true eigenvalues and eigenvectors were known, and stochastic noise, reaching the level of a sufficient number of particles and other challenges in stochastic propagations, could be ignored so the focus was solely on how many iterations were needed to converge.

The model system studied was the three-dimensional uniform electron gas (UEG) with two electrons of opposite spin in 1850 spinorbitals which has a Hilbert space size of 925. As mentioned in chapter 2, the Fock value for spinorbital m is given by[68]

$$\langle m|\hat{F}|m\rangle = \frac{1}{2}|\mathbf{k}|^2 - \sum_{\substack{n \text{ in } \mathbf{0} \\ m \neq n \\ \text{same spin}}} \left\langle nm \left| \frac{1}{r_{12}} \right| mn \right\rangle \left(+\frac{1}{2}V_{\text{Mad.}} \right) \quad (6.7)$$

where the last term in round brackets including the Madelung constant per electron $V_{\text{Mad.}}$ is added to spinorbitals m occupied in the reference only. $V_{\text{Mad.}} \approx -2.837297 \times (\frac{3}{4\pi r_s^3 N_{\text{el.}}})^{1/3}$ as determined by Schoof et al.[157, 158] with Wigner-Seitz radius r_s . Using the HANDE QMC code[5], an FCI calculation was performed and Hamiltonian matrix elements were calculated. The initial guess for the wavefunction was a vector with 1 at the \mathbf{D}_0 position and 0 otherwise. This corresponds to a standard FCIQMC calculation with initially one Monte Carlo particle at the reference determinant. The shift S was set to the projected energy at every iteration. The time step was set to the reciprocal of the highest eigenvalue of $\mathbf{A}^{-1}\tilde{\mathbf{H}}$ where in the original propagation \mathbf{A} is the identity

¹To approximate $\tilde{\mathbf{H}}$, we need to evaluate $\langle \mathbf{D}_j | \hat{H} - E | \mathbf{D}_j \rangle$ as part of the death step for any type of propagation, so there is no extra cost in the death step. For the spawn step, $\langle \mathbf{D}_i | \hat{H} - E | \mathbf{D}_i \rangle$ is needed. Since \mathbf{D}_i and \mathbf{D}_j differ by at most a double excitation, $\langle \mathbf{D}_j | \hat{H} - E | \mathbf{D}_j \rangle$ can be used as a starting point and the difference can be calculated. This is an $\mathcal{O}(N_{\text{el.}})$ operation (Personal communication with Nick Blunt).



(a) Propagation for $r_s = 0.5a_0$.(b) Propagation for $r_s = 20a_0$.

Fig. 6.1 (See published paper for update: slightly higher $\delta\tau$ can be used for original propagation, does not change conclusions.) Deterministic propagation of the 3D UEG with two electrons of opposite spins in 1850 spinorbitals with original, full Newton, Jacobi pre-conditioned and quasi-Newton propagation for $r_s = 0.5a_0$ (a) and $r_s = 20a_0$ (b). Jacobi $\delta\epsilon$ auto. sets the first diagonal element of the approximated Hessian to its second element whereas $\delta\epsilon = 0$ leaves the diagonal untouched. The shift is set to the projected energy at each iteration and the time step (except for the quasi-Newton run with $\delta\tau = 0.66$) for each propagation is the reciprocal of the highest eigenvalue of the propagation matrix $\mathbf{A}^{-1}\tilde{\mathbf{H}}$. In case of full Newton and Jacobi propagations, $S = S_{\min.}$ if $|S| < |S_{\min.}|$. The full Newton curve with $S_{\min.} = -10^{-2}$ and the Jacobi one with $\delta\epsilon = 0$ cannot be distinguished at this scale.

and in the other propagations it is the Hessian $\tilde{\mathbf{H}}$ or an approximation thereof. This was inspired by the fact that the highest allowed time step in FCIQMC is twice the reciprocal of the highest eigenvalue of $\tilde{\mathbf{H}}$ [33, 76] although this might not apply to all propagations exactly. The full Newton propagation used a Hessian with elements $\langle \mathbf{D}_i | \hat{H} - 0.99S | \mathbf{D}_j \rangle$ with a factor of 0.99 since its inverse would otherwise tend to be singular as $S \rightarrow E_{\text{corr.}}$. In the first iteration, where the shift and projected energy are zero, the first diagonal element is set to a small number such as 10^{-2} , 10^{-3} or 10^{-7} (see figure), in the case of the full Newton and Jacobi propagations. If *auto.* mode is chosen when using the quasi-Newton propagation, the first diagonal element of the approximated Hessian would be zero, so it is set to the second diagonal element. When using the Jacobi propagation, E in the propagation is set to the shift S and a threshold $\delta\epsilon$ is applied or the first element set to the second (auto. $\delta\epsilon$). Figure shows the propagation for $r_s = 0.5a_0$ and $r_s = 20a_0$. For quasi-Newton, two time steps are shown; one found as described above (≈ 0.477), and the other being 0.66 which is higher.

Clearly, in terms of convergence with respect to iterations, the original propagation is outperformed by the others which perform similarly to each other. As demonstrated by the full Newton propagation the initial guess for $S = S_{\text{min.}}$ can obviously affect convergence. The higher time step used for quasi-Newton performs slightly better than the automatically found time step but it is still similar in behaviour. The more correlated the UEG system gets, the higher r_s , the smaller the range in Fock eigenvalues so the more similar the original propagation is to the quasi-Newton propagation.

6.4 Stochastic Propagation

Next, the quasi-Newton propagation is compared with the original propagation in FCIQMC. The quasi-Newton propagation can be straightforwardly implemented into FCIQMC as the only changes are in the *spawn* and *death* steps. In the case of the *spawn* step, the probability that a spawn is accepted is divided by Δ_i where

$$\Delta_{\mathbf{i}} = \begin{cases} \Delta'_{\mathbf{i}} & \text{if } \Delta'_{\mathbf{i}} \geq \delta\epsilon \\ \Delta_v & \text{otherwise} \end{cases} \quad (6.8)$$

with

$$\Delta'_{\mathbf{i}} = \sum_{m \text{ in } j} \langle m | \hat{F} | m \rangle - \sum_{m' \text{ in } \mathbf{0}} \langle m' | \hat{F} | m' \rangle. \quad (6.9)$$

$\delta\epsilon$ is a threshold and Δ_v an alternative value chosen which could be set to 1 (see later part on the chromium dimer) or, as in this stochastic UEG study here, to $\delta\epsilon$. Similarly to the deterministic investigation, $\delta\epsilon$ can be chosen to be the difference between the sum of Fock energies of the reference and first excited determinant to maximise the time step possible. In the original *death* step, the *death* probability of a particle on determinant $|\mathbf{D}_{\mathbf{i}}\rangle$ is written as[33]

$$p_{\text{death}}(|\mathbf{D}_{\mathbf{i}}\rangle) \propto \delta\tau \langle \mathbf{D}_{\mathbf{i}} | \hat{H} - S | \mathbf{D}_{\mathbf{i}} \rangle. \quad (6.10)$$

If a quasi-Newton modification were also performed to the *death* step, the resulting *death* probability would be $\frac{p_{\text{death}}(|\mathbf{D}_{\mathbf{i}}\rangle)}{\Delta_{\mathbf{i}}}$. We consider the hypothetical case where the estimate of the wavefunction is a multiple of the true wavefunction, but S is not equal to the true energy the wavefunction would stay at the true solution as all determinants are affected equally by the error in S in the *death* step. However, in the case of quasi-Newton, due to the determinant dependence of $\Delta_{\mathbf{i}}$, the estimate of the wavefunction would move away from the true solution. A modified *death* step (inspired by the coupled cluster Monte Carlo modification of Franklin et al. [65]) is

$$p_{\text{death}}(|\mathbf{D}_{\mathbf{i}}\rangle) \propto \delta\tau \left(\frac{\langle \mathbf{D}_{\mathbf{i}} | \hat{H} - E_{\text{proj.}} | \mathbf{D}_{\mathbf{i}} \rangle}{\Delta_{\mathbf{i}}} + \rho(E_{\text{proj.}} - S) \right), \quad (6.11)$$

with the projected energy $E_{\text{proj.}}$ and ρ as a constant population control factor to add an extra degree of freedom. We have assumed that E_{HF} has already been subtracted of the Hamiltonian matrix diagonal. At the true solution, $E_{\text{proj.}}$ takes the correct value so the net effect of the first term in equation 6.11 when applied to the whole population

is zero, and the latter term merely scales the whole population, so the wavefunction remains at the true solution.

Using the spin non-polarised three dimensional (3D) UEG again, this time with 1850 spinorbitals, 14 electrons, and $r_s = 0.5a_0$, the stochastic propagations using FCIQMC with quasi-Newton and the original propagation were compared. The instantaneous projected energies were binned with respect to the cumulative number of particles, $N_{\text{tot.}}$, to reach those instantaneous projected energies and the mean in each bin for each calculation run calculated. The same calculation was then run at least 20 times with different random number generator seeds. The means of these independent bin means are shown in figure 6.2 with their standard deviations and standard errors across the different runs as error bars. Empty bins did not contribute to the mean or its errors. The bin positions are the same for all calculations. Note that not all calculations ran for the same number of iterations, some ended early. The cumulative number of particles $N_{\text{tot.}}$ is a measure of the cost of the calculation that is more implementation- and platform-independent than the compute time for example, as an iteration in the FCIQMC algorithm scales approximately linearly in the number of particles at that time step². A pre-calculated $\mathcal{O}(1)$ version of a *uniform Power-Pitzer* excitation generator adapted to the UEG was used[3, 78]. Floating-point amplitudes[85, 86] were employed with a spawn cutoff of 0.01. Figure 6.2 shows that the instantaneous projected energy converges significantly faster when using the quasi-Newton propagation. The time steps for the quasi-Newton propagation are 10–40 times greater than time steps of the original propagations shown. $\delta\epsilon \approx 11.8E_h$, the Fock value difference between the ground and first excited determinant of the same symmetry. As expected, using a lower initial population decreases the initial cost of converging to a certain energy but increases the noise. Population control has not been applied here, we have just focussed on convergence, not evaluating the final energy.

²Each particle does one spawn attempt here.




Fig. 6.2 Convergence of the instantaneous projected energy as a function of cumulative number of Monte Carlo particles $N_{\text{tot.}}$ as cost measure for the 3D UEG with 1850 spinorbitals, 14 electrons, and $r_s = 0.5a_0$ in FCIQMC. The instantaneous projected energy was binned with respect to the cumulative particle number and the bin means calculated. Each calculations was done at least 20 times in independent runs and the means of those bin means are shown with their standard deviations and standard errors as outer and inner error bars respectively, placed at the number of cumulative particles that is at the middle of the bin. Some runs ended early so not every data point was determined by the same number of independent bins. The estimate for the true energy is taken from Ref. 2. Twice its error is shown in the line spread but it is too small to be visible. $N_{\text{tot.}}(\tau = 0)$ the initial population. $\rho = 1.0$, $\delta\epsilon \approx 11.8E_h$, $\Delta_v = \delta\epsilon$ and the shift was not varied.

6.5 Application to the Chromium Dimer

Finally, the quasi-Newton propagation was tested on an archetypical quantum chemistry problem, the chromium dimer, at a bond length of 1.5\AA . The basis set considered is Ahlrichs' SV[62] where first a CAS of 24 electrons correlated in 30 spatial orbitals was applied and then the full system was studied with initiator FCIQMC. The Hartree–Fock orbitals and their integrals were evaluated with the Psi4 code[94, 95]. The weighted heat-bath excitation generator[77] (adapted[3]) has been used. Again, floating-point amplitudes[85, 86] were employed with a spawn cutoff of 0.01. Booth et al.[34] have previously applied FCIQMC to the chromium dimer with a CAS and DMRG results exist for both smaller CAS[159–161] and full[161] system, also in Ahlrich's SV basis[62]³. For the smaller CAS system, figure 6.3 shows various initiator convergence curves, displaying energy as a function of population size, for quasi-Newton and original propagation. The quasi-Newton propagation was tested at $\delta\tau = 0.002, 0.008$ and 0.02 , whereas the original was only stable or did not converge very slowly at $\delta\tau = 0.002$ out of these time steps (given the set initial population etc.). The range of the result by Booth et al.[34] is shown. Reblocking analysis was used to estimate errors on quoted energy values[80]. All initiator curves tend to this result and the threshold $\delta\epsilon$ did not seem to have a noticeable effect.

³Refs 160 and 34 state that they have used Ahlrich's SV(P) or SVP basis set. In summary, given that their results agree very well with ours, we conclude that we most likely used the same, SV basis set, details given here. The basis we used (Ahlrich's SV basis set) can be found at EMSL Basis Set Exchange Library, <https://bse.pnl.gov/bse/portal> [accessed 22.05.2019], under "Ahlrichs VDZ" and selecting "Cr" as the element. It has {63311/53/41} functions[62]. SV(P)/SVP then contains a polarizing p function (coefficients 0.1206750 and 1.0000000) as well and that basis set can be found under "Ahlrichs pVDZ". The Hartree–Fock, CCSD and CCSD(T) energies in a CAS of 24 electrons in 30 orbitals (freezing the lowest occupied orbitals) were compared using the Psi4 code. The Hartree–Fock was $-2085.57297 E_h$ in the SV basis and $-2085.60285 E_h$ in the SV(P)/SVP basis. Our full active space SV CCSD(T) energy, $-2086.39864 E_h$, agrees with Olivares-Amaya et al.[161]. In this section, the correlation energies of other studies were calculated by subtracting the Hartree–Fock energy in a SV basis (no polarising p) off the total energy quoted in the various studies. The difference in correlation energies between the SV and the SV(P)/SVP basis sets with respect to the SV Hartree–Fock energy in this (24e, 30o) CAS was -0.03 and -0.05 for CCSD and CCSD(T) respectively. This difference is an order of magnitude larger than energy differences to those studies in this chromium investigation here. We therefore concluded that the basis set used was SV in Refs 160 and 34 as well.




Fig. 6.3 Initiator curve of Cr_2 in a (24 electrons, 30 orbitals) active space in SV basis[62] at a bond length of 1.5\AA . Three DMRG results[159–161] are shown with horizontal lines. An initiator curve point from Booth et al.[34] is included. $\Delta_v = 1$, $\rho = 1$ here.




Fig. 6.4 Convergence of instantaneous projected energy in the all-electron chromium dimer in the SV basis[62] at a bond length of 1.5\AA of quasi-Newton (QN) and original propagation at a target population of 5×10^5 (population where shift starts varying) and initial population of 100. The QN results were run with $\rho = 0$ up to iteration 5000 ($\delta\tau = 0.0015$) and then set to $\rho = 1$, always using $\Delta_v = 1$ and $\delta\epsilon = 10^{-5}$. The instantaneous projected energy was binned with respect to the cumulative particle number and the bin means calculated. Each calculations was done at least 100 times in independent runs and the means of those bin means are shown with their standard errors as error bars, placed at the number of cumulative particles that is at the middle of the bin. Only calculations where the inst. $E_{\text{proj.}} < 0.1$ and $> -3.8E_h$ always were included. The horizontal lines (least negative $E_{\text{proj.}}$ Orig. at $\delta\tau = 0.0007$, then QN and most negative $E_{\text{proj.}}$ is Orig. at $\delta\tau = 0.0005$) indicate the mean $E_{\text{proj.}}$ and its error found by taking the mean energy of all calculations of that type with at least 5×10^5 iterations. The left most vertical shows when ρ was changed in the QN calculation and the others show when the shift was allowed to vary in the respective calculation. The vertical lines do not show error bars.

Convergence of the full all-electron system with a Hilbert space size of 10^{22} was then studied for a particular target population comparing quasi-Newton to original propagation (figure 6.4). In figure 6.4, the convergence of original (at $\delta\tau = 0.0007$) and quasi-Newton propagation defined as the point of overlap with the expected value is

comparable. However, the quasi-Newton propagation is slightly faster convergent, even according to that definition, and the cost to get within $\pm 0.005 E_h$, even if not stable, is significantly less costly than with the original propagation.

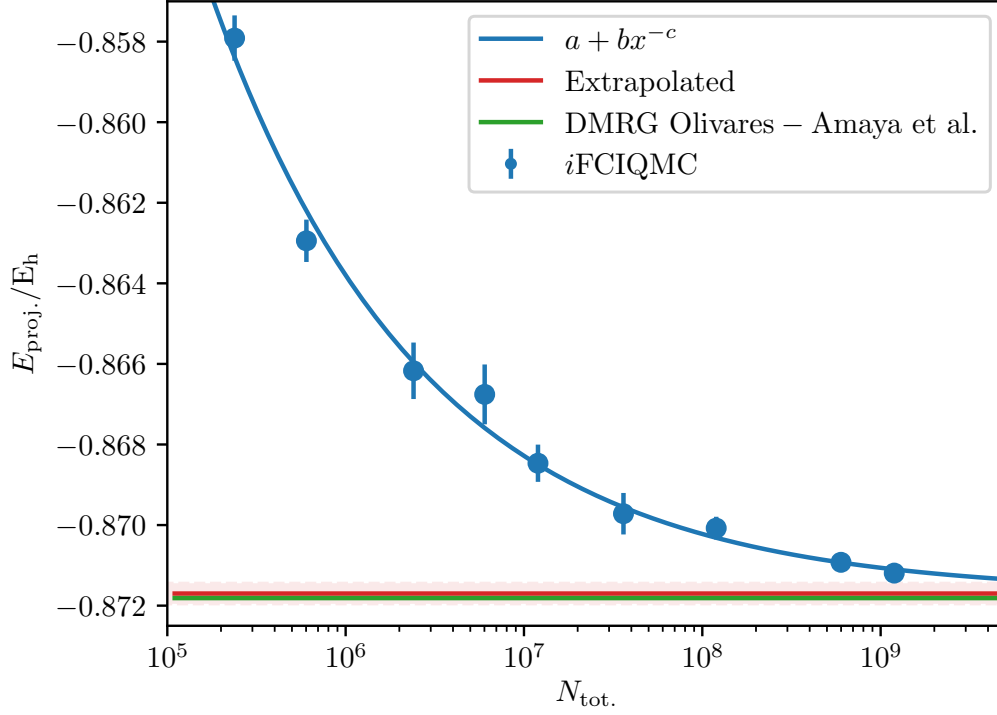


Fig. 6.5 Initiator curve of Cr_2 in the full-electron chromium dimer in the SV basis[62] at a bond length of 1.5\AA . A DMRG[161] result is shown, slightly below the extrapolated estimate.

An initiator quasi-Newton study with populations up to just above 10^9 was done and a function of the form $a + bx^{-c}$ was fitted to the data set, see figure 6.5. The determined convergence value is $-0.8717(3) E_h$ which agrees with DMRG[161], $-0.871813 E_h$. The maximum number of particles is of order 10^9 , a factor of 10^{13} reduction from the complete Hilbert space. As with the smaller CAS study, this shows that FCIQMC with quasi-Newton propagation gives reliable energies.

6.6 Conclusion and Further Work

We have shown that the quasi-Newton propagation introduced here (applicable to both CCMC and FCIQMC) can accelerate the convergence of the (instantaneous) projected energy compared to the original propagation. It scales more favourably ($\mathcal{O}(1)$ instead of $\mathcal{O}(N_{\text{el}})$) than the Jacobi propagation while having a comparable benefit. In conjunction with an excitation generator that does not scale with system size, such as the *heat bath Power Pitzer ref.* excitation generator[3] in the case of CCMC, not adding extra scaling to the algorithm is important in large electronic systems. Using the quasi-Newton propagation, we quoted the first (initiator) FCIQMC result on the chromium dimer in the full SV basis set[62]. The next steps would be to combine this with a data analysis method as in Ref. [82] that can give an energy estimate using less data and the instantaneous projected energy as it was shown that the instantaneous projected energy converges faster with quasi-Newton. The shift still takes time to converge, which is an issue that needs to be tackled for population control, especially should the shift be needed for data analysis.

Chapter 7

Concluding Remarks and Future Outlook

This thesis has shown steps necessary on the way to calculate accurate energies in periodic solid systems with CCMC and FCIQMC, after Booth et al.[41] have published initial “realistic” solid results with FCIQMC, even though the sizes of the systems studied were small. First, a model solid system – easier to handle than a “realistic” solid – was studied at different degrees of electron correlation with CCMC and the level of coupled cluster needed for accurate energies was determined[2], see chapter 4. This information can help in future to approximately estimate what level is required for a particular solid system. Both CCMC and FCIQMC are still limited severely in the system size they can study by their computational cost and potentially memory requirements. Solids where a small number of \mathbf{k} points is not sufficient are not feasible yet. Improving the excitation generators[77, 78, 3], see chapter 5, and the convergence to the ground state energy[81, 4], see chapter 6, are steps towards reducing the computational costs to make CCMC and FCIQMC possible in such systems so that (systematically improvable) accurate energies are possible. Being able to use high performance computing resources is important for this venture, see chapter 3. Very recently, we published[5] some initial solids result of diamond using CCMC and

Summary of solid study in latest HANDE QMC publication (Ref. [5]):

(Reproduced in part with permission from [5]. Copyright 2019 American Chemical Society. ACS Articles on Request author-directed link: <http://pubs.acs.org/articlesonrequest/AOR-bhuCYVv5KiPXUyNFTwsd> and doi: <https://doi.org/10.1021/acs.jctc.8b01217>)

WARNING: This study was very preliminary, results were not fully benchmarked, etc, first and some settings in the calculations might not be ideal/fully understood yet, so this should be understood as a demonstration rather than a scientific study! This demonstration (see end of this interlude for contributions) is the first publication of CCMC applied to solids, demonstrating HANDE QMC’s ability to tackle “realistic” solids systems.

Using DFT orbitals (DZVP basis, GTH pseudopotential[162–164]), CCMC up to CCSDTQ level and FCIQMC were run up to various \mathbf{k} point levels. This was then compared to CCSD results with PySCF[93] using HF orbitals and previous literature results by McClain et al[43], who also employed PySCF, see figure 7.1. Orbitals in this study were density fitted[45].

Potential explanations given for the disagreements between different CCSD results at \mathbf{k} points other than the Γ point were the varying treatments of the exchange integrals (to be published) and orbitals (HF or DFT optimisation) whose effects are visible between the CCMC and CCSD-PySCF result, as well as the shifting of the mesh to include the Γ point which was not done by McClain et al.[43].

It is again stressed that this study was very approximate as mentioned in the beginning here and it used a small number of \mathbf{k} points and basis set.




Fig. 7.1 Convergence of the correlation energy — defined as the difference of (finite) total energy and HF energy to find a consistent measure when having both DFT and HF orbitals — with the number of \mathbf{k} points for diamond at lattice constant 3.567\AA in a DZVP basis with a GTH pseudopotential[162–164] (see below for contributions), the same basis set and pseudopotential as in the study by McClain et al.[43]. The CCMC and FCIQMC results use DFT orbitals whereas CCSD-PySCF employed HF orbitals. The CCMC and FCIQMC results were then corrected by adding the DFT energy and subtracting the HF energy (beware of different exchange treatments). The mesh for those calculations has been shifted so that the Γ point is part of it. Only a part of McClain et al.’s data is shown.

Contributions: Results and figure mostly by Jiří Etrych, a summer research student in the group in 2018 that I co-supervised. I contributed with discussions, supervision, much of the text in the corresponding section in the paper (edited by the other authors) and ran the FCIQMC result, and slightly modified the figure to fit in this thesis. Alex Thom had the original ideas for the study and helped with discussions and further ideas. There were also comments from the other authors of Ref. [5]. The code to enable solids calculations in HANDE QMC code was mainly written by Charlie Scott and other HANDE developers contributed, including me. Interface to PySCF (to be published) to dump required integrals mainly written by Alex Thom and Jiří Etrych.

FCIQMC up to a 222 \mathbf{k} points grid¹, showing that calculations in (albeit small) solids are now doable with CCMC as well. A quick summary of that study is given on the previous page in the area separated by horizontal lines. Non-uniform excitation generators were employed but the convergence was not accelerated yet at that time. There are various other challenges with solids calculations, such as pre-computing the one- and two-body integrals quickly for large systems, e.g. using codes such as PySCF[93], and making sure that self-interaction is treated correctly in those integrals. To go beyond small basis sets and number of \mathbf{k} points with CCMC and FCIQMC, further algorithmic improvements are necessary, too.

These improvements include further developments in the CCMC and FCIQMC algorithms, such as transcorrelated FCIQMC/similarity-transformed FCIQMC[129, 165, 166] that can help overcome the electron cusp divergence. Combining (*i*)FCIQMC with other quantum chemistry methods, e.g. using selected configuration interaction approaches (e.g. Refs [167–169], see Ref.[170] for more) to determine the initiator determinants[170], can also be a fruitful path. Of course, there can also be further reformulations of stochastic coupled cluster or stochastic configuration interaction, e.g. very recently Scott et al.[171] designed a stochastic coupled cluster algorithm that is closer to deterministic diagrammatic coupled cluster than CCMC, decreasing memory requirements.

Another path to success can be the type of computing resources. DMC has been tested on field-programmable gate arrays (FPGAs) for example[172, 173]. Graphical processing units (GPUs) have also been used with quantum chemistry methods, see e.g. Refs [174, 175].

Further advanced statistical or machine learning methods can also extend the reach of quantum chemistry, see e.g. work in Refs [176–183]. This includes kernel fitting techniques, e.g. Ref.[176], and deep learning strategies, e.g. Refs [181, 183], for example.

¹See end of interlude on previous page for author contributions, results and study mainly by Jiří Etrych, a summer research student in the group in 2018 that I co-supervised.

Other ways to reach accurate energies besides further quantum Monte Carlo methods[32] include the development of deterministic approaches, converging to the ground state energy in the limit of including the whole space for example. This could be selected configuration interaction methods, such as heat bath configuration interaction[169], or (deterministic) coupled cluster. As mentioned in the introduction, recently, coupled cluster has been used to study solid systems, e.g. in Refs [40–43, 12], and various advancements have been made to speed up the calculations[40, 44, 46, 45, 12]. However, the size of periodic solid that can be tackled (with a sufficient high level coupled cluster truncation level) is still limited, necessitating further development.

These are some directions that the quest of how to tackle solids with CCMC or FCIQMC can take. As mentioned above, acceleration can be achieved by the design and/or implementation of stochastic or machine learning advances as well as modifications to quantum chemistry methods that can be applied to CCMC/FCIQMC as well and by alternative hardware. Suitable solid systems for CCMC/FCIQMC are correlated enough that CCSD/CCSD(T) is not sufficient, otherwise deterministic approaches are more mature and faster for the problem at hand, unless the stochastic nature of CCMC becomes necessary due to high memory costs. Overall, it can be said that while quantum chemistry is expensive it is also indispensable when accurate energies are required, which implies that it is important to develop quantum chemistry further, increasing the number of problems that can be solved. The work in this thesis contributes a step towards that goal.

Bibliography

- [1] J. S. Spencer, V. A. Neufeld, W. A. Vigor, et al. *Large scale parallelization in stochastic coupled cluster*. J. Chem. Phys. **149**, 20, 204103 (2018).
- [2] V. A. Neufeld and A. J. W. Thom. *A study of the dense uniform electron gas with high orders of coupled cluster*. J. Chem. Phys. **147**, 19, 194105 (2017).
- [3] V. A. Neufeld and A. J. W. Thom. *Exciting Determinants in Quantum Monte Carlo: Loading the Dice with Fast, Low-Memory Weights*. J. Chem. Theory Comput. **15**, 1, 127–140 (2019).
- [4] V. A. Neufeld and A. J. W. Thom. *Accelerating Stochastic Quantum Chemistry*. arXiv [physics.chem-ph, 1910.05210] (2019).
- [5] J. S. Spencer, N. S. Blunt, S. Choi, et al. *The HANDE-QMC Project: Open-Source Stochastic Quantum Chemistry from the Ground State Up*. J. Chem. Theory Comput. **15**, 3, 1728–1742 (2019).
- [6] J. A. Elliott. *Novel approaches to multiscale modelling in materials science*. Int. Mater. Rev. **56**, 4, 207–225 (2011).
- [7] D. Frenkel and B. Smit. *Understanding molecular simulation: From algorithms to applications*, volume 1. Elsevier Science, 2nd edition.
- [8] P. D. Bristowe and P. J. Hasnip. *General aspects of materials modelling*. Materials science and technology (New York, N.Y.). Maney Pub., for the Institute of Materials, Minerals and Mining, London (2005).
- [9] R. T. Fenner. *Continuum Mechanics Problems*, pages 19–45. Imperial College Press (2013).
- [10] J. Yang, W. Hu, D. Usvyat, et al. *Ab initio determination of the crystalline benzene lattice energy to sub-kilojoule/mole accuracy*. Science (80-.). **345**, 6197, 640–643 (2014).
- [11] E. Mostaani, N. D. Drummond, and V. I. Fal’ko. *Quantum Monte Carlo calculation of the binding energy of bilayer graphene*. Phys. Rev. Lett. **115**, 11, 115501 (2015).
- [12] T. Gruber, K. Liao, T. Tsatsoulis, et al. *Applying the Coupled-Cluster Ansatz to Solids and Surfaces in the Thermodynamic Limit*. Phys. Rev. X **8**, 2, 021043 (2018).

- [13] V. L. Deringer, M. A. Caro, R. Jana, et al. *Computational Surface Chemistry of Tetrahedral Amorphous Carbon by Combining Machine Learning and Density Functional Theory*. Chem. Mater. **30**, 21, 7438–7445 (2018).
- [14] Y. S. Al-Hamdani, D. Alfè, O. A. von Lilienfeld, et al. *Water on BN doped benzene: A hard test for exchange-correlation functionals and the impact of exact exchange on weak binding*. J. Chem. Phys. **141**, 18, 18C530 (2014).
- [15] T. Gruber and A. Grüneis. *Ab initio calculations of carbon and boron nitride allotropes and their structural phase transitions using periodic coupled cluster theory*. Phys. Rev. B **98**, 13, 134108 (2018).
- [16] M. Veit, S. K. Jain, S. Bonakala, et al. *Equation of State of Fluid Methane from First Principles with Machine Learning Potentials*. J. Chem. Theory Comput. **15**, 4, 2574–2586 (2019).
- [17] W. Kohn. *Nobel Lecture: Electronic structure of matter—wave functions and density functionals*. Rev. Mod. Phys. **71**, 5, 1253–1266 (1999).
- [18] W. Kohn and L. J. Sham. *Self-Consistent Equations Including Exchange and Correlation Effects*. Phys. Rev. **140**, 4A, A1133–A1138 (1965).
- [19] P. Hohenberg and W. Kohn. *Inhomogeneous Electron Gas*. Phys. Rev. **136**, 3B, B864–B871 (1964).
- [20] M. Arita, S. Arapan, D. R. Bowler, et al. *Large-scale DFT simulations with a linear-scaling DFT code CONQUEST on K-computer*. J. Adv. Simul. Sci. Eng. **1**, 1, 87–97 (2014).
- [21] A. Szabo and N. S. Ostlund. *Modern quantum chemistry: introduction to advanced electronic structure theory*. McGraw-Hill Publishing Company/ Dover Publications, inc./, New York, first rev. edition (1989).
- [22] T. Helgaker, P. Jørgensen, and J. Olsen. *Hartree-Fock Theory*. In *Mol. Electron. Theory*, pages 433–522. John Wiley & Sons, Ltd, Chichester, UK (2000).
- [23] C. Møller and M. S. Plesset. *Note on an approximation treatment for many-electron systems*. Phys. Rev. **46**, 7, 618–622 (1934).
- [24] F. Coester and H. Kümmel. *Short-range correlations in nuclear wave functions*. Nucl. Phys. **17**, 477–485 (1960).
- [25] J. Čížek. *On the correlation problem in atomic and molecular systems. Calculation of wavefunction components in Ursell-type expansion using quantum-field theoretical methods*. J. Chem. Phys. **45**, 11, 4256–4266 (1966).
- [26] J. Čížek and J. Paldus. *Correlation problems in atomic and molecular systems III. Rederivation of the coupled-pair many-electron theory using the traditional quantum chemical methods*. Int. J. Quantum Chem. **5**, 4, 359–379 (1971).
- [27] R. J. Bartlett and M. Musiał. *Coupled-cluster theory in quantum chemistry*. Rev. Mod. Phys. **79**, 1, 291–352 (2007).

- [28] T. Helgaker, P. Jørgensen, and J. Olsen. *Coupled-Cluster Theory*. In *Mol. Electron. Theory*, pages 648–723. John Wiley & Sons, Ltd, Chichester, UK (2000).
- [29] T. Helgaker, P. Jørgensen, and J. Olsen. *Configuration-Interaction Theory*. In *Mol. Electron. Theory*, pages 523–597. John Wiley & Sons, Ltd, Chichester, UK (2000).
- [30] K. Raghavachari, G. W. Trucks, J. A. Pople, et al. *A fifth-order perturbation comparison of electron correlation theories*. Chem. Phys. Lett. **157**, 6, 479–483 (1989).
- [31] T. J. Lee and G. E. Scuseria. *Achieving Chemical Accuracy with Coupled-Cluster Theory*. In *Quantum Mech. Electron. Struct. Calc. with Chem. Accuracy*, pages 47–108. Springer Netherlands, Dordrecht (1995).
- [32] W. M. C. Foulkes, L. Mitas, R. J. Needs, et al. *Quantum Monte Carlo simulations of solids*. Rev. Mod. Phys. **73**, 1, 33–83 (2001).
- [33] 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**, 5, 054106 (2009).
- [34] G. H. Booth, S. D. Smart, and A. Alavi. *Linear-scaling and parallelisable algorithms for stochastic quantum chemistry*. Mol. Phys. **112**, 14, 1855–1869 (2014).
- [35] 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**, 4, 041103 (2010).
- [36] J. J. Shepherd, G. Booth, A. Grüneis, et al. *Full configuration interaction perspective on the homogeneous electron gas*. Phys. Rev. B **85**, 8, 081103 (2012).
- [37] J. C. Greer. *Estimating full configuration interaction limits from a Monte Carlo selection of the expansion space*. J. Chem. Phys. **103**, 5, 1821–1828 (1995).
- [38] Y. Ohtsuka and S. Nagase. *Projector Monte Carlo method based on configuration state functions. Test applications to the H₄ system and dissociation of LiH*. Chem. Phys. Lett. **463**, 4-6, 431–434 (2008).
- [39] A. J. W. Thom. *Stochastic Coupled Cluster Theory*. Phys. Rev. Lett. **105**, 26, 263004 (2010).
- [40] A. Grüneis, G. H. Booth, M. Marsman, et al. *Natural Orbitals for Wave Function Based Correlated Calculations Using a Plane Wave Basis Set*. J. Chem. Theory Comput. **7**, 9, 2780–2785 (2011).
- [41] G. H. Booth, A. Grüneis, G. Kresse, et al. *Towards an exact description of electronic wavefunctions in real solids*. Nature **493**, 7432, 365–370 (2013).
- [42] A. Grüneis. *A coupled cluster and Møller-Plesset perturbation theory study of the pressure induced phase transition in the LiH crystal*. J. Chem. Phys. **143**, 10, 102817 (2015).

- [43] J. McClain, Q. Sun, G. K.-L. Chan, et al. *Gaussian-Based Coupled-Cluster Theory for the Ground-State and Band Structure of Solids*. J. Chem. Theory Comput. **13**, 3, 1209–1218 (2017).
- [44] G. H. Booth, T. Tsatsoulis, G. K. L. Chan, et al. *From plane waves to local Gaussians for the simulation of correlated periodic systems*. J. Chem. Phys. **145**, 8, 084111 (2016).
- [45] Q. Sun, T. C. Berkelbach, J. D. McClain, et al. *Gaussian and plane-wave mixed density fitting for periodic systems*. J. Chem. Phys. **147**, 16, 164119 (2017).
- [46] K. Liao and A. Grüneis. *Communication: Finite size correction in periodic coupled cluster theory calculations of solids*. J. Chem. Phys. **145**, 14, 141102 (2016).
- [47] I. Y. Zhang and A. Grüneis. *Coupled Cluster Theory in Materials Science*. Front. Mater. **6** (2019).
- [48] T. Helgaker, P. Jørgensen, and J. Olsen. *Molecular electronic-structure theory*. John Wiley & Sons, Chichester (2000).
- [49] P. W. Atkins and R. Friedman. *Molecular quantum mechanics*. Oxford University Press, fifth edition (2011).
- [50] J. A. Pople. *Nobel Lecture: Quantum chemical models*. Rev. Mod. Phys. **71**, 5, 1267–1274 (1999).
- [51] P. A. M. Dirac. *A new notation for quantum mechanics*. Math. Proc. Cambridge Philos. Soc. **35**, 3, 416–418 (1939).
- [52] E. Schrödinger. *Quantisierung als Eigenwertproblem*. Ann. Phys. **384**, 4, 361–376 (1926).
- [53] W. Pauli. *Über den Zusammenhang des Abschlusses der Elektronengruppen im Atom mit der Komplexstruktur der Spektren*. Zeitschrift für Phys. **31**, 1, 765–783 (1925).
- [54] 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**, 8, 084108 (2016).
- [55] M. Born and R. Oppenheimer. *Zur Quantentheorie der Molekeln*. Ann. Phys. **389**, 20, 457–484 (1927).
- [56] P. J. Mohr, B. N. Taylor, and D. B. Newell. *CODATA Recommended Values of the Fundamental Physical Constants: 2010*. J. Phys. Chem. Ref. Data **41**, 4, 043109 (2012).
- [57] J. C. Slater. *The Theory of Complex Spectra*. Phys. Rev. **34**, 10, 1293–1322 (1929).
- [58] E. U. Condon. *The Theory of Complex Spectra*. Phys. Rev. **36**, 7, 1121–1133 (1930).

- [59] J. Kohanoff. *Quantum many-body theory: chemical approaches*. In *Electron. Struct. Calc. Solids Mol.*, pages 28–50. Cambridge University Press, Cambridge (2006).
- [60] W. J. Hehre, R. F. Stewart, and J. A. Pople. *Self-Consistent Molecular-Orbital Methods. I. Use of Gaussian Expansions of Slater-Type Atomic Orbitals*. J. Chem. Phys. **51**, 6, 2657–2664 (1969).
- [61] W. J. Hehre, R. Ditchfield, and J. A. Pople. *Self-Consistent Molecular Orbital Methods. XII. Further Extensions of Gaussian-Type Basis Sets for Use in Molecular Orbital Studies of Organic Molecules*. J. Chem. Phys. **56**, 5, 2257–2261 (1972).
- [62] A. Schäfer, H. Horn, and R. Ahlrichs. *Fully optimized contracted Gaussian basis sets for atoms Li to Kr*. J. Chem. Phys. **97**, 4, 2571–2577 (1992).
- [63] T. H. Dunning. *Gaussian basis sets for use in correlated molecular calculations. I. The atoms boron through neon and hydrogen*. J. Chem. Phys. **90**, 2, 1007–1023 (1989).
- [64] T. Helgaker, P. Jørgensen, and J. Olsen. *Gaussian Basis Sets*. In *Mol. Electron. Theory*, pages 287–335. John Wiley & Sons, Ltd, Chichester, UK (2014).
- [65] R. S. T. Franklin, J. S. Spencer, A. Zocante, et al. *Linked coupled cluster Monte Carlo*. J. Chem. Phys. **144**, 4, 044111 (2016).
- [66] R. M. Martin. *Uniform electron gas and simple metals*. In *Electron. Struct.*, pages 100–118. Cambridge University Press, Cambridge (2004).
- [67] G. Giuliani and G. Vignale. *Introduction to the electron liquid*. In *Quantum Theory Electron Liq.*, pages 1–68. Cambridge University Press, Cambridge (2005).
- [68] J. J. Shepherd. *A Quantum Chemical Perspective on the Homogeneous Electron Gas*. Ph.D. thesis, University of Cambridge (2013).
- [69] P.-F. Loos and P. M. W. Gill. *The uniform electron gas*. WIREs Comput. Mol. Sci. **6**, 4, 410–429 (2016).
- [70] A. Altland and B. Simons. *Second quantization*. In *Condens. Matter F. Theory*, pages 39–93. Cambridge University Press (2006).
- [71] J. E. Gubernatis, N. Kawashima, and P. Werner. *Quantum Monte Carlo methods : algorithms for lattice models*. Cambridge University Press (2016).
- [72] N. Metropolis and S. Ulam. *The Monte Carlo Method*. J. Am. Stat. Assoc. **44**, 247, 335–341 (1949).
- [73] M. Saito and M. Matsumoto. *SIMD-Oriented Fast Mersenne Twister: a 128-bit Pseudorandom Number Generator*. In A. Keller, S. Heinrich, and H. Niederreiter (Editors), *Monte Carlo Quasi-Monte Carlo Methods 2006*, pages 607–622. Springer Berlin Heidelberg, Berlin, Heidelberg (2008).

-
- [74] G. H. Givens and J. A. Hoeting. *Computational Statistics*. John Wiley & Sons, Inc., Hoboken, NJ, USA, second edition (2013).
- [75] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, et al. *Equation of State Calculations by Fast Computing Machines*. J. Chem. Phys. **21**, 6, 1087–1092 (1953).
- [76] N. Trivedi and D. M. Ceperley. *Ground-state correlations of quantum antiferromagnets: A Green-function Monte Carlo study*. Phys. Rev. B **41**, 7, 4552–4569 (1990).
- [77] A. A. Holmes, H. J. Changlani, and C. J. Umrigar. *Efficient Heat-Bath Sampling in Fock Space*. J. Chem. Theory Comput. **12**, 4, 1561–1571 (2016).
- [78] S. D. Smart, G. H. Booth, and A. Alavi. *Excitation generation in full configuration interaction quantum Monte Carlo based on Cauchy-Schwarz distributions*. unpublished .
- [79] J. S. Spencer, N. S. Blunt, and W. M. Foulkes. *The sign problem and population dynamics in the full configuration interaction quantum Monte Carlo method*. J. Chem. Phys. **136**, 5, 054110 (2012).
- [80] H. Flyvbjerg and H. G. Petersen. *Error estimates on averages of correlated data*. J. Chem. Phys. **91**, 1, 461–466 (1989).
- [81] N. S. Blunt, A. J. W. Thom, and C. J. C. Scott. *Preconditioning and Perturbative Estimators in Full Configuration Interaction Quantum Monte Carlo*. J. Chem. Theory Comput. **15**, 6, 3537–3551 (2019).
- [82] T. Ichibha, K. Hongo, R. Maezono, et al. *Making the most of data: Quantum Monte Carlo Post-Analysis Revisited*. arXiv[physics.comp-ph, 1904.09934] (2019).
- [83] D. Cleland, G. H. Booth, C. Overy, et al. *Taming the First-Row Diatomics: A Full Configuration Interaction Quantum Monte Carlo Study*. J. Chem. Theory Comput. **8**, 11, 4138–4152 (2012).
- [84] G. H. Booth, D. Cleland, A. J. W. Thom, et al. *Breaking the carbon dimer: The challenges of multiple bond dissociation with full configuration interaction quantum Monte Carlo methods*. J. Chem. Phys. **135**, 8, 084104 (2011).
- [85] F. R. Petruzielo, A. A. Holmes, H. J. Changlani, et al. *Semistochastic Projector Monte Carlo Method*. Phys. Rev. Lett. **109**, 23, 230201 (2012).
- [86] C. Overy, G. H. Booth, N. S. Blunt, et al. *Unbiased reduced density matrices and electronic properties from full configuration interaction quantum Monte Carlo*. J. Chem. Phys. **141**, 24, 244117 (2014).
- [87] N. S. Blunt, S. D. Smart, J. A. F. Kersten, et al. *Semi-stochastic full configuration interaction quantum Monte Carlo: Developments and application*. J. Chem. Phys. **142**, 18, 184107 (2015).

- [88] H. Luo and A. Alavi. *Combining the Transcorrelated Method with Full Configuration Interaction Quantum Monte Carlo: Application to the Homogeneous Electron Gas*. J. Chem. Theory Comput. **14**, 3, 1403–1411 (2018).
- [89] L. R. Schwarz, A. Alavi, and G. H. Booth. *Projector Quantum Monte Carlo Method for Nonlinear Wave Functions*. Phys. Rev. Lett. **118**, 17, 176403 (2017).
- [90] N. S. Blunt. *Communication: An efficient and accurate perturbative correction to initiator full configuration interaction quantum Monte Carlo*. J. Chem. Phys. **148**, 22, 221101 (2018).
- [91] M.-A. Filip, C. J. C. Scott, and A. J. W. Thom. *Multireference Stochastic Coupled Cluster*. J. Chem. Theory Comput. **15**, 12, 6625–6635 (2019).
- [92] C. J. C. Scott and A. J. W. Thom. *Stochastic coupled cluster theory: Efficient sampling of the coupled cluster expansion*. J. Chem. Phys. **147**, 12, 124105 (2017).
- [93] Q. Sun, T. C. Berkelbach, N. S. Blunt, et al. *PySCF: the Python-based simulations of chemistry framework*. Wiley Interdiscip. Rev. Comput. Mol. Sci. **8**, 1, e1340 (2018).
- [94] J. M. Turney, A. C. Simmonett, R. M. Parrish, et al. *Psi4: an open-source ab initio electronic structure program*. Wiley Interdiscip. Rev. Comput. Mol. Sci. **2**, 4, 556–565 (2012).
- [95] R. M. Parrish, L. A. Burns, D. G. A. Smith, et al. *Psi4 1.1: An Open-Source Electronic Structure Program Emphasizing Automation, Advanced Libraries, and Interoperability*. J. Chem. Theory Comput. **13**, 7, 3185–3197 (2017).
- [96] W. A. Vigor, J. S. Spencer, M. J. Bearpark, et al. *Understanding and improving the efficiency of full configuration interaction quantum Monte Carlo*. J. Chem. Phys. **144**, 9, 094110 (2016).
- [97] W. A. Vigor, J. S. Spencer, M. J. Bearpark, et al. *Minimising biases in full configuration interaction quantum Monte Carlo*. J. Chem. Phys. **142**, 10, 104101 (2015).
- [98] C. J. Umrigar, M. P. Nightingale, and K. J. Runge. *A diffusion Monte Carlo algorithm with very small time-step errors*. J. Chem. Phys. **99**, 4, 2865–2890 (1993).
- [99] J. S. Spencer, N. S. Blunt, W. A. Vigor, et al. *Open-Source Development Experiences in Scientific Software: The HANDE Quantum Monte Carlo Project*. J. Open Res. Softw. **3**, 1, 1–6 (2015).
- [100] D. Bohm and D. Pines. *A Collective Description of Electron Interactions. I. Magnetic Interactions*. Phys. Rev. **82**, 5, 625–634 (1951).
- [101] D. Pines and D. Bohm. *A Collective Description of Electron Interactions: II. Collective vs Individual Particle Aspects of the Interactions*. Phys. Rev. **85**, 2, 338–353 (1952).

- [102] D. Bohm and D. Pines. *A Collective Description of Electron Interactions: III. Coulomb Interactions in a Degenerate Electron Gas*. Phys. Rev. **92**, 3, 609–625 (1953).
- [103] D. L. Freeman. *Coupled-cluster expansion applied to the electron gas: Inclusion of ring and exchange effects*. Phys. Rev. B **15**, 12, 5512–5521 (1977).
- [104] G. E. Scuseria, T. M. Henderson, and D. C. Sorensen. *The ground state correlation energy of the random phase approximation from a ring coupled cluster doubles approach*. J. Chem. Phys. **129**, 23, 231101 (2008).
- [105] J. J. Shepherd and A. Grüneis. *Correlation Energy Divergences in Metallic Systems*. arXiv[physics.chem-ph, 1208.6103] (2012).
- [106] J. J. Shepherd and A. Grüneis. *Many-body quantum chemistry for the electron gas: Convergent perturbative theories*. Phys. Rev. Lett. **110**, 22, 226401 (2013).
- [107] J. J. Shepherd, A. Grüneis, G. H. Booth, et al. *Convergence of many-body wave-function expansions using a plane-wave basis: From homogeneous electron gas to solid state systems*. Phys. Rev. B **86**, 3, 035111 (2012).
- [108] 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**, 24, 244101 (2012).
- [109] D. M. Ceperley and B. J. Alder. *Ground State of the Electron Gas by a Stochastic Method*. Phys. Rev. Lett. **45**, 7, 566–569 (1980).
- [110] G. Ortiz and P. Ballone. *Correlation energy, structure factor, radial distribution function, and momentum distribution of the spin-polarized uniform electron gas*. Phys. Rev. B **50**, 3, 1391–1405 (1994).
- [111] G. Ortiz and P. Ballone. *Erratum: Correlation energy, structure factor, radial distribution function, and momentum distribution of the spin-polarized uniform electron gas [Phys. Rev. B 50 , 1391 (1994)]*. Phys. Rev. B **56**, 15, 9970–9970 (1997).
- [112] G. Ortiz, M. Harris, and P. Ballone. *Zero Temperature Phases of the Electron Gas*. Phys. Rev. Lett. **82**, 26, 5317–5320 (1999).
- [113] Y. Kwon, D. M. Ceperley, and R. M. Martin. *Effects of backflow correlation in the three-dimensional electron gas: Quantum Monte Carlo study*. Phys. Rev. B **58**, 11, 6800–6806 (1998).
- [114] M. Holzmann, D. M. Ceperley, C. Pierleoni, et al. *Backflow correlations for the electron gas and metallic hydrogen*. Phys. Rev. E **68**, 4, 046707 (2003).
- [115] P. López Ríos, A. Ma, N. D. Drummond, et al. *Inhomogeneous backflow transformations in quantum Monte Carlo calculations*. Phys. Rev. E - Stat. Nonlinear, Soft Matter Phys. **74**, 6, 066701 (2006).

- [116] N. D. Drummond, R. J. Needs, A. Sorouri, et al. *Finite-size errors in continuum quantum Monte Carlo calculations*. Phys. Rev. B - Condens. Matter Mater. Phys. **78**, 12, 125106 (2008).
- [117] 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**, 8, 085121 (2013).
- [118] R. F. Bishop and K. H. Lührmann. *Electron correlations: I. Ground-state results in the high-density regime*. Phys. Rev. B **17**, 10, 3757–3780 (1978).
- [119] R. F. Bishop and K. H. Lührmann. *Electron correlations. II. Ground-state results at low and metallic densities*. Phys. Rev. B **26**, 10, 5523–5557 (1982).
- [120] A. Roggero, A. Mukherjee, and F. Pederiva. *Quantum Monte Carlo with coupled-cluster wave functions*. Phys. Rev. B **88**, 11, 115138 (2013).
- [121] J. McClain, J. Lischner, T. Watson, et al. *Spectral functions of the uniform electron gas via coupled-cluster theory and comparison to the GW and related approximations*. Phys. Rev. B **93**, 23, 235139 (2016).
- [122] J. J. Shepherd. *Communication: Convergence of many-body wave-function expansions using a plane-wave basis in the thermodynamic limit*. J. Chem. Phys. **145**, 3, 031104 (2016).
- [123] A. L. L. East and W. D. Allen. *The heat of formation of NCO*. J. Chem. Phys. **99**, 6, 4638–4650 (1993).
- [124] J. J. Shepherd, T. M. Henderson, and G. E. Scuseria. *Coupled cluster channels in the homogeneous electron gas*. J. Chem. Phys. **140**, 12, 124102 (2014).
- [125] J. J. Shepherd, T. M. Henderson, and G. E. Scuseria. *Range-separated Brueckner coupled cluster doubles theory*. Phys. Rev. Lett. **112**, 13, 133002 (2014).
- [126] G. McPherson. *Statistics in Scientific Investigation*. Springer Texts in Statistics. Springer New York, New York, NY (1990).
- [127] J. D. Hunter. *Matplotlib: A 2D Graphics Environment*. Comput. Sci. Eng. **9**, 3, 90–95 (2007).
- [128] L. K. Wagner and D. M. Ceperley. *Discovering correlated fermions using quantum Monte Carlo*. Reports Prog. Phys. **79**, 9, 094501 (2016).
- [129] H. Luo and A. Alavi. *Combining the Transcorrelated Method with Full Configuration Interaction Quantum Monte Carlo: Application to the Homogeneous Electron Gas*. J. Chem. Theory Comput. **14**, 3, 1403–1411 (2018).
- [130] N. S. Blunt, T. W. Rogers, J. S. Spencer, et al. *Density-matrix quantum Monte Carlo method*. Phys. Rev. B **89**, 24, 245124 (2014).
- [131] F. D. Malone, N. S. Blunt, J. J. Shepherd, et al. *Interaction picture density matrix quantum Monte Carlo*. J. Chem. Phys. **143**, 4, 044116 (2015).

- [132] F. D. Malone, N. S. Blunt, E. W. Brown, et al. *Accurate Exchange-Correlation Energies for the Warm Dense Electron Gas*. Phys. Rev. Lett. **117**, 11, 115701 (2016).
- [133] J. Lee, F. D. Malone, and M. A. Morales. *An auxiliary-Field quantum Monte Carlo perspective on the ground state of the dense uniform electron gas: An investigation with Hartree-Fock trial wavefunctions*. J. Chem. Phys. **151**, 6, 064122 (2019).
- [134] S. Zhang, J. Carlson, and J. E. Gubernatis. *Constrained path Monte Carlo method for fermion ground states*. Phys. Rev. B **55**, 12, 7464–7477 (1997).
- [135] S. Zhang and H. Krakauer. *Quantum Monte Carlo Method using Phase-Free Random Walks with Slater Determinants*. Phys. Rev. Lett. **90**, 13, 136401 (2003).
- [136] J. D. Power and R. M. Pitzer. *Inequalities For Electron Repulsion Integrals*. Chem. Phys. Lett. **24**, 4, 478–483 (1974).
- [137] J. M. Foster and S. F. Boys. *Canonical Configurational Interaction Procedure*. Rev. Mod. Phys. **32**, 2, 300–302 (1960).
- [138] A. J. W. Thom and A. Alavi. *A combinatorial approach to the electron correlation problem*. J. Chem. Phys. **123**, 20, 204106 (2005).
- [139] A. J. W. Thom and A. Alavi. *Stochastic Perturbation Theory: A Low-Scaling Approach to Correlated Electronic Energies*. Phys. Rev. Lett. **99**, 14, 143001 (2007).
- [140] M. Kolodrubetz and B. K. Clark. *Partial node configuration-interaction Monte Carlo as applied to the Fermi polaron*. Phys. Rev. B **86**, 7, 075109 (2012).
- [141] F. Pederiva, A. Roggero, and K. E. Schmidt. *Variational and Diffusion Monte Carlo Approaches to the Nuclear Few- and Many-Body Problem*. In M. Hjorth-Jensen, M. P. Lombardo, and U. van Kolck (Editors), *An Adv. Course Comput. Nucl. Phys. Bridg. Scales from Quarks to Neutron Stars*, pages 401–476. Springer International Publishing, Cham (2017).
- [142] A. J. Walker. *New fast method for generating discrete random numbers with arbitrary frequency distributions*. Electron. Lett. **10**, 8, 127 (1974).
- [143] A. J. Walker. *An Efficient Method for Generating Discrete Random Variables with General Distributions*. ACM Trans. Math. Softw. **3**, 3, 253–256 (1977).
- [144] R. A. Kronmal and A. V. Peterson. *On the Alias Method for Generating Random Variables from a Discrete Distribution*. Am. Stat. **33**, 4, 214–218 (1979).
- [145] D. E. Knuth. *3.4.1. Other Types of Random Quantities - Numerical Distributions*. Addison-Wesley, Reading, Mass. ; Harlow, 3rd ed. edition.
- [146] N. S. Blunt, G. H. Booth, and A. Alavi. *Density matrices in full configuration interaction quantum Monte Carlo: Excited states, transition dipole moments, and parallel distribution*. J. Chem. Phys. **146**, 24, 244105 (2017).

- [147] A. A. Holmes. *Advances in deterministic, stochastic, and semistochastic quantum chemistry*. Ph.D. thesis, Cornell University (2017).
- [148] L. R. Schwarz. *Projector Quantum Monte Carlo Methods for Linear and Non-linear Wavefunction Ansatzes*. Ph.D. thesis, University of Cambridge (2017).
- [149] C. C. J. Roothaan. *New Developments in Molecular Orbital Theory*. Rev. Mod. Phys. **23**, 2, 69–89 (1951).
- [150] E. R. Davidson. *The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real-symmetric matrices*. J. Comput. Phys. **17**, 1, 87–94 (1975).
- [151] T. Zhang and F. A. Evangelista. *A Deterministic Projector Configuration Interaction Approach for the Ground State of Quantum Many-Body Systems*. J. Chem. Theory Comput. **12**, 9, 4326–4337 (2016).
- [152] I. Sabzevari and S. Sharma. *Improved Speed and Scaling in Orbital Space Variational Monte Carlo*. J. Chem. Theory Comput. **14**, 12, 6276–6286 (2018).
- [153] L. Otis and E. Neuscamman. *Complementary first and second derivative methods for ansatz optimization in variational Monte Carlo*. Phys. Chem. Chem. Phys. **21**, 27, 14491–14510 (2019).
- [154] J. E. Deustua, J. Shen, and P. Piecuch. *Converging High-Level Coupled-Cluster Energetics by Monte Carlo Sampling and Moment Expansions*. Phys. Rev. Lett. **119**, 22, 223003 (2017).
- [155] J. E. Deustua, I. Magoulas, J. Shen, et al. *Communication: Approaching exact quantum chemistry by cluster analysis of full configuration interaction quantum Monte Carlo wave functions*. J. Chem. Phys. **149**, 15, 151101 (2018).
- [156] E. R. Davidson and W. J. Thompson. *Monster Matrices: Their Eigenvalues and Eigenvectors*. Comput. Phys. **7**, 5, 519 (1993).
- [157] T. Schoof, S. Groth, J. Vorberger, et al. *Ab Initio Thermodynamic Results for the Degenerate Electron Gas at Finite Temperature*. Phys. Rev. Lett. **115**, 13, 130402 (2015).
- [158] L. M. Fraser, W. M. C. Foulkes, G. Rajagopal, et al. *Finite-size effects and Coulomb interactions in quantum Monte Carlo calculations for homogeneous systems with periodic boundary conditions*. Phys. Rev. B **53**, 4, 1814–1832 (1996).
- [159] Y. Kurashige and T. Yanai. *High-performance ab initio density matrix renormalization group method: Applicability to large-scale multireference problems for metal compounds*. J. Chem. Phys. **130**, 23, 234114 (2009).
- [160] S. Sharma and G. K.-L. Chan. *Spin-adapted density matrix renormalization group algorithms for quantum chemistry*. J. Chem. Phys. **136**, 12, 124121 (2012).
- [161] R. Olivares-Amaya, W. Hu, N. Nakatani, et al. *The ab-initio density matrix renormalization group in practice*. J. Chem. Phys. **142**, 3, 034102 (2015).

- [162] S. Goedecker, M. Teter, and J. Hutter. *Separable dual-space Gaussian pseudopotentials*. Phys. Rev. B **54**, 3, 1703–1710 (1996).
- [163] C. Hartwigsen, S. Goedecker, and J. Hutter. *Relativistic separable dual-space Gaussian pseudopotentials from H to Rn*. Phys. Rev. B **58**, 7, 3641–3662 (1998).
- [164] J. VandeVondele, M. Krack, F. Mohamed, et al. *Quickstep: Fast and accurate density functional calculations using a mixed Gaussian and plane waves approach*. Comput. Phys. Commun. **167**, 2, 103–128 (2005).
- [165] W. Dobrazutz, H. Luo, and A. Alavi. *Compact numerical solutions to the two-dimensional repulsive Hubbard model obtained via nonunitary similarity transformations*. Phys. Rev. B **99**, 7, 075119 (2019).
- [166] A. J. Cohen, H. Luo, K. Guthrie, et al. *Similarity transformation of the electronic Schrödinger equation via Jastrow factorization*. J. Chem. Phys. **151**, 6, 061101 (2019).
- [167] B. Huron, J. P. Malrieu, and P. Rancurel. *Iterative perturbation calculations of ground and excited state energies from multiconfigurational zeroth-order wavefunctions*. J. Chem. Phys. **58**, 12, 5745–5759 (1973).
- [168] R. J. Buenker and S. D. Peyerimhoff. *Individualized configuration selection in CI calculations with subsequent energy extrapolation*. Theor. Chim. Acta **35**, 1, 33–58 (1974).
- [169] A. A. Holmes, N. M. Tubman, and C. J. Umrigar. *Heat-Bath Configuration Interaction: An Efficient Selected Configuration Interaction Algorithm Inspired by Heat-Bath Sampling*. J. Chem. Theory Comput. **12**, 8, 3674–3680 (2016).
- [170] N. S. Blunt. *A hybrid approach to extending selected configuration interaction and full configuration interaction quantum Monte Carlo*. J. Chem. Phys. **151**, 17, 174103 (2019).
- [171] C. J. C. Scott, R. Di Remigio, T. D. Crawford, et al. *Diagrammatic Coupled Cluster Monte Carlo*. J. Phys. Chem. Lett. **10**, 5, 925–935 (2019).
- [172] A. Gothandaraman, G. D. Peterson, G. Warren, et al. *FPGA acceleration of a quantum Monte Carlo application*. Parallel Comput. **34**, 4-5, 278–291 (2008).
- [173] S. Cardamone, J. R. R. Kimmitt, H. G. A. Burton, et al. *Field-programmable gate arrays and quantum Monte Carlo: Power efficient coprocessing for scalable high-performance computing*. Int. J. Quantum Chem. **119**, 12, e25853 (2019).
- [174] N. Luehr, I. S. Ufimtsev, and T. J. Martínez. *Dynamic Precision for Electron Repulsion Integral Evaluation on Graphical Processing Units (GPUs)*. J. Chem. Theory Comput. **7**, 4, 949–954 (2011).
- [175] A. E. DePrince and J. R. Hammond. *Coupled Cluster Theory on Graphics Processing Units I. The Coupled Cluster Doubles Method*. J. Chem. Theory Comput. **7**, 5, 1287–1295 (2011).

- [176] A. P. Bartók, M. C. Payne, R. Kondor, et al. *Gaussian Approximation Potentials: The Accuracy of Quantum Mechanics, without the Electrons*. Phys. Rev. Lett. **104**, 13, 136403 (2010).
- [177] M. Rupp, A. Tkatchenko, K.-R. Müller, et al. *Fast and Accurate Modeling of Molecular Atomization Energies with Machine Learning*. Phys. Rev. Lett. **108**, 5, 058301 (2012).
- [178] G. B. Goh, N. O. Hodas, and A. Vishnu. *Deep learning for computational chemistry*. J. Comput. Chem. **38**, 16, 1291–1307 (2017).
- [179] J. T. Margraf and K. Reuter. *Making the Coupled Cluster Correlation Energy Machine-Learnable*. J. Phys. Chem. A **122**, 30, 6343–6348 (2018).
- [180] J. P. Coe. *Machine Learning Configuration Interaction*. J. Chem. Theory Comput. **14**, 11, 5739–5749 (2018).
- [181] K. T. Schütt, M. Gastegger, A. Tkatchenko, et al. *Unifying machine learning and quantum chemistry with a deep neural network for molecular wavefunctions*. Nat. Commun. **10**, 1, 5024 (2019).
- [182] J. S. Smith, B. T. Nebgen, R. Zubatyuk, et al. *Approaching coupled cluster accuracy with a general-purpose neural network potential through transfer learning*. Nat. Commun. **10**, 1, 2903 (2019).
- [183] D. Pfau, J. S. Spencer, A. G. d. G. Matthews, et al. *Ab-Initio Solution of the Many-Electron Schrödinger Equation with Deep Neural Networks*. arXiv [physics.chem-ph, 1909.02487] (2019).

